

IBM[®] DB2 Universal Database[™]



Guide to GUI Tools for Administration and Development

Version 8

IBM[®] DB2 Universal Database[™]



Guide to GUI Tools for Administration and Development

Version 8

Before using this information and the product it supports, be sure to read the general information under *Notices*.

This document contains proprietary information of IBM. It is provided under a license agreement and is protected by copyright law. The information contained in this publication does not include any product warranties, and any statements provided in this manual should not be interpreted as such.

You can order IBM publications online or through your local IBM representative.

- To order publications online, go to the IBM Publications Center at www.ibm.com/shop/publications/order
- To find your local IBM representative, go to the IBM Directory of Worldwide Contacts at www.ibm.com/planetwide

To order DB2 publications from DB2 Marketing and Sales in the United States or Canada, call 1-800-IBM-4YOU (426-4968).

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© **Copyright International Business Machines Corporation 2002. All rights reserved.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

About this book.	v	Task Center	37
Who should use this book	v	Information Catalog Center	39
How this book is structured.	vi	Health Center	40
Before you begin working with the GUI.	vi	Journal.	42
Chapter 1. DB2 concepts	1	License Center	44
Database object hierarchy	1	Development Center	45
Systems	2	Information Center.	46
Instances	3	Memory Visualizer.	48
Databases	4	Spatial Extender	50
Database partition groups.	5	SQL Assist	50
Table spaces	5	Visual Explain	52
Long data	6	Access plan	53
DB2 objects	7	Access plan graph	53
Database manager	7	Configuration Assistant	55
Catalogs.	7	Indoubt Transaction Manager	56
Packages	8	Web tools	57
Aliases	8	Web Command Center	57
Authorities.	9	Web Health Center.	58
Performance elements.	10	Chapter 4. Using the tools	61
Buffer pools	10	Backup and restore.	61
Configuration parameters	11	Backup.	61
Data elements	12	Restore.	62
Registry variables	12	Backing up a database online for the first	
Distributed databases	13	time.	63
Replication	13	Restoring a database	64
Federated systems	14	Database health.	66
Chapter 2. GUI navigation.	15	Reacting to alerts	66
Basic GUI navigation concepts explained	15	Replication	67
Chapter 3. DB2 GUI tools	21	Replicating data.	67
Control Center	22	Appendix A. DB2 Universal Database	
Tool invocation	23	technical information	73
Storage Management view	24	Overview of DB2 Universal Database	
Which objects you can administer.	25	technical information	73
Remote object administration in the Control		FixPaks for DB2 documentation	73
Center	26	Categories of DB2 technical information.	73
Federated systems in the Control Center	26	Printing DB2 books from PDF files	81
z/OS and OS/390 objects in the Control		Ordering printed DB2 books	82
Center	28	Accessing online help.	82
Replication Center	30	Finding topics by accessing the DB2	
Satellite Administration Center.	33	Information Center from a browser	84
Data Warehouse Center	34	Finding product information by accessing the	
Command Center	35	DB2 Information Center from the	
		administration tools	86

Viewing technical documentation online directly from the DB2 HTML Documentation CD	87	Compatibility with Assistive Technologies	94
Updating the HTML documentation installed on your machine	88	Accessible Documentation	94
Copying files from the DB2 HTML Documentation CD to a Web server	90	DB2 tutorials.	95
Troubleshooting DB2 documentation search with Netscape 4.x	91	DB2 Information Center accessed from a browser	96
Searching the DB2 documentation.	92	Appendix B. Notices.	99
Online DB2 troubleshooting information	93	Trademarks	102
Accessibility	93	Index	105
Keyboard Input and Navigation	94	Contacting IBM	107
Accessible Display	94	Product information	107
Alternative Alert Cues	94		

About this book

As you move into the Internet age of broad-band computing, you need tools that can manage ever increasing volumes of data, handle the tremendous growth in the number of users, deliver better-than-ever performance, and support the next generation of applications. You need to do more. You need to do it more safely. And you need to do it more easily. DB2 Universal Database™ provides a suite of graphical user interface (GUI) administration and development tools to help you accomplish your goals. The DB2® GUI tools support tasks like database management, database health monitoring, and task scheduling allowing you to view and work with database objects in a hierarchical tree, optimize queries, and organize task flow. Reading this guide offers you the opportunity to learn about data storage, data access, and database administration for optimum performance using the DB2 GUI tools.

For information on the command-driven DB2 interface, see the *Command Reference* and the *SQL Reference, Volume 1*.

Who should use this book

This book is for database users who might be familiar with other database products but are new to DB2. It is for experienced DB2 administrators who want information on the new GUI features and enhancements available with DB2 Universal Database, Version 8. It is also intended for command line users, introducing them to the powerful set of GUI tools in DB2 and, perhaps, inducing them to try one or two of them.

For new users, a graphical user interface (GUI) is a program interface that takes advantage of the computer's graphics capabilities to make the program easier to use. A well-designed GUI can free the user from learning complex command languages. DB2 provides more than fifteen GUI tools to help you get your job done easily and efficiently.

Experienced DB2 users will find new and enhanced tools for administering their DB2 environments. For example, the new Health Center identifies key performance and resource allocation problems and provides recommended actions that can help resolve the problems.

Users who find that they work more effectively with a command line interface are encouraged to give the GUI a try. They can derive some comfort from knowing that the GUI can show the SQL or commands that the tools generate.

Note: The contents of the sample dialogs pictured in this book might not exactly match your dialogs as content is environment-specific. For example, system names are related to the name of the computer where DB2 is installed.

How this book is structured

This book describes how to use the DB2 GUI tools. It also explains why you may want to use a specific tool or combination of tools. It focuses on giving you an understanding of some key strategies for using the GUI tools to perform real-world tasks effectively. It does not present these strategies as a step-by-step guided tutorial, but, rather it encourages you to explore the GUI as an active learner. As you try things, you can use the comprehensive online help for assistance.

Chapter 1, “DB2 concepts” on page 1 describes specific DB2 concepts and explains database terms that have slightly different meanings for users of DB2. It does not provide definitions for common database terms. It will help new users of DB2 become familiar with the DB2 vocabulary.

Chapter 2, “GUI navigation” on page 15 contains information on basic DB2 GUI navigation. It explains how to work with menus, views, and filters. It also describes the differences between dialogs, wizards, and launchpads.

Chapter 3, “DB2 GUI tools” on page 21 presents an overview of the DB2 administration and development tools that are available to you, and it explains how you can use them to get your job done.

Chapter 4, “Using the tools” on page 61 provides some scenario-based descriptions of effective use of the GUI tools for your day-to-day tasks.

Before you begin working with the GUI

For the administrative tools such as the Control Center to work properly, you must have a screen resolution of at least 800 by 600 pixels, and be using a display palette of at least 32 colors.

In order to run the GUI tools, you should have 256MB of available memory.

On AIX 5L machines, AIX APAR IY29345 is required in order to work with the GUI tools.

Chapter 1. DB2 concepts

DB2 Universal Database is a relational database management system (RDBMS) that lets you create, update, and administer relational databases using Structured Query Language (SQL) statements entered by a user or contained in an application program. In these respects, it resembles most other RDBMS products available on the market today. However, in much the same way that the English spoken in Australia differs from that spoken in the United States, DB2 has a slightly different dialect than that of other RDBMS products.

This chapter describes concepts specific to DB2 and explains database terms that have slightly different meanings for users of DB2. It also identifies the GUI tools that can be used when working with an object and the associated tasks that can be performed. It does not provide definitions for common database terms.

Database object hierarchy

Figure 1 on page 2 illustrates the DB2 hierarchy of database objects.

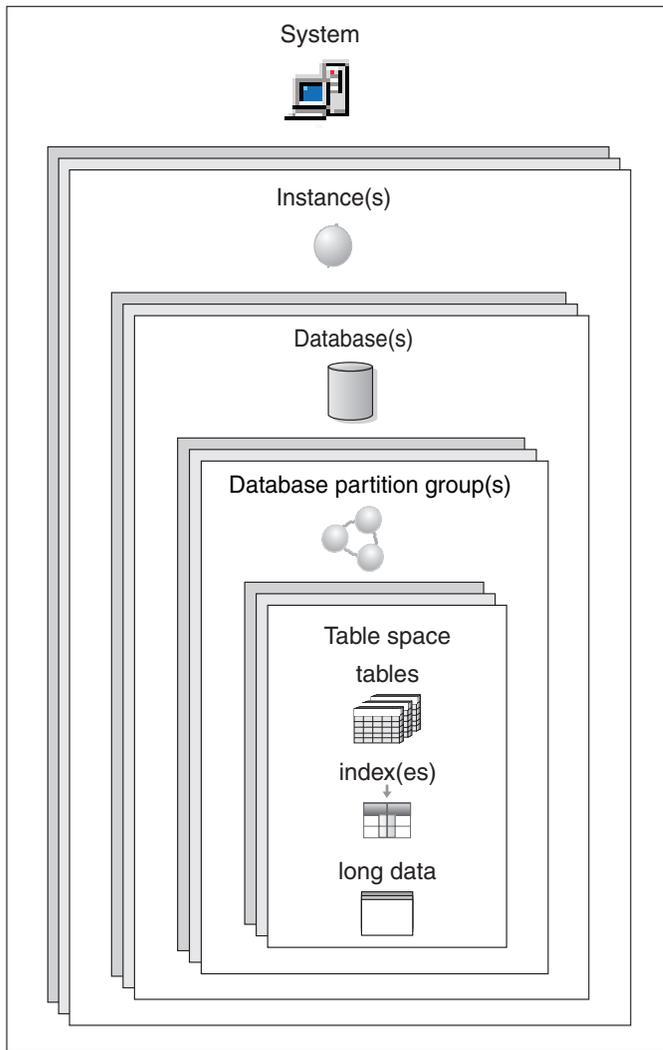


Figure 1. Hierarchy of database objects

Systems

The highest-level object in the DB2 hierarchy is a system. A system represents an installation of DB2. A system can have one or more DB2 instances, each of which can manage one or more databases. The databases may be partitioned with their table spaces residing in database partition groups. The table spaces, in turn, store table data.

A list of systems is maintained by the Control Center: it records the information needed to communicate with each system (such as its network address, operating system, and communication protocol).

You can:

- Add a system using the Configuration Assistant
- Add a system to the Control Center
- Remove a system using the Configuration Assistant
- Remove a system from the Control Center

Related concepts:

- “Configuration Assistant” on page 55
- “Control Center” on page 22
- “Instances” on page 3
- “Database partition groups” on page 5
- “Table spaces” on page 5

Instances

An instance is a logical database manager environment where you catalog databases and set configuration parameters. You can create multiple instances on the same physical server providing a unique database server environment for each instance. With multiple instances, you can:

- Use one instance for a development environment and another instance for a production environment.
- Tune an instance for a particular environment.
- Restrict access to sensitive information.
- Control the assignment of SYSADM, SYSCTRL, and SYSMANT authority for each instance.
- Optimize the database manager configuration for each instance.
- Limit the impact of an instance failure. In the event of an instance failure, only that instance is affected. Other instances can continue to function normally.

Multiple instances require:

- Additional system resources (virtual memory and disk space) for each instance
- More administration

On Windows[®] operating systems, ensure that no instance name is the same as a service name.

You can:

- Add an instance using the Control Center or Configuration Assistant
- Remove an instance using the Control Center or Configuration Assistant
- Configure an instance using the Control Center or Configuration Assistant

Related concepts:

- “Configuration Assistant” on page 55
- “Control Center” on page 22
- “Systems” on page 2
- “Database manager” on page 7
- “Configuration parameters” on page 11

Databases

A relational database presents data as a collection of tables. A table consists of a defined set of columns and any number of rows. The data in each table is logically related, and relationships can be defined between tables. Data can be viewed and manipulated based on mathematical principles and operations called relations (such as, INSERT, SELECT, and UPDATE).

A database is self-describing in that it contains, in addition to data, a description of its own structure. It includes a set of system catalog tables, which describe the logical and physical structure of the data; a configuration file, which contains the parameter values associated with the database; and a recovery log, which records ongoing transactions and transactions that can be archived.

Databases can be local or remote. A local database is physically located on the workstation in use, while a database on another machine is considered remote.

You can:

- Create a database using the Control Center
- Add a database using the Control Center or Configuration Assistant
- Drop a database from the Control Center
- Backup a database using the Control Center
- Restore a database using the Control Center
- Configure a database using the Control Center
- Connect to a database using the Control Center
- Monitor a database with the event monitor.
- Monitor database health with the Health Center

Related concepts:

- “Configuration Assistant” on page 55
- “Control Center” on page 22
- “Health Center” on page 40
- “Configuration parameters” on page 11
- “Catalogs” on page 7

Database partition groups

In a single-partition database environment, a database partition group exists but it is transparent. Database partition groups are much more important in multi-partition database environments.

A database partition group (called a nodegroup in previous versions) is a named set of database partitions that belong to a database. A database can consist of a single database partition on one server or multiple database partitions spread over a number of servers. Each database partition has its own user data, indexes, configuration files, and transaction logs.

A database partition group can contain as few as one database partition or span all the database partitions that make up the database. For convenience, a default database partition group comprising all database partitions is created in every new partitioned database.

When a database partition group is created or modified, a partitioning map is associated with it. A partitioning map, in conjunction with a partitioning key and a hashing algorithm, is used by the database manager to determine which database partition in the database partition group will store a given row of data.

You can use the Control Center to:

- Create a database partition group
- Add database partitions to a database partition group
- Drop database partitions from a database partition group
- Redistribute table data within a database partition group

Related concepts:

- “Control Center” on page 22
- “Database manager” on page 7
- “Table spaces” on page 5
- “Databases” on page 4

Table spaces

The physical space within a database is organized into a collection of table spaces. Each table space consists of a collection of containers, each of which is an allocation of physical storage (for example, a directory on a machine, a physical file, or a device such as a hard drive). The system attempts to spread the data across the containers of its table spaces.

Each table is assigned to a table space. For improved performance, or for purposes of having table space backups, you can store the indexes for a table

in a second table space and the large objects for the table in a third table space. More than one table can be assigned to the same table space. In a partitioned database environment, table spaces reside in database partition groups.

Assigning tables to table spaces and mapping table spaces to physical storage gives you some control over the performance of your database. For example, you can use your fastest storage devices for your most frequently used tables and store less frequently used data on slower devices.

There are two types of table spaces:

- System managed space (SMS) table spaces in which the operating system file system manager allocates and manages the space where a table is stored. The user decides on the location of the files, DB2[®] controls their names, and the file system is responsible for managing them. Container size is dynamic and determined by the available space on the device.
- Database managed space (DMS) table spaces in which the database manager controls the storage space. This storage model consists of a limited number of devices whose space is managed by DB2. The database administrator decides which devices to use, and DB2 manages the space on those devices. The administrator also specifies the amount of space that can be allocated and DB2 will not exceed that limit.

You can use the Control Center to:

- Create a table space
- Drop a table space
- Alter the characteristics of a table space
- Grant and revoke privileges on a table space
- Show the objects related to a table space
- Monitor a table space with the event monitor
- Monitor table space health with the Health Center

Related concepts:

- “Control Center” on page 22
- “Health Center” on page 40
- “Database partition groups” on page 5

Long data

Long field data (data type LONG VARCHAR or LONG VARGRAPHIC) and large objects (data type BLOB, CLOB, or DBCLOB) are usually stored in a separate table object in a table space. These tables are structured differently than the storage space for other data types. Special restrictions apply to SQL expressions referencing these data types.

Related concepts:

- “Control Center” on page 22
- “Table spaces” on page 5

Related reference:

- “Character strings” in the *SQL Reference, Volume 1*
- “Graphic strings” in the *SQL Reference, Volume 1*

DB2 objects

Database manager

The database manager includes the database engine and the facilities to access data, such as the command line processor and the application interfaces.

The database engine manages data, controls all access to data, generates packages, generates optimized paths, provides transaction management, ensures data integrity and data security, and provides concurrency control.

A DB2[®] instance is a logical database manager environment where databases can be cataloged and configuration parameters set. Multiple instances of the database manager may be created on a single server and can be used for different purposes, for example as separate test and production environments.

DB2 program files are physically stored in one location on a particular computer. Each instance of the database manager that is created shares these files so the program files are not duplicated for each instance that is created.

You can use the Configuration Assistant and Control Center to view, update, and reset database manager configuration parameters.

Related concepts:

- “Configuration Assistant” on page 55
- “Control Center” on page 22
- “Instances” on page 3
- “Configuration parameters” on page 11
- “Packages” on page 8

Catalogs

Catalogs are special tables that contain information about all the objects within a database. This includes objects such as tables, views, and indexes, as well as security controls like constraints and database-specific authorities. These system catalog tables are created when the database is created.

When an object is created, altered, or dropped, DB2 inserts, updates, or deletes the rows of the catalog that describe the object and how that object relates to other objects. For example, when you use the Control Center to create a table DB2® adds rows to the system catalog tables.

Catalogs are primarily for read-only purposes because they are maintained by DB2. Their data is available through normal SQL query facilities. However there is a special set of catalog views that are updatable. This set of views is defined on the catalogs that are used to update database statistics. Manual update of these views provides a means through which to influence the system optimizer or to perform experiments on test databases.

You can:

- Use the Control Center to update the system catalog statistics on the data in a table and its indexes
- Use Visual Explain to examine the results of manual updates on optimization experiments

Related concepts:

- “Control Center” on page 22
- “Visual Explain” on page 52

Packages

A package is an object produced during program preparation that contains all of the compiled SQL statements and control structures used at run time in a single source file. A package is made up of sections. A section contains the compiled form of an SQL statement. Although every section corresponds to one statement, not every statement has a section. The sections created for static SQL are comparable to the bound, or operational, form of SQL statements. The sections created for dynamic SQL are comparable to placeholder control structures used at run time.

You can use the Control Center to grant and revoke privileges on a package.

Aliases

An alias is an alternative name for a database, table, view, or even another alias.

Database aliases

Database alias names are local synonyms given to local and remote databases.

You cannot create a database on a DB2 server if that database name already exists on the server. A DB2 client, however, can have connections to two different databases with the same name, but on different servers, if each has its own alias.

Alias names must be unique within the system database directory in which all aliases are stored for the individual instance of the database manager. When you create a new database, the alias defaults to the database name if you do not specify an alias.

You can create an alias for a database when adding a database using the Configuration Assistant or Control Center.

Table-related aliases

Table-related aliases are somewhat different from database aliases. Assigning an alias to a database can avoid potential client connection problems in environments where different servers might have databases with the same name. Using a table or view alias lets you refer to an object indirectly so that an SQL statement can be independent of the actual name of that object. Using an alias in an SQL statement is equivalent to using the object's name. The advantage of using an alias instead of the object's name is the ease with which you can change the object to which the alias applies. Programs that use aliases can be easily redirected from one object to another without any changes to their code.

You can:

- Create an alias for a table or a view using the Control Center
- Drop an alias for a table or a view using the Control Center
- Show alias relationships for tables and views using the Control Center

Related concepts:

- “Configuration Assistant” on page 55
- “Control Center” on page 22

Authorities

Authorities provide a way both to group privileges and to control maintenance and utility operations for instances, databases, and database objects. Users can have administrative authorities that give full privileges on a set of database objects, or they can have system authorities that give full privileges on managing the system but do not allow access to the data.

SYSADM is the highest level of administrative authority. It includes all the privileges on all databases within the DB2[®] instance as well as the authority to grant and revoke all the other authorities and privileges.

DBADM provides administrative authority for a specific database. It allows the user to access and modify all the objects within that database. A user with DBADM authority can grant and revoke privileges on the database but cannot grant or revoke DBADM authority.

SYSCTRL is the authority for controlling the resources used by the database manager (for example, creating and deleting databases), but it does not allow access to the data within the databases.

SYSMAINT is the authority for performing maintenance operations, such as starting and stopping the DB2 server and backing up and restoring databases. It does not allow access to the data within the databases.

LOAD authority at the database level, combined with INSERT privilege on a table, allows the user to load data into that table.

Database-specific authorities are stored in the database catalogs; system authorities are stored in the database manager configuration file for the instance.

You can use the Control Center to grant and revoke database authorities.

Related concepts:

- “Control Center” on page 22
- “Instances” on page 3

Performance elements

Performance is the way a computer system behaves during a particular workload. Performance is measured in terms of system response time, throughput, and availability. This section describes some of DB2's performance related components.

Buffer pools

A buffer pool is an area of memory into which database pages are read, modified, and held during processing.

Buffer pools improve database performance. If a needed page of data is already in the buffer pool, that page is accessed faster than if that page had to be read directly from disk. The database manager has agents whose tasks are to retrieve data pages from disk and place them in the buffer pool (prefetchers), and to write modified data pages from the buffer pool back to disk (page cleaners).

The reading and writing of data pages to and from disk is called disk input/output (I/O). Avoiding the wait associated with disk I/O is the primary way to improve the performance of the database. How you create the buffer pool and configure the database manager and the agents associated with the buffer pool controls the performance of the database. Using SQL and configuration parameters, you can control the size of the buffer pool, the number of prefetchers and page cleaners that move data pages into and out of the buffer pool, the size of the data pages, and the number of data pages that can be moved at one time.

A database must always have at least one buffer pool. DB2[®] provides a default buffer pool (IBMDEFAULTBP), which you can use as it is, or you can alter it.

You can use the Control Center to:

- Create a buffer pool
- Alter a buffer pool
- Drop a buffer pool, except the default buffer pool

Related concepts:

- “Control Center” on page 22
- “Database manager” on page 7
- “Configuration parameters” on page 11

Configuration parameters

Configuration parameters affect the operating characteristics of a database or database manager. They are stored in configuration files.

The database manager configuration file is created when a DB2[®] instance is created. The parameters it contains affect system resources at the instance level, independent of any one database that is part of that instance. Values for many of these parameters can be changed from the system default values to improve performance or increase capacity, depending on your system configuration. Database manager configuration parameters exist on servers and clients. The parameters available on a client are a subset of the database management configuration parameters that can be set on the server.

A database configuration file is created when a database is created, and it resides where that database resides. There is one configuration file per database, or in a partitioned database environment, one configuration file for each database partition. The file parameters specify, among other things, the amount of resource to be allocated to that database. Values for many of the parameters can be changed to improve performance or increase capacity. Different changes may be required, depending on the type of activity in a specific database.

You can:

- View, update, and reset database manager configuration parameters for a server using the Control Center, the Configuration Assistant, and the Command Center
- View, update, and reset database manager configuration parameters for a client using the Configuration Assistant and the Command Center
- View, update, and reset database configuration parameters using the Control Center

Related concepts:

- “Configuration Assistant” on page 55
- “Command Center” on page 35
- “Control Center” on page 22
- “Instances” on page 3
- “Database manager” on page 7

Data elements

Data elements are used to store information about the operation and performance of the database manager and that of the applications using it. This information can be useful for tuning performance and diagnosing problems. Snapshots can be taken of the current state of activity in the database manager for a particular object or group of objects. In addition, you can also gather information about transient events that are difficult to monitor through snapshots, such as deadlocks and transaction completions.

You can use the Control Center’s event monitor to work with data elements.

Related concepts:

- “Control Center” on page 22
- “Database manager” on page 7

Registry variables

Registry variables define the DB2 operating environment. These variables are stored in the DB2® profile registries. There are two levels of registry variables:

Global

Variable settings are system-wide.

Instance

Variable settings that are for a particular instance. Values defined for an instance override their settings at the global level.

There are four profile registries:

- The DB2 Instance Level Profile Registry. The majority of the DB2 variables are placed within this registry. The variable settings for a particular instance are kept in this registry. Values defined in this level override their settings in the global level.
- The DB2 Global Level Profile Registry. If a variable is not set for a particular instance, this registry is used. This registry has the machine-wide variable settings. In DB2 UDB Enterprise Server Edition (ESE), one global-level profile exists at each machine.
- The DB2 Instance Node Level Profile Registry. This registry level contains variable settings that are specific to a database partition server (node) in a partitioned database environment. Values defined in this level override their settings at the instance and global levels.
- The DB2 Instance Profile Registry. This registry contains a list of all instance names recognized by this system.

In a Lightweight Directory Access Protocol (LDAP) environment, the DB2 profile registry variables can be set at the user level which allows a user to customize their own DB2 environment.

You can use the Configuration Assistant to work with global and instance registry variables.

Related concepts:

- “Configuration Assistant” on page 55
- “Instances” on page 3

Distributed databases

This section describes some of capabilities of DB2 when working with distributed databases.

Replication

Replication is a process of maintaining a defined set of data in more than one location. It involves applying designated changes from one location (a source) to another (a target), and synchronizing the data in both locations. The source and target can be in logical servers (such as a DB2[®] database or a DB2 for z/OS[™] and OS/390[®] subsystem or a data-sharing group) that are on the same computer or on different computers in a distributed network.

The replication environment that you need depends on when you want data updated and how you want transactions handled.

You can use the Replication Center to administer replication environments.

Related concepts:

- “Replication Center” on page 30

Related tasks:

- “Replicating data” on page 67
- “Setting up for replication” in the *Replication Guide and Reference*

Federated systems

A DB2® federated system is a special type of distributed database management system. A federated system consists of a DB2 instance that operates as a server, a database that serves as the federated database, one or more data sources, and clients (users and applications) who access the federated database and data sources. A federated system allows you to send distributed requests to multiple data sources within a single SQL statement.

To end users and client applications, data sources appear as a single collective database in DB2. Users and applications interface with the federated database that is within the federated server.

You can:

- Use the Control Center to set up a federated system.
- Use the Command Center to configure data sources.

Related concepts:

- “Command Center” on page 35
- “Instances” on page 3
- “Federated systems in the Control Center” on page 26
- “Federated systems” in the *Federated Systems Guide*

Chapter 2. GUI navigation

The graphical user interface (GUI) tools enable you to perform your job easily; however, you need to understand how the GUI works and become familiar with its features. This chapter explains basic GUI navigation features.

Basic GUI navigation concepts explained

Here are some basic GUI concepts you should be familiar with.

Start Menu:

The *Start menu* is the starting point for using the GUI Administration tools in Windows® operating environments. From the **Start** menu, select **IBM DB2** —> **General Administration Tools** —> **Control Center**. You can also start other centers or tools such as the Journal, the Replication Center, and the Task Center from the Start menu as well as select the command line tools, the development tools, the monitoring tools, and the Information Center.

To start the GUI Administration tools from the Linux desktop environment, open the *IBM® DB2®* folder on your Gnome or KDE desktop and then select **Control Center**. On UNIX-based systems, enter the **db2cc** command from a command line.

Windows:

Windows are panels that open and enable you to enter information pertaining to the action you want to perform, for example, you can type information into fields within a window. In the DB2 interface, fields that must be filled in are surrounded by a thick red border.

Notebooks:

A *notebook* is used to logically organize data into groups when there is too much information to fit on one page. The resulting pages are organized with tabs that reflect the particular page content.

Wizards:

Wizards are integrated into the administration tools. They assist you in completing a single task by stepping you through the task. To select a wizard, from the Control Center window, select **Tools** —> **Wizards**. The **Wizards**

window opens. Select the wizard you want to use. Select the object for which you want help and follow the instructions to complete the task. The wizard task overview on the first page of the wizard lists any prerequisite steps and briefly describes every page of the wizard. Other pages of the wizard may contain links to conceptual or reference information to help you understand the function of the wizard. From a wizard, you can launch other wizards, windows or notebooks.

Advisors:

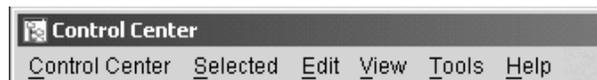
Advisors are integrated into the administration tools. They assist you with more complex tasks, such as tuning and performance tasks, by gathering information and recommending options that you may not have considered. You can accept or reject the advice of the advisor. Advisors can be called from the GUI as well as from APIs and the command line interface. To select an advisor, from the Control Center window, select **Tools** —> **Wizards**. The **Wizards** window opens. Select the advisor you want to use. Select the object for which you want help and follow the instructions to complete the task. Conceptual and reference information is available to help you understand the function of the advisor.

Launchpads:

Launchpads are integrated into the administration tools. They assist you in completing high-level tasks by stepping you through a set of tasks in the right order. Launchpads can call wizards or other dialogs to accomplish the high-level task. To select a launchpad, from the **Control Center** window, select **Tools** —> **Wizards**. The **Wizards** window opens. Select the launchpad you want to use. Select the object for which you want help and follow the instructions to complete the task. Pages of the launchpad may contain links to conceptual or reference information to help you understand the function of the launchpad.

Menu bars:

The *Menu bar* is a special panel that is displayed at the top of the window.



It contains menu options that allow you to access drop-down menus. From the drop-down menus, you can select menu items. Items in the menu bar include actions that affect the objects in the center you are using. For example, menu items under **Control Center** include actions that affect the entire Control Center. Menu items under **Selected** dynamically change to include only those actions that apply to currently selected objects in the contents pane.

Menu items under **Edit** include actions that let you work with objects in the contents pane. Menu items under **View** include actions that let you customize the display of objects in the contents pane. Menu items under **Help** include actions that display help information and allow you to work with the Information Center and Tutorials.

Toolbars:

Toolbars are panels that contain icons representing functions you can perform. Toolbars are located below the menu bar. To see a brief description of a tool, place your cursor over its icon and hover help will tell you what function each icon represents. Toolbars provide quick access to the functions you can perform. The functions can also be selected in the **View** menu.



A *Contents pane toolbar* is located below the contents pane. It allows you to tailor the information in the contents pane.



Object trees:

Object trees display the system and database objects graphically in the left navigation pane, enabling you to see the hierarchical relationship between different objects and to work with those objects. You can expand the object tree to expose the objects that are contained within the object. The exposed objects are displayed as folders beneath the object. Each folder represents an object type. If the object tree is collapsed, the folders contained within the object no longer appear in the view of the object tree.

Objects on your local workstation and any remote systems that can be connected to and from your local system are displayed in the object tree. Some objects, such as tables and views, do not contain other objects and appear at the bottom of the object tree.

When you select an object in the object tree, the objects that reside in the object are displayed in the contents pane. To invoke an action on an object, right-click the object to open a pop-up menu of available actions. You can also invoke an action on the object by selecting the object then clicking the **Selected** menu option from the menu bar.

Contents Pane (Details view):

The *Contents pane* is also called the *Details view*. It allows flexibility in representing large amounts of complex data in table format. From the Control Center, you can create and save customized views by selecting **View** from the Contents pane toolbar at the bottom of the Contents pane. From the **View** dropdown, you can filter, sort and customize columns and save these views as the default view or save the view using another name. This enables you to group the key elements and save them for future use. It also allows you to group rows of data with the same value in a specific column.

Infopops:

An *infopop* is a pop-up window that is displayed when a control in a window or notebook has focus and you press F1. Holding the mouse cursor over a control in a window or notebook also causes the infopop to display. Infopops contain a description of the field or control. They may also list restrictions and requirements or provide instructions. Infopops are disabled or re-enabled from the **General** page of the **Tools Settings** notebook. To disable or enable infopops for the Configuration Assistant, select the **Display Infopops** checkbox menu in the **Help** menu from the Configuration Assistant tool.

Mandatory fields:

Mandatory fields are fields for which you must supply information. Mandatory fields are displayed with thick red borders around the field. When information is entered in the mandatory field, the red border is no longer displayed.

Filtering:

Filtering enables you to work with a subset of displayed objects in the Control Center. Two forms of filtering exists. The first form of filtering allows you to create a customized view of objects which you would like to appear in the Contents pane of the Control Center. You select the subset of objects by right clicking the object folder in the object tree and selecting **Filter** —> **Create**. The Filter notebook opens allowing you to select which columns you would like to have in your customized view. Once you have created a filtered view of the objects in the Contents pane, the filter icon appears beneath the filtered object in the object tree to indicate that a filter has been set. To remove the filter and display all the objects in the Contents pane, right click the object folder and select **Filter** —> **Remove**.

Once you have filtered the objects in the Contents pane, you can customize your view even further by selecting the filter icon from the Contents pane

toolbar at the bottom of Contents pane.



The **Filter window** opens and displays available columns for filtering. You specify which columns you wish to view by supplying the information in the Filter notebook. To enable the filter, select the **Enable filter check box**.

Help:

Help is available from icons on the toolbar or by using the Help menu.



displays help for getting started with the Control Center.



opens the Information Center so you can search for help on tasks, commands, and information in the DB2 library.

The Help menu displays menu items for displaying the online help index, general information about the Control Center, and keyboard help. The menu also displays the list of tutorials available with DB2.

Chapter 3. DB2 GUI tools

This chapter describes the DB2 GUI tools. It explains what each tool is used for and why you might want to use it. It tells you how to invoke a tool and presents basic usage information.

You can select the following tools from the toolbar:

- “Control Center” on page 22
- “Replication Center” on page 30
- “Satellite Administration Center” on page 33
- “Data Warehouse Center” on page 34
- “Command Center” on page 35
- “Task Center” on page 37
- “Information Catalog Center” on page 39
- “Health Center” on page 40
- “Journal” on page 42
- “License Center” on page 44
- “Development Center” on page 45
- “Information Center” on page 46

You can invoke the following tools from within another tool:

- “Memory Visualizer” on page 48
- “Spatial Extender” on page 50
- “SQL Assist” on page 50
- “Visual Explain” on page 52

You can invoke the following tools by entering a command or (on a Windows system) selecting a menu choice:

- “Configuration Assistant” on page 55
- “Indoubt Transaction Manager” on page 56

You can invoke the following tools from a browser:

- “Web Command Center” on page 57
- “Web Health Center” on page 58

Control Center

You can start the Control Center in the following ways:

- Select Control Center from the Tools menu of another tool.
- Click the  icon from the toolbar of another tool.
- Enter the `db2cc` command.
- On Windows® systems, click the **Start** button and select **Programs** → **IBM DB2** → **General Administration Tools** → **Control Center**.
- On Linux systems, open the **IBM DB2** folder on the desktop and select **Control Center**.

Use the Control Center to manage systems, DB2® Universal Database instances, DB2 Universal Database™ for OS/390® and z/OS™ subsystems, databases, and database objects such as tables and views.

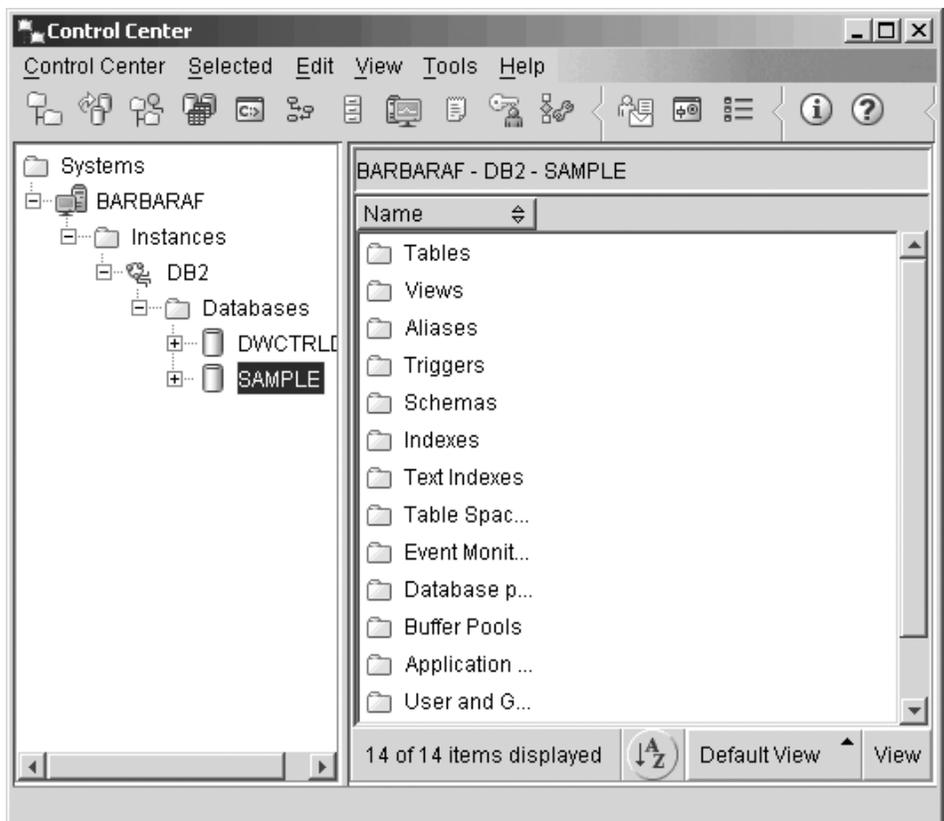


Figure 2. Control Center window

In the Control Center, you can administer all of your systems, instances, databases, and database objects. From the Control Center, you can also open other centers and tools to help you optimize queries, jobs, and scripts; perform data warehousing tasks; create stored procedures; and work with DB2 commands.

The following are some of the key tasks that you can perform with the Control Center:

- Add DB2 systems, federated systems, DB2 for z/OS and OS/390 systems, instances, databases, and database objects to the object tree.
- Manage database objects. You can create, alter, and drop databases, table spaces, tables, views, indexes, triggers, and schemas. You can also manage users.
- Manage data. You can load, import, export, and reorganize data. You can also gather statistics.
- Perform preventive maintenance by backing up and restoring databases or table spaces.
- Configure and tune instances and databases.
- Manage database connections, such as DB2 Connect™ servers and subsystems.
- Manage DB2 for z/OS and OS/390 subsystems
- Manage applications.
- Analyze queries using Visual Explain to look at access plans.
- Launch other tools such as the Command Center and the Health Center.

In many cases, wizards and launchpads are available to help you perform these tasks more quickly and easily.

Note: To schedule tasks through the Control Center, you must first create a tools catalog. If you did not create a tools catalog when you installed DB2, you can select **Tools** —> **Tools Settings** —> **Scheduler Settings** and create one now.

Tool invocation

You can invoke the following tools by selecting their icon on the toolbar:



Control Center



Replication Center



Satellite Administration Center



Data Warehouse Center (with the Data Warehousing option)



Command Center



Task Center



Information Catalog Center (with the Data Warehousing option)



Health Center



Journal



License Center



Development Center



Information Center

To start the Visual Explain tool, right-click a database name and select either **Show Explained Statements History** or **Explain SQL**.

To start the Memory Visualizer, right-click an instance and select **View Memory Usage**.

Storage Management view

To use the Storage Management view, right-click a database, database partition group, or table space and select **Manage Storage**.

Use the Storage Management view to monitor the storage state of a database. The Storage Management view is the graphical interface to the Storage Management tables and stored procedures. In the Storage Management view, you can take storage snapshots for a database, a database partition group, or a table space.

When a table space snapshot is taken, statistical information is collected from the system catalogs and database monitor for tables, indexes, and containers defined under the scope of the given table space. When a database or database partition group snapshot is taken, statistical information is collected for all the table spaces defined in the given database or database partition

group. When a database snapshot is taken, statistical information is collected for all the database partition groups within the database.

Different types of storage snapshots can be used to help you monitor different aspects of storage:

- Space usage can be monitored through snapshots of table spaces.
- On partitioned databases only, data skew (database distribution) can be monitored best through snapshots of database partition groups.
- Cluster ratio of indexes can be captured through both database partition group snapshots and table space snapshots.

The cluster ratio of indexes is presented through the detail view of the index folder. The Storage Management view also enables you to set thresholds for data skew, space usage, and index cluster ratio. If a target object exceeds a specified threshold, the icons beside the object and its parent object in the Storage Management view are marked with a warning flag or an alarm flag.

Which objects you can administer

In the Control Center, you can administer many DB2 objects, including:

- Systems
- Instances
- Databases
- Table spaces
- Tables
- Views
- Indexes
- Schemas
- Triggers
- Stored procedures
- User-defined types
- User-defined functions
- User-defined methods
- Packages
- Aliases
- Event monitors
- Buffer pools
- Users and groups

To see which actions you can perform on an object, select the object and click the right mouse button. A pop-up menu lists the available actions.

Related concepts:

- “Command Center” on page 35
- “Data Warehouse Center” on page 34
- “Health Center” on page 40
- “Development Center” on page 45
- “Information Center” on page 46
- “Journal” on page 42
- “License Center” on page 44
- “Memory Visualizer” on page 48
- “Replication Center” on page 30
- “Task Center” on page 37
- “Visual Explain” on page 52
- “Federated systems in the Control Center” on page 26
- “z/OS and OS/390 objects in the Control Center” on page 28
- “Satellite Administration Center” on page 33
- “Basic GUI navigation concepts explained” on page 15

Related tasks:

- “Backing up a database online for the first time” on page 63
- “Getting started with the Control Center: Control Center help”
- “Using wizards to perform tasks quickly and easily: Control Center help”

Remote object administration in the Control Center

The following sections describe how to administer remote systems through the Control Center:

- “Federated systems in the Control Center”
- “z/OS and OS/390 objects in the Control Center” on page 28

Federated systems in the Control Center

Use the Control Center to set up a federated system. You can create the wrappers, supply the server definitions, identify user mappings, and create nicknames for the data source objects.

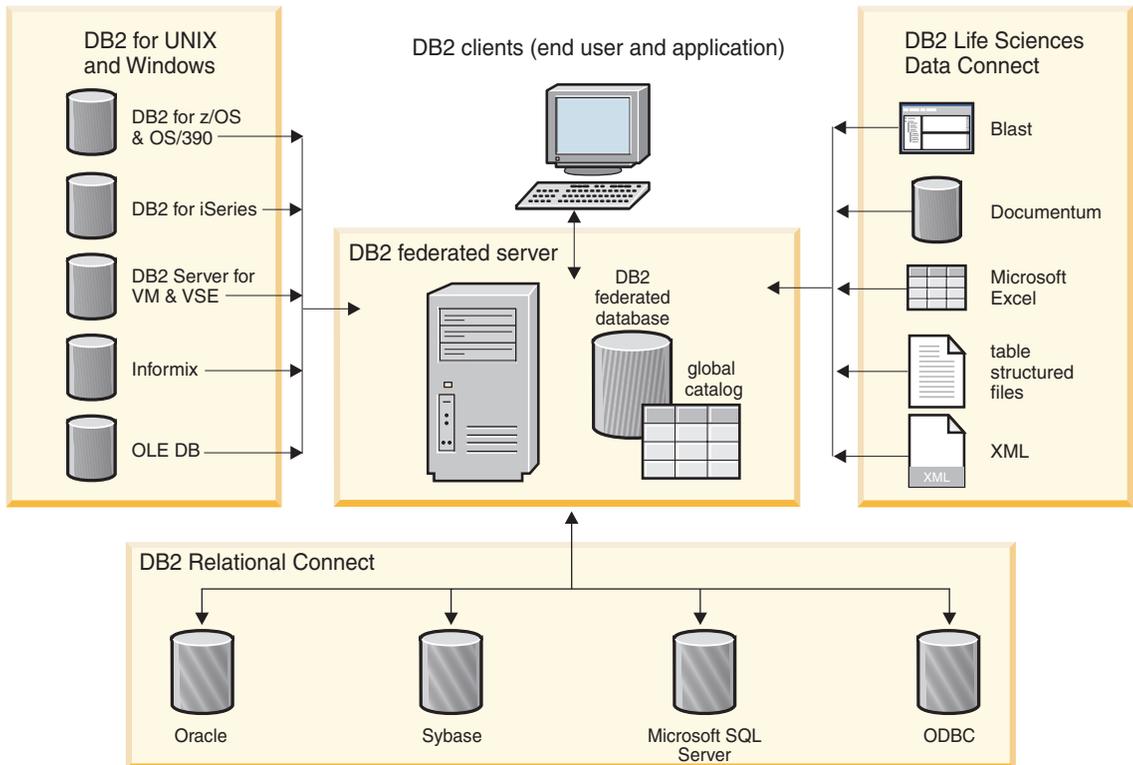


Figure 3. Components of a federated system

Through the Control Center, you can also:

- Create a remote table on a data source, using transparent DDL
- Add columns to a remote table that was created using transparent DDL
- Show all of the remote tables on a data source, whether they were created using transparent DDL or created natively on the data source.

The following optional tasks are available through the command line processor:

- Create data type mappings
- Create function mappings
- Create index specifications

To add a new data source to your federated system:

1. Connect to the federated DB2® database.
2. Create a wrapper for each data source type that you want to include in your federated system.

3. Create a server definition for each remote data source that you want to access. This defines the data source to the federated system.
4. Create a user mapping between your authorization ID to access the federated database and your authorization ID to access the data source.
5. Create a nickname for each data source object (such as a table or view) at the data source that you want to access through DB2.
6. Query the data source object using DB2 database SQL.

Related concepts:

- “Command Center” on page 35
- “Control Center” on page 22
- “Federated systems” on page 14
- “Federated systems” in the *Federated Systems Guide*

z/OS and OS/390 objects in the Control Center

You can manage the following DB2[®] Universal Database for OS/390[®] and z/OS[™] objects in the Control Center:

- Aliases
- Application objects (collections, distinct types, packages, plans, procedures, and user-defined functions)
- Buffer pools
- Catalog tables
- Database users
- Databases
- Data sets
- Indexes
- Locations
- Schemas
- Storage groups
- Synonyms
- Table spaces
- Tables
- Triggers
- Utility objects (data set templates, object lists, and utility procedures)
- Views

To add a DB2 for OS/390 and z/OS Version 6 or later subsystem, use either the Configuration Assistant or the command line processor (CLP) to configure

a connection. With the Configuration Assistant, you can use the Add Database wizard, add a connection manually, or import a connection by using a profile.

To create a new object, right-click the object folder in the object tree. To see which actions you can perform on a particular object, right-click the object name in the Contents pane. A pop-up menu lists the available actions. For example, you can create, alter, or drop a view, as well as see its contents, modify the privileges on it, and show a list of other objects that are related to it.

You can execute many DB2 for OS/390 and z/OS database commands. For example, these are some of the actions that you can perform on a table space:

- Create
- Drop
- Locate
- Load and unload
- Reorganize
- Recover
- Check data
- Run statistics
- Copy
- Quiesce
- Set privileges

With the Create Cloning Session wizard, you can generate the job control language (JCL) jobs that will homogeneously copy (clone) an entire DB2 for OS/390 and z/OS Version 7 or later subsystem. You then manually submit the cloning jobs for execution.

You can also perform many other tasks, including:

- Start or stop a subsystem
- Generate, edit, and save JCL for common tasks such as reorganizing a table space
- Create and work with user-defined functions and stored procedures
- Update table space statistics

Related concepts:

- “Configuration Assistant” on page 55
- “Control Center” on page 22

Related tasks:

- “Getting started with DB2 Universal Database for z/OS and OS/390 objects: CC390 help”

Replication Center

You can start the Replication Center in the following ways:

- Select Replication Center from the Tools menu of another tool.
- Click the  icon from the toolbar of another tool.
- Enter the **db2rc** command.
- On Windows[®] systems, click the **Start** button and select **Programs —> IBM DB2 —> General Administration Tools —> Replication Center**.

Use the Replication Center to administer replication between a DB2[®] database and another relational database (DB2 or non-DB2). From the Replication Center, you can define replication environments, apply designated changes from one location to another, and synchronize the data in both locations.

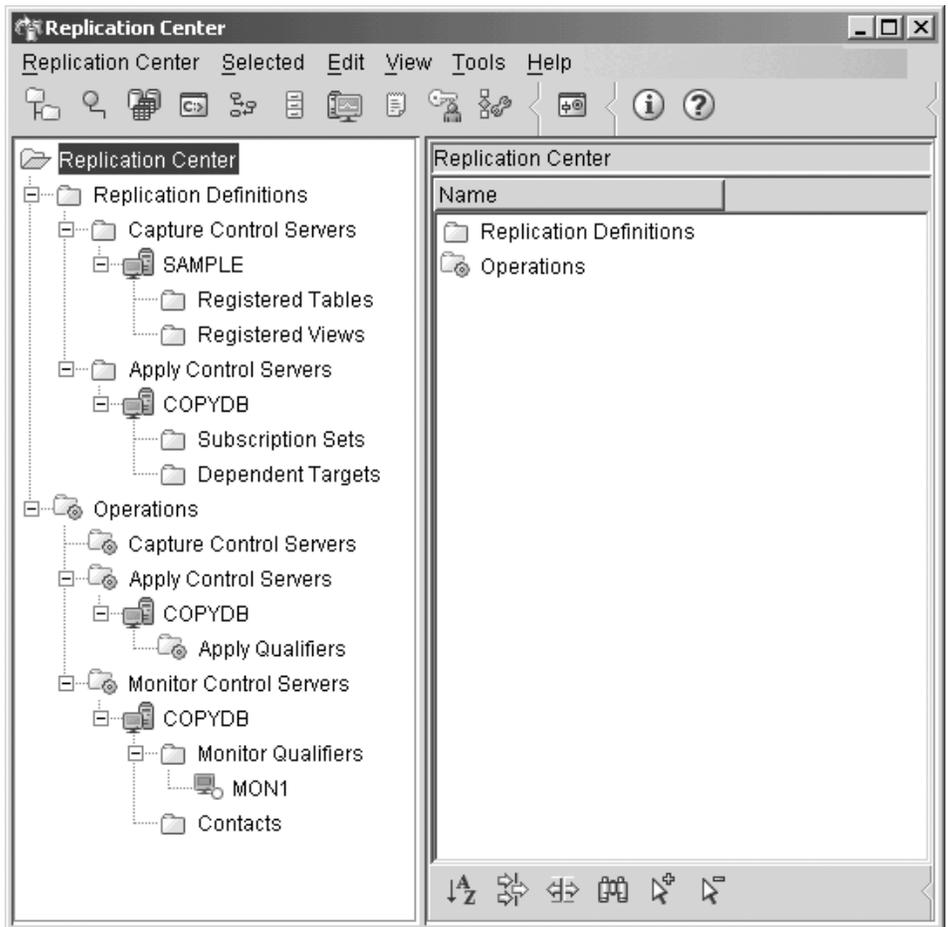


Figure 4. Replication Center window

The Replication Center launchpad guides you through some of the basic replication functions.

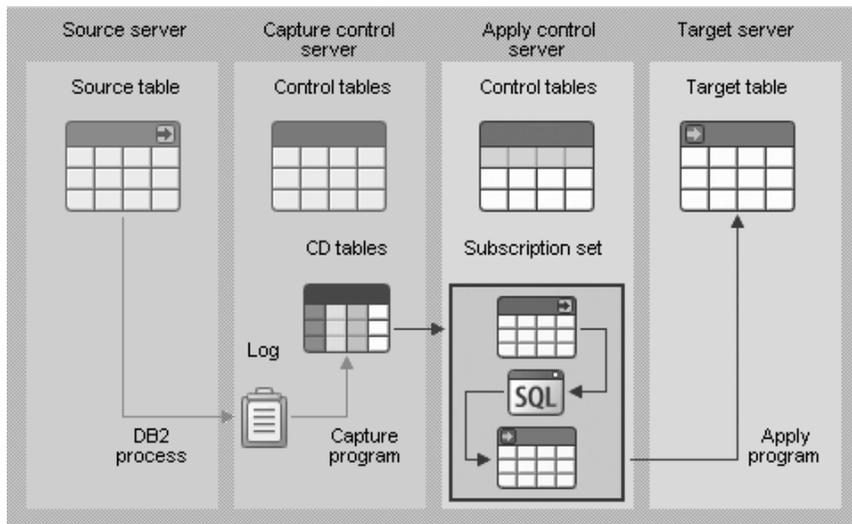


Figure 5. Replication functions

The following are some of the key tasks that you can perform with the Replication Center:

- Create replication control tables
- Register replication sources
- Create subscription sets
- Operate the Capture program
- Operate the Apply program
- Monitor the replication process
- Perform basic troubleshooting for replication

Related concepts:

- “Replication” on page 13

Related tasks:

- “Replicating data” on page 67
- “Setting up for replication” in the *Replication Guide and Reference*
- “Getting started with the Replication Center: Replication Center help”

Satellite Administration Center

You can start the Satellite Administration Center in the following ways:

- Select Satellite Administration Center from the Tools menu of another tool.

- Click the  icon from the toolbar of another tool.

Use the Satellite Administration Center to set up and administer a group of DB2 servers that perform the same business function. These servers, known as satellites, all run the same application and have the same DB2 configuration (database definition) to support this application.

In the Satellite Administration Center, you create a group and then define the satellites as members of this group. You can administer the group as a single entity, instead of having to administer each satellite separately. If you later acquire additional DB2 servers that will perform the same business function, you can add them to the group as additional satellites. The administration solution provided by the satellite environment is fully scalable.

Information about the satellite environment is stored in the satellite control database (SATCTLDB). This database records, among other things, which satellites are in the environment and which group each satellite belongs to. The DB2 server that contains this database is known as the satellite control server.

You administer a group of satellites by creating batches, using the Satellite Administration Center. The batches contain scripts that set up and maintain the database definition that supports the application on the satellite. Each satellite regularly connects to its satellite control server, and downloads any scripts that apply to it. The satellite executes these scripts locally, and uploads the results back to the satellite control database.

This process of downloading batch scripts, executing them, and reporting the results of the batch execution is known as synchronization. A satellite synchronizes to maintain its consistency with the other satellites that belong to its group.

Although the satellites in a group run the same application, they do not necessarily run the same version of the application. For example, if you are in the process of upgrading to version 4 of your application, you might have some satellites at the version 3 level and some at version 4.

The application version is set on each satellite. When you want to deploy a new version of the application, you create a new set of batches that set up and maintain the new database definition that supports the new version of

the application. As you move each satellite over to the new version of the application, you update the application version on the satellite.

When a satellite synchronizes, it uploads its application version. The satellite control server ensures that the satellite can only download and execute the batches that correspond to its application version. This allows you to stage the deployment of a new version of the application. Satellites with the new version of the application will not download and execute the batches of the previous application version, and satellites with the previous version of the application will not download and execute the batches for the new application version. Each satellite will have the database definition that is required to support its version of the application.

Before you can use the Satellite Administration Center, you must do the following:

- Install the satellite control server.
- Create the satellite control database.
- Catalog the satellite control database in the Control Center, if it is not already catalogued.

Note: You should only administer the satellite environment from the Satellite Administration Center.

Related concepts:

- “Satellite Control Server” in the *Installing and Administering a Satellite Environment*
- “Satellites” in the *Installing and Administering a Satellite Environment*
- “Control Center” on page 22

Related tasks:

- “Getting started with the Satellite Administration Center : Satellite Administration Center help”

Data Warehouse Center

If you have installed the Data Warehousing option, you can start the Data Warehouse Center in the following ways:

- Select Data Warehouse Center from the Tools menu of another tool.



- Click the  icon from the toolbar of another tool.
- Enter the **db2dwc** command.

- On Windows[®] systems, click the **Start** button and select **Programs** → **IBM DB2** → **Business Intelligence Tools** → **Data Warehouse Center**.

Use the Data Warehouse Center to automate the extraction and transformation of data for your data warehouse. You can use the Data Warehouse Center to define the processes that move and transform data for the warehouse. You can also use the Data Warehouse Center to schedule, maintain, and monitor these processes.

The following are some of the key tasks that you can perform with the Data Warehouse Center:

- Set up the warehouse
- Create a star schema
- Set up DB2[®] and non-DB2 data sources
- Set up access to the warehouse
- Define and run processes
- Load and transform data
- Calculate statistics
- Manage metadata and the control database

To learn more about data warehousing, see the following tutorials:

- *Business Intelligence Tutorial: Introduction to the Data Warehouse*
- *Business Intelligence Tutorial: Extended Lessons in Data Warehousing*

Related concepts:

- “What is data warehousing?” in the *Data Warehouse Center Administration Guide*
- “DB2 tutorials” on page 95

Related tasks:

- “Getting Started with the Data Warehouse Center: Data Warehouse Center help”

Command Center

You can start the Command Center in the following ways:

- Select Command Center from the Tools menu of another tool.
- Click the  icon from the toolbar of another tool.
- On Windows[®] systems, click the **Start** button and select **Programs** → **IBM DB2** → **Command Line Tools** → **Command Center**.

Use the Command Center to execute DB2® commands and SQL statements; to execute MVS console commands; to work with command scripts; and to view a graphical representation of the access plan for explained SQL statements.

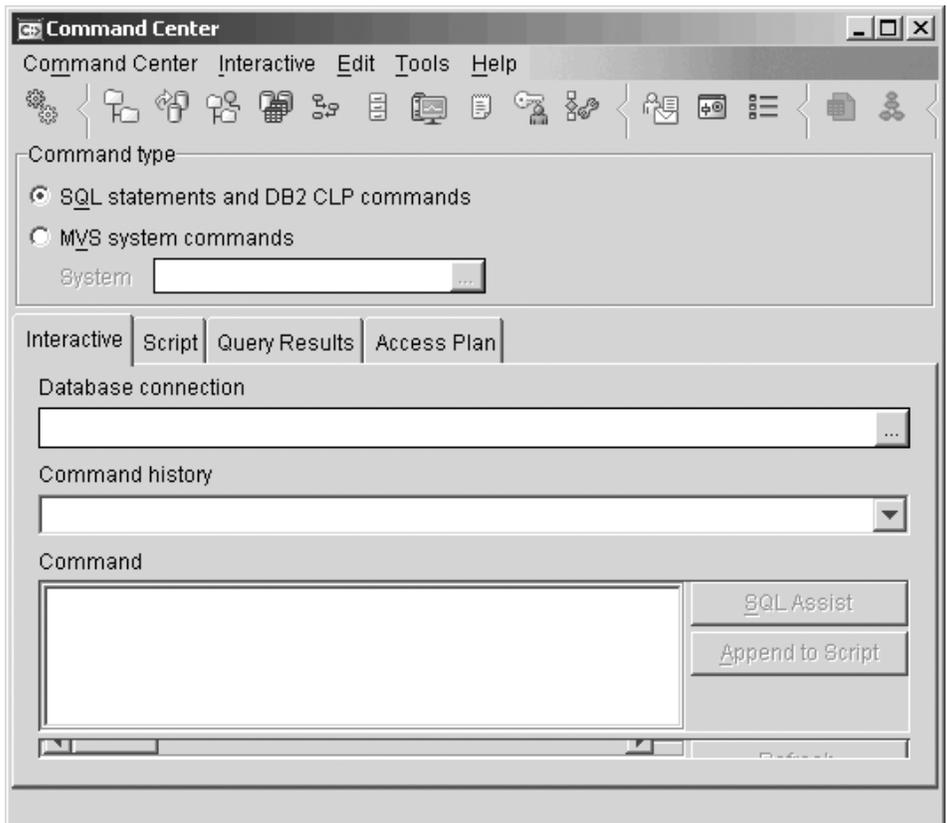


Figure 6. Command Center window

On the **Interactive** page, you can perform the following actions:

- Execute an SQL statement or DB2 CLP command. You do not need to precede the command by *DB2*.
- Run a command or statement, by clicking the  icon (at the left of the toolbar).

On the **Script** page, you can perform the following actions:

- Execute commands in sequence.
- Create and save a script. You can optionally store a saved script in the Task Center, where you can schedule the script to run at a specific time.
- Run an existing script.

- Schedule a task.

On the **Query Results** page, you can see the results of the queries. You can also save the results or edit the contents of the table.

On the **Access Plan** page, you can see the access plan for any explainable statement that you specified on the Interactive page or the Script page. DB2 generates the access plan when it compiles the SQL statement. You can use this information to tune your queries for better performance.

Note: If you specify more than one statement on the Script page, an access plan is created only for the first statement.

To invoke the SQL Assist tool, click the SQL Assist button on the Interactive page. To invoke the Visual Explain tool, execute an explainable statement on the Interactive page or the Script page.

The Web version of the Command Center includes most of these features but it *does not* currently include Visual Explain or SQL Assist.

Related concepts:

- “SQL Assist” on page 50
- “Task Center” on page 37
- “Visual Explain” on page 52
- “Web Command Center” on page 57

Task Center

You can start the Task Center in the following ways:

- Select Task Center from the Tools menu of another tool.
- Click the  icon from the toolbar of another tool.
- Enter the **db2tc** command.
- On Windows® systems, click the **Start** button and select **Programs** → **IBM DB2** → **General Administration Tools** → **Task Center**.

Note: To use the Task Center, you must first create a tools catalog. If you did not create a tools catalog when you installed DB2, you can select **Tools** → **Tools Settings** → **Scheduler Settings** and create one now.

Use the Task Center to run tasks, either immediately or according to a schedule, and to notify people about the status of completed tasks. The Task Center includes functionality from the Script Center in previous versions of DB2, plus additional functionality.

A task is a script, together with the associated success conditions, schedules, and notifications. You can create a task within the Task Center, create a script within another tool and save it to the Task Center, import an existing script, or save the options from a DB2® dialog or wizard such as the Load wizard. A script can contain DB2, SQL, or operating system commands.

For each task, you can do the following:

- Schedule the task
- Specify success and failure conditions
- Specify actions that should be performed when this task completes successfully or when it fails
- Specify e-mail addresses (including pagers) that should be notified when this task completes successfully or when it fails

You can specify conditional coding by creating task actions. Each task action consists of a task and the action that should be performed on the task. For example, task 1 could have the following task actions:

- If task 1 is successful, task action A enables the schedule for task 2.
- If task 1 fails, task action B runs task 3.

You can also create a grouping task, which combines several tasks into a single logical unit of work. When the grouping task meets the success or failure conditions that you define, any follow-on tasks are run. For example, you could combine three backup scripts into a grouping task and then specify a reorganization as a follow-on task that will be executed if all of the backup scripts execute successfully.

To see the most recent information, click the  button. You can view execution results as follows:

- To see the results from the last time a task was executed, right-click the task and select **Show Results**.
- To see the results from every time a task was executed, open the Journal. In the Journal, you can right-click the task execution that you are interested in and select **Show Results**.

To see help text for a message that appears in the execution results, open the Command Center and enter the error message preceded by a question mark. For example, enter `? SQL0100` to get more information about message SQL0100W.

Several customizable saved views are available via the pull-up menu at the bottom of the Task Center.

Related concepts:

- “Command Center” on page 35
- “Journal” on page 42

Related tasks:

- “Getting started with the Task Center: Task Center help”

Information Catalog Center

If you have installed the Data Warehousing option and initialized a catalog, you can start the Information Catalog Center in the following ways:

- Select Information Catalog Center from the Tools menu of another tool.
- Click the  icon from the toolbar of another tool.
- Enter the `db2icc` command.
- On Windows[®] systems, click the **Start** button and select **Programs** → **IBM DB2** → **Business Intelligence Tools** → **Information Catalog Center**.

Note: To use the Information Catalog Center, you must first create a tools catalog. If you did not create a tools catalog when you installed DB2, you can select **Tools** → **Tools Settings** → **Scheduler Settings** and create one now.

Use the Information Catalog Center to organize and search for business information. The Information Catalog Center streamlines the search for business information by providing a single information catalog with a robust searching capability. It does for your organization what an electronic card catalog does for a library, and more.

The information catalog contains descriptive data about source information. For example, the information catalog can include the type of information, a description of the information, what it contains, who owns and updates it, and where and how to get to it.

You can search for specific objects, view any relationships that an object participates in, view an object's lineage, and create comments for objects. If you have the authority, you can also create new objects in the information catalog.

After you find the information that you need, you can start a spreadsheet program, word processor, graphics tool, or other application from the Information Catalog Center. The Information Catalog Center remains in the background while you work with the retrieved information in the application where the information was originally created.

To learn more about the Information Catalog Center, see the *Information Catalog Center Tutorial*.

Related concepts:

- “DB2 tutorials” on page 95
- “Data Warehouse Center” on page 34

Related tasks:

- “Getting started with the Information Catalog Center” in the *Information Catalog Center Administration Guide*
- “Information Catalog Center -- Overview : Information Catalog Center help”

Health Center

You can start the Health Center in the following ways:

- Select Health Center from the Tools menu of another tool.
- Click the  icon from the toolbar of another tool.
- Enter the **db2hc** command.
- On Windows® systems, click the **Start** button and select **Programs** —> **IBM DB2** —> **Monitoring Tools** —> **Health Center**.
- Click the health center status beacon when it appears on a DB2® window. On the Tools Settings window, you can configure this beacon to appear on every DB2 window whenever an alert occurs.

Use the Health Center to monitor the state of the database environment and make any necessary changes.

When you use DB2, a health monitor continuously monitors a set of health indicators. If the current value of a health indicator is outside the acceptable operating range defined by its warning and alarm thresholds, the health

monitor generates a health alert. DB2 comes with a set of predefined threshold values, which you can later customize. For example, you can customize the alarm and warning thresholds for the amount of space used in a table space.

Depending on the configuration of the DB2 instance, the following actions can occur when the health monitor generates an alert:

- An entry is written in the administration notification log, which you can read from the Journal.
- The health center status beacon appears in the lower right corner of the DB2 GUI Tools window.
- A script or task is executed.
- An e-mail or pager message is sent to the contacts that you specify for this instance.

The following are some of the key tasks that you can perform with the Health Center:

- View the status of the database environment. Beside each object in the navigation tree, an icon indicates the most severe alert for the object (or for any objects contained by that object). For example, a green diamond icon beside an instance means that the instance and the databases contained in the instance do not have any alerts.
- View the alerts for an instance or a database. When you select an object in the navigation tree, the alerts for that object are shown in the pane to the right.
- View detailed information about an alert, and recommended actions. When you double-click an alert, a notebook appears. The first page shows the details for the alert. The second page shows any recommended actions.
- Configure the health monitor settings for a specific object, and the default settings for an object type or for all objects within an instance.
- Select which contacts will be notified of alerts with an e-mail or pager message.
- Review the history of alerts for an instance.

You can use the toggle buttons at the top of the navigation bar to filter the alerts according to their severity:

-  Displays only alarms (the most severe type of alert)
-  Displays alarms and warnings
-  Displays alarms, warnings, and attentions

-  Displays all alerts

The Web version of the Health Center includes a subset of these features.

Related concepts:

- “Journal” on page 42
- “Memory Visualizer” on page 48
- “Web Health Center” on page 58
- “Indoubt Transaction Manager” on page 56

Related tasks:

- “Reacting to alerts” on page 66

Journal

You can start the Journal in the following ways:

- Select Journal from the Tools menu of another tool.
- Click the  icon from the toolbar of another tool.
- On Windows[®] systems, click the **Start** button and select **Programs** → **IBM DB2** → **General Administration Tools** → **Journal**.

The Journal displays historical information about tasks, database actions and operations, Control Center actions, messages, and alerts.

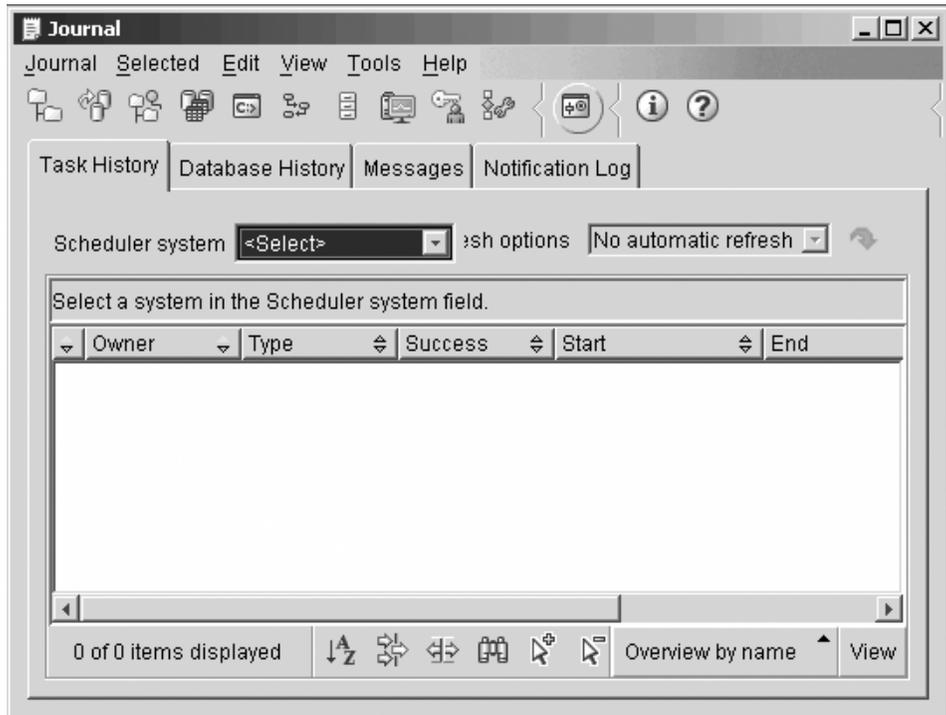


Figure 7. Journal window

To see the most recent information, click the  button.

The **Task History** page shows the results of tasks that were previously executed. You can use this information to estimate how long future tasks will run.

The Task History page contains one row for each execution of a task. The Task Center contains only one row for each task no matter how many times the task is executed. Each row in the Task Center could be directly related to multiple rows in the Task History page of the Journal.

For each completed execution of a task, you can perform the following actions:

- View the execution results
- View the task that was executed
- Edit the task that was executed
- View the task execution statistics
- Remove the task execution object from the Journal

To perform one of these actions, right-click a completed task execution and select the corresponding action from the pop-up menu that appears.

The **Database History** page shows information from the recovery history file. This file is updated when various operations are performed, including:

- Backup
- Restore
- Roll forward
- Load
- Reorg

This information could be useful if you need to restore a database or tablespace.

The **Messages** page shows messages that were previously issued from the Control Center and other GUI tools.

The **Notification Log** page shows information from the administration notification log. You can filter the list. For example, you can show only alerts from the health monitor.

Several customizable saved views are available via the pull-up menu at the bottom of the Journal.

Related concepts:

- “Health Center” on page 40
- “Task Center” on page 37

License Center

You can start the License Center in the following ways:

- Select License Center from the Tools menu of another tool.
- Click the  icon from the toolbar of another tool.

Use the License Center to display license status and usage information for DB2® products that are installed on your system. You can also use the License Center to configure your system for license monitoring.

The following are some of the key tasks that you can perform with the License Center:

- Add or remove licenses or registered users

- Change the enforcement or license type policies, the number of concurrent users, the number of licensed processors, or the number of internet processor licenses
- View processor status information, concurrent users policy information, license information, and user statistics or details

The License page displays the following information:

- Product name
- Version information
- Expiry date
- Registered user policy
- Number of entitled DB2 Connect™ users
- Current number of DB2 Connect users
- Enforcement policy
- Number of processors and number of licensed processors (for DB2 Universal Database™ Enterprise Server Edition)

Related tasks:

- “Getting started with the License Center : License Center help”

Development Center

You can start the Development Center in the following ways:

- Select Development Center from the Tools menu of another tool.
- Click the  icon from the toolbar of another tool.
- Enter the **db2dc** command.
- On Windows® systems, click the **Start** button and select **Programs** → **IBM DB2** → **Development Tools** → **Development Center**.

Use the Development Center to develop stored procedures, user-defined functions (UDFs), and structured types. The Development Center includes functionality from the Stored Procedure Builder in previous versions of DB2, plus additional functionality.

A set of wizards makes it easy to perform your development tasks. The Development Center provides a single development environment that supports the entire DB2® family ranging from the workstation to z/OS.

With the Development Center, you can:

- Create, build, and deploy Java™ and SQL stored procedures

- Create, build, and deploy user-defined functions:
 - SQL scalar UDFs
 - SQL table UDFs
 - UDFs that read MQSeries® messages
 - UDFs that access OLE DB data sources
 - UDFs that extract data from XML documents
- Debug SQL stored procedures using the integrated debugger
- Create and build structured types
- See the contents of the server for each database connection that is in your project or that you have explicitly added to the Server View
- View and work with other database objects such as tables, triggers, and views
- Export and import routines and project information

The Development Center also provides a DB2 Development Add-In for each of the following development environments:

- Microsoft® Visual C++
- Microsoft Visual Basic
- Microsoft Visual InterDev

With the add-ins, you can easily access the features of the Development Center and other DB2 centers from your Microsoft development environment. This makes it easier for you to develop and incorporate stored procedures and UDFs into your DB2 application development.

Related concepts:

- “SQL Assist” on page 50

Related tasks:

- “Getting started with the Development Center: Development Center help”

Information Center

You can start the Information Center in the following ways:

- Click the  icon from the toolbar of another tool.
- Enter the **db2ic** command.
- On Windows® systems, click the **Start** button and select **Programs** → **IBM DB2** → **Information** → **Information Center**.

Use the Information Center to find information about DB2 Universal Database. You can select topics from the DB2 documentation or the DB2 online help. You can also update your local HTML documentation, as described in “Updating the HTML documentation installed on your machine” on page 88.

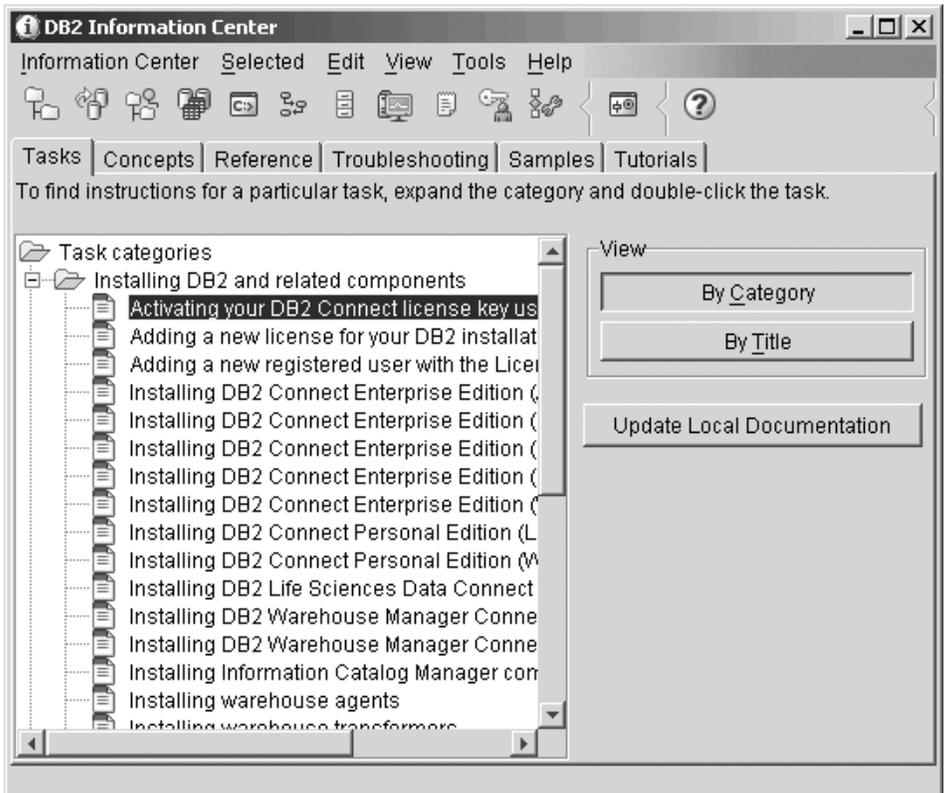


Figure 8. Information Center window

When you select a topic from the DB2 documentation, the page opens up in a browser window.

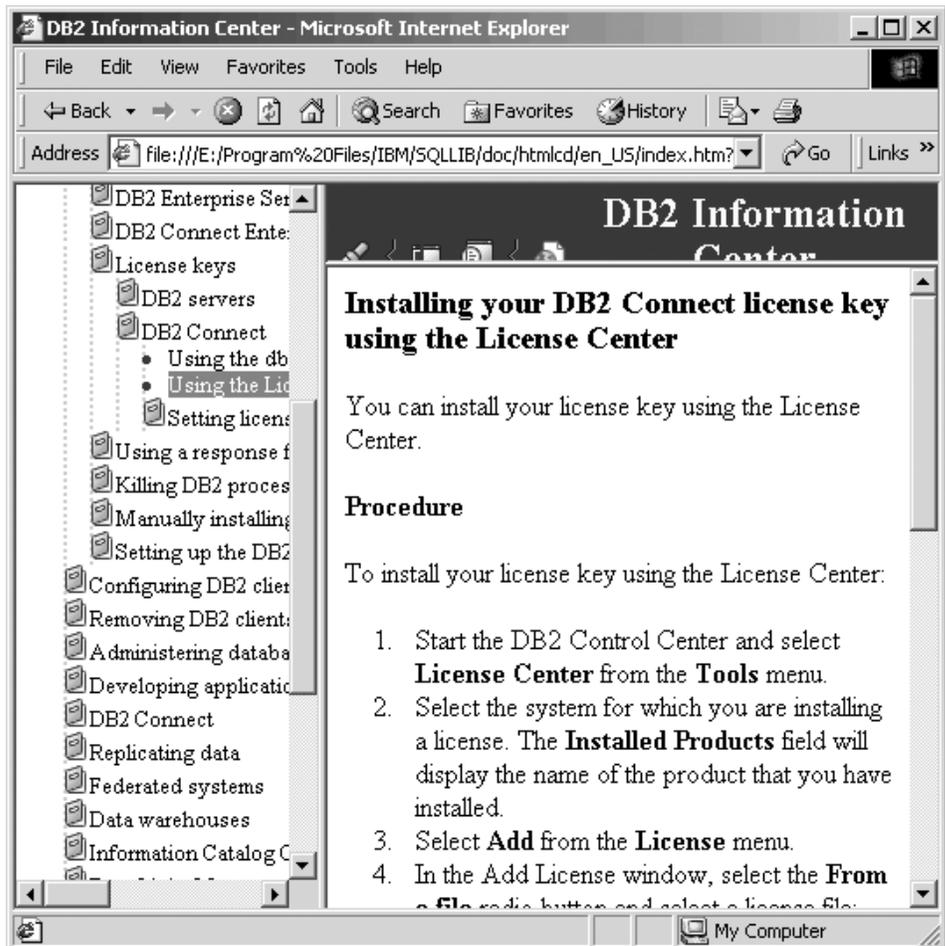


Figure 9. Information Center HTML page

From an Information Center page in a browser window, you can find information in any of the following ways:

- Expand the appropriate category in the navigation tree
- Select the index item in the navigation tree, then select the page that you want
- Click **Search** in the navigation toolbar, and enter one or more search terms

Memory Visualizer

You can start the Memory Visualizer in the following ways:

- In the Control Center, right-click an instance and select **View memory usage**.

- On Windows[®] systems, click the **Start** button and select **Programs** → **IBM DB2** → **Monitoring Tools** → **Memory Visualizer**.

Use the Memory Visualizer to monitor the memory allocation of a DB2[®] Version 8 instance.

The top pane of the window shows memory components organized in a navigation tree. Historical values and alarm and warning thresholds are shown to the right of each component. The lower pane shows a memory usage plot.

The high-level memory components include:

- Database manager shared memory
- Database global memory
- Application global memory
- Agent/Application shared memory
- Agent private memory

Each high-level component is divided into lower-level components that determine how the memory is allocated and deallocated. For example, memory is allocated and deallocated when the database manager starts, when a database is activated, and when an application connects to a database.

To display a memory usage plot, select the Show Plot check box for the desired component or components. The usage plots for different components are shown with a different color and shape. For each change to a configuration parameter, the usage plot also shows the original value, the new value, and the time when the value changed.

The following are some of the key tasks that you can perform with the Memory Visualizer:

- View overall memory usage
- Specify which memory information to display and which information to hide for a DB2 instance and its databases
- Update the configuration parameters for an individual memory component to prevent it from using too much or too little memory
- Save the memory allocation data
- Load memory allocation data from a file into a Memory Visualizer window

Related concepts:

- “Health Center” on page 40
- “Indoubt Transaction Manager” on page 56

Spatial Extender

Use DB2 Spatial Extender to increase your business intelligence capabilities by adding spatial data to your database. You can query and analyze location information and information about geographical features as well as standard business data. For example, a real estate agent could determine the distance from a house to the nearest school. A retail chain could determine the best location for a new store based on demographic data and the locations of competing stores.

You can analyze several types of geographical information, including:

- The shape of natural features such as rivers, lakes, and mountains
- User-defined areas, such as the marketing area serviced by a particular business unit
- Specific locations, such as the intersection where an auto accident occurred

To use DB2 Spatial Extender, do the following:

1. Purchase and install DB2 Spatial Extender.
2. In the Control Center, right-click a database and select **Spatial Extender** → **Enable**. This prepares the database to accept spatial data types and work with geographic information systems.
3. In the Control Center, right-click a spatially enabled database, table, or view. Select **Spatial Extender**, then select the function that you want to perform. You can also invoke DB2 Spatial Extender functions directly from the command line processor.

Related concepts:

- “Control Center” on page 22
- “Spatial data” in the *Spatial Extender User’s Guide and Reference*

SQL Assist

To start the SQL Assist tool, click the SQL Assist button in another tool. The SQL Assist button exists at various locations in the Control Center, Command Center, Replication Center, Development Center, and Data Warehouse Center.

Each tool allows you to create SQL statements that are appropriate in the context of that tool. In some tools, you can create only SELECT statements. In other tools, you can also create INSERT, UPDATE, and DELETE statements.

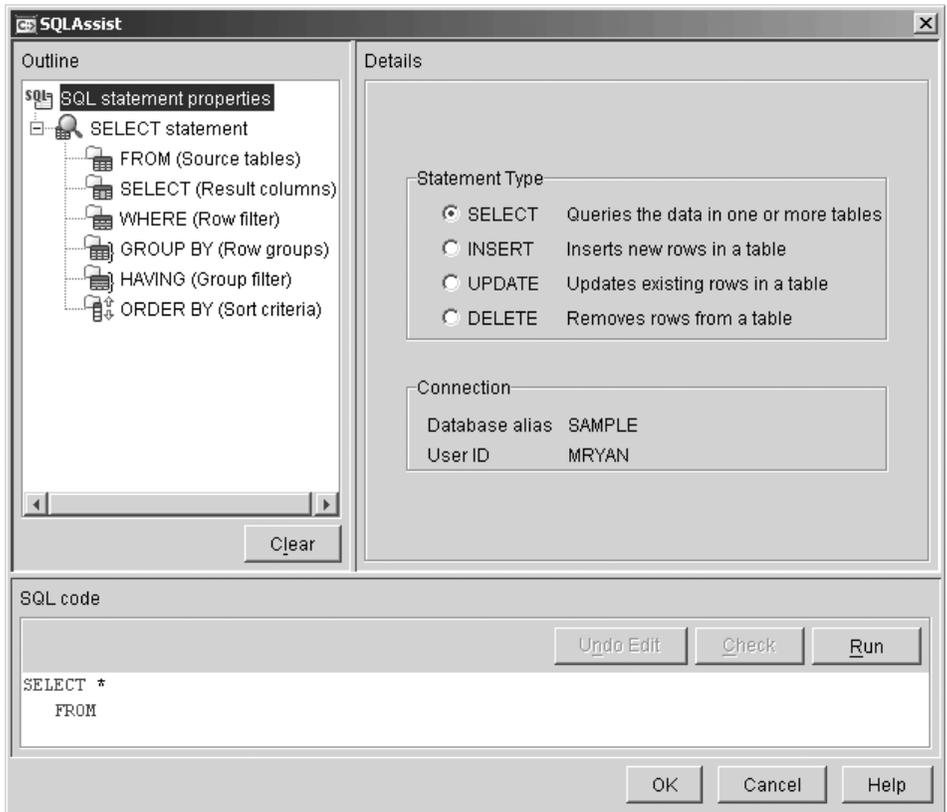


Figure 10. SQL Assist window

The **Outline** area contains a high-level representation of the current SQL statement type. Select an element of the SQL statement.

The **Details** area displays the appropriate panel for the element that you selected. You build your SQL statement by working with this panel. The changes that you make will be reflected in the **SQL code** area.

The **SQL code** area displays the current SQL statement. Colors are used to highlight the syntax. You can edit the SQL statement or paste an existing SQL statement into this area. If you make any changes in this area, you will not be able to take any actions in the **Details** area until you validate the syntax of the modified SQL.

To validate the syntax, click the **Check** button. If the syntax is correct, the **Details** area will be enabled again, and its content will be updated to reflect your changes. To run the SQL statement, click the **Run** button.

Note: The Data Warehouse Center currently invokes a back-level version of the SQL Assist tool (from DB2® version 7), which does not support all of the functionality that is described in this book. For more information about the back-level SQL Assist tool, see the online help from within that tool.

Related concepts:

- “Command Center” on page 35
- “Control Center” on page 22
- “Data Warehouse Center” on page 34
- “Development Center” on page 45
- “Replication Center” on page 30

Visual Explain

You can start Visual Explain in the following ways:

- In the Control Center, right-click a database name and select either **Show Explained Statements History** or **Explain SQL**.
- In the Command Center, execute an explainable statement on the Interactive page or the Script page.

Use Visual Explain to view the access plan for explained SQL statements as a graph. You can use the information in the graph to tune your SQL queries for better performance.

An access plan graph shows details of:

- Tables (and their associated columns) and indexes
- Operators (such as table scans, sorts, and joins)
- Table spaces and functions
- Cost (in timerons)

You can use Visual Explain to perform the following tasks:

- View the statistics that were used at the time of optimization. You can compare these statistics to the current catalog statistics to help you determine whether rebinding the package might improve performance.
- Determine whether or not an index was used to access a table. If an index was not used, Visual Explain can help you determine which columns might benefit from being indexed.
- View the effects of performing various types of tuning by comparing the before and after versions of the access plan graph for a query.

- Obtain information about each operation in the access plan, including the total estimated cost and number of rows retrieved (cardinality).

Access plan

Certain data is necessary to resolve an explainable SQL statement. An access plan specifies an order of operations for accessing this data. It lets you view statistics for selected tables, indexes, or columns; properties for operators; global information such as table space and function statistics; and configuration parameters that are relevant to optimization. With Visual Explain, you can view the access plan for an SQL statement in graphical form.

The optimizer produces an access plan whenever an explainable SQL statement is compiled. This happens at prepare or bind time for static statements, and at run time for dynamic statements.

An access plan is an estimate based on the available information, such as:

- Statistics in system catalog tables, which you can update by using the RUNSTATS command
- Configuration parameters
- Bind options
- The query optimization class

The cost information for an access plan is the optimizer's best estimate of the resource usage for a query. The actual elapsed time for a query may vary depending on factors that are independent of DB2® (for example, the number of other applications that are running at the same time).

Access plan graph

An access plan graph includes information from several sources, as shown in the following illustration. Based on various inputs, the optimizer chooses an access plan. Visual Explain displays this access plan in an access plan graph. Each node in the graph represents a table, an index, or an operation on a table or index. The links between the nodes represent the flow of data.

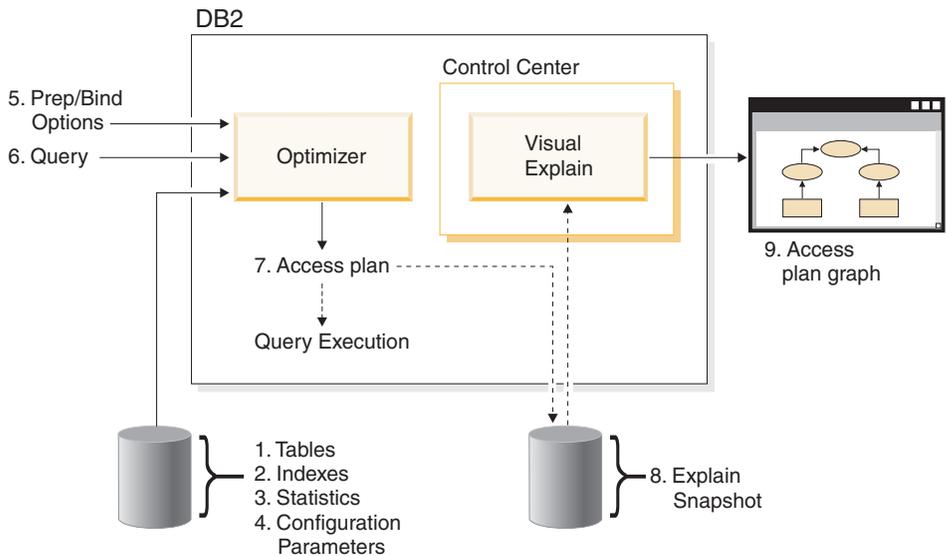


Figure 11. Visual Explain procedure

The previous illustration shows the following steps. Broken lines indicate steps that are required for Visual Explain.

1. Tune your table design and reorganize table data
2. Create appropriate indexes
3. Use the RUNSTATS command to provide the optimizer with current statistics
4. Choose appropriate configuration parameters
5. Choose appropriate bind options
6. Design queries to retrieve only the required data
7. Create an access plan
8. Create explain snapshots
9. Display and use an access plan graph

For example, to use Visual Explain, first update the current statistics using the RUNSTATS command on the tables and indexes that the statement uses. These statistics, the configuration parameters, bind options, and the query itself are used by the optimizer to create an access plan and an explain snapshot when the package is bound. Visual Explain uses the resulting explain snapshot to display the access plan graph for the statement.

To learn more about Visual Explain, see the *Visual Explain Tutorial*.

Related concepts:

- “Explain tools” in the *Administration Guide: Performance*
- “DB2 tutorials” on page 95
- “Command Center” on page 35
- “Control Center” on page 22

Configuration Assistant

You can start the Configuration Assistant in the following ways:

- Enter the **db2ca** command.
- On Windows® systems, click the **Start** button and select **Programs** → **IBM DB2** → **Set-up Tools** → **Configuration Assistant**.
- On Linux systems, open the **IBM DB2** folder on the desktop and select **Configuration Assistant**.

Use the Configuration Assistant to configure your clients. You can also use it as a light-weight alternative to the Control Center, in situations where you do not want to install the complete set of GUI tools.

You must configure your DB2 clients so they can work with the available objects, as follows:

- To access an instance or database on another server or system, DB2® must catalog that system in the node directory of the client.
- To access a database, DB2 must catalog database information in the database directory of the client.

From the Configuration Assistant, you can perform various tasks including:

- Add new database objects
- Work with existing database objects
- Bind applications
- Set database manager configuration parameters
- Import and export configuration information
- Set DB2 registry parameters
- Change passwords
- Test connections
- Configure CLI parameters

The graphical interface makes these complex tasks easier through:

- Wizards that help you perform certain tasks
- Dynamic fields that are activated based on your input choices
- Hints that help you make configuration decisions

- Discovery that can retrieve information about selected database objects

The Configuration Assistant displays a list of the databases to which your applications can connect. Each database is identified by its database alias. You can use the Add Database wizard to add databases to the list. You can use the Change Database wizard to alter the information that is associated with databases in the list.

From the View menu, you can select an advanced view, which uses a notebook to organize connection information by object: Systems, Instance Nodes, Databases, Database Connection Services (DCS), and Data Sources. You can use the notebook pages to perform object-specific actions.

Related tasks:

- “Getting started with the Configuration Assistant: Configuration Assistant help”

Indoubt Transaction Manager

You can start the Indoubt Transaction Manager in the following ways:

- Enter the **db2indbt** command.
- On Windows[®] systems, click the **Start** button and select **Programs** —> **IBM DB2** —> **Monitoring Tools** —> **Indoubt Transaction Manager**.

Use the Indoubt Transaction Manager to work with global transactions that are in an indoubt state. For example, a broken communication line could leave a transaction prepared but not yet committed or rolled back.

In most cases, you should simply wait for the Transaction Manager to resynchronize. If you cannot wait, you can use the Indoubt Transaction Manager with great care as a last resort.

You can perform the following actions on a transaction:

- Commit
- Rollback
- Forget (cancel a previous Commit or Rollback)

Related concepts:

- “Health Center” on page 40
- “Memory Visualizer” on page 48

Web tools

The following DB2 Web Tools provide access to the DB2 environment through a secure Web-enabled interface:

- “Web Command Center”
- “Web Health Center” on page 58

Web Command Center

Use the Web version of the Command Center to execute SQL statements, DB2 commands, and OS commands against a remote DB2 server, either interactively or in scripts. This is particularly useful for database administrators and experienced users who need to be able to work with the database at any time, from any place. The Web version of the Command Center includes most of the features of the full Command Center, but it *does not* currently include Visual Explain or SQL Assist.

To access the DB2 server, you use an HTML 4.0-compliant Web browser that runs on any computing device from a desktop computer to a Personal Digital Assistant (PDA). The browser is a client, which communicates with a middle-tier application server. The application server communicates with the DB2 server. Secure Socket Layer (SSL), the same technology that is used for most Internet banking and secure e-commerce transactions, is used for all communication.

The client can be any computing device that supports an HTML 4.0-compliant Web browser. The browser does not need to support Java, JavaScript, or cookies. The middle tier must have DB2[®] Web Tools installed. DB2 Web Tools requires an application server, such as IBM[®] WebSphere[®] or BEA WebLogic.

If you are the database administrator of a large center, you could use the Web tools as follows:

- While you are away from the office, attending a conference, you receive an urgent message on your PDA about a database problem.
- Using the PDA, you securely log on to a portal where DB2 Web Tools is installed.
- You use the Web versions of the Health Center and the Command Center to investigate the problem.
- You use the Web version of the Command Center to execute SQL, DB2, and OS commands that fix the problem.

Related concepts:

- “DB2 Web Command Center (DB2 Web Tools)” in the *Installation and Configuration Supplement*
- “Command Center” on page 35

- “Web Health Center” on page 58

Related tasks:

- “Deploying DB2 Web Tools on WebSphere application servers” in the *Installation and Configuration Supplement*
- “Deploying DB2 Web Tools on WebLogic application servers” in the *Installation and Configuration Supplement*
- “Deploying DB2 Web Tools on other application servers” in the *Installation and Configuration Supplement*

Web Health Center

Use the Web version of the Health Center to monitor the health of remote databases. This is particularly useful for database administrators and experienced users who need to be able to work with the database at any time, from any place.

The Web version of the Health Center includes a subset of the features of the full Health Center. You can view the active alerts for an instance, view the details, history, and resolution recommendations for an alert, and apply some recommendations through the Web version of the Command Center.

To access the DB2 server, you use an HTML 4.0-compliant Web browser that runs on any computing device from a desktop computer to a Personal Digital Assistant (PDA). The browser is a client, which communicates with a middle-tier application server. The application server communicates with the DB2 server. Secure Socket Layer (SSL), the same technology that is used for most Internet banking and secure e-commerce transactions, is used for all communication.

The client can be any computing device that supports an HTML 4.0-compliant Web browser. The browser does not need to support Java, JavaScript, or cookies. The middle tier must have DB2[®] Web Tools installed. DB2 Web Tools requires an application server, such as IBM[®] WebSphere[®] or BEA WebLogic.

If you are the database administrator of a large center, you could use the Web tools as follows:

- While you are away from the office, attending a conference, you receive an urgent message on your PDA about a database problem.
- Using the PDA, you securely log on to a portal where DB2 Web Tools is installed.
- You use the Web versions of the Health Center and the Command Center to investigate the problem.
- You use the Web version of the Command Center to execute SQL, DB2, and OS commands that fix the problem.

Related concepts:

- “DB2 Web Health Center (DB2 Web Tools)” in the *Installation and Configuration Supplement*
- “Health Center” on page 40
- “Web Command Center” on page 57

Related tasks:

- “Deploying DB2 Web Tools on WebSphere application servers” in the *Installation and Configuration Supplement*
- “Deploying DB2 Web Tools on WebLogic application servers” in the *Installation and Configuration Supplement*
- “Deploying DB2 Web Tools on other application servers” in the *Installation and Configuration Supplement*

Chapter 4. Using the tools

The introduction to this book made mention of your business moving into the Internet age of computing and suggested that DB2® could help you achieve the transition. The previous chapters have prepared you for this trip:

- Chapter 1 taught you the DB2 language
- Chapter 2 instructed you on navigating through the DB2 GUI landscape
- Chapter 3 introduced you to the various DB2 GUI destinations

This chapter takes that knowledge and guides you through some real-world tasks that can be part of your company's ongoing business evolution.

These tasks are not presented as detailed step-by-step tutorials, but, rather they are designed to encourage you to explore the GUI as an active learner. Don't be afraid to try an item if it intrigues you. If you need assistance, you can use the comprehensive online help.

Backup and restore

A database can become unusable because of hardware or software failure (or both). You should have a backup and restore strategy in place to protect your database against the possibility of these failure situations.

Backup

The concept of a database *backup* is the same as any other data backup: taking a copy of the data and storing it on a different medium in case of failure or damage to the original.

A backup operation can be either *online* or *offline*. If the backup operation is online, other applications or processes can connect to the database, as well as read and modify data while the backup operation is running. If the backup operation is running offline, other applications *cannot* connect to the database.

DB2 uses logs to keep records of database changes. If a database needs to be restored to a point beyond the last full, offline backup, logs are required to roll the data forward to the point of failure.

There are two types of DB2 logging: *circular* and *archive*. Each provides a different level of recovery capability:

- *Circular* logging is the default behavior when a new database is created. (The *logretain* and *userexit* database configuration parameters are set to NO.) With this type of logging, only full, offline backups of the database are allowed. As the name suggests, circular logging uses a "ring" of online logs

to record database changes. The logs are used and retained only to the point of ensuring the integrity of current transactions. Circular logging does not allow you to roll a database forward through transactions performed after the last full backup operation. All changes occurring since the last backup operation are lost.

- *Archive* logging is used specifically for rollforward recovery. Enabling the *logretain* and/or the *userexit* database configuration parameter will result in archive logging. The advantage of choosing archive logging is that rollforward recovery can rebuild a database either to the end of the logs, or to a specific point in time. The archived log files can be used to recover changes made after the last backup was taken. The logs can also be used to undo a problem that occurred between the time of the last backup and the end of the logs.

Online backups are only supported if the database is configured for archive logging. During an online backup operation, all activities against the database are logged. When an online backup image is restored, the logs must be rolled forward at least to the point in time at which the backup operation completed. For this to happen, the logs must be archived and made available when the database is restored. After an online backup is complete, DB2 forces the currently active log to close and, as a result, it will be archived. This ensures that your online backup has a complete set of archived logs available for recovery.

Restore

Restore rebuilds a damaged or corrupted database that has been backed up. The restored database is in the same state it was in when the backup was made.

When restoring to an *existing* database, you should not be connected to the database that is to be restored: restore automatically establishes a connection to the specified database, and this connection is terminated at the completion of the restore operation. When restoring to a *new* database, an instance attachment is required to create the database. When restoring to a *new remote* database, you must first attach to the instance where the new database will reside. Then, create the new database.

A database restore requires an exclusive connection: that is, no applications can be running against the database when the operation starts, and the restore prevents other applications from accessing the database until the restore operation completes successfully.

If, at the time of the backup operation, the database was enabled for rollforward recovery (archive logging), the database can be brought to the state it was in prior to the occurrence of the damage or corruption by invoking the rollforward utility after successful completion of a restore.

Backing up a database online for the first time

Use this scenario to learn how to perform an online backup of a database that has never been backed up or, previously, has only been backed up offline.

Business problem

In your company, full backups of your operational database are performed weekly, on Sundays. The backups have always taken place offline and overnight, a time when database usage is low.

As the database has grown, the time required for the full backup has increased steadily. Now, the backup process is taking longer than the low-usage window allows. You want to start performing these weekly full backups online rather than offline, so the database can be accessed while the backup is taking place. You also run a monthly full offline backup for an extra layer of safety.

Before you begin: You should estimate the storage space and time requirements for the backup operation, ensuring that the task can be carried out within those constraints. This type of work would be done outside of the DB2 environment.

Step 1: Setting the logging type

By default, databases are created with circular logging. You must change this to archive logging.

When you change to archive logging you may also need to change the storage device being used for your logs because, in the long term, archival logs will require more space than circular logs.

You should schedule your logging switch to occur overnight, because this switch will initiate an offline database backup. You should also inform users that they will not have access to the database during the offline backup.

This step can be completed in the Control Center using the **Configure Database Logging Wizard** (Configure Database Logging is an option available for each database).

Step 2: Creating an online backup job

After the logging type has been changed to archive and the required offline backup occurs, you can schedule your first online backup.

This step can be completed in the Control Center using the **Backup Wizard** (Backup is an option available for each database).

Step 3: Backing up the database

After you have scheduled your online backup, you can view the task execution order in the Task Center. The tasks associated with each completion of the **Backup Wizard** will have the same prefix. The prefix consists of the wizard name and a task creation timestamp (for example, Backup - Mar5, 2002 3:49:28 PM).

Database administrators should be notified that Create, Drop, and Alter Table activities cannot be performed during the online backup.

The results of the backup can be checked in the Task Center. You can use this information to refine future weekly backups.

Related concepts:

- “Developing a Backup and Recovery Strategy” in the *Data Recovery and High Availability Guide and Reference*
- “Control Center” on page 22
- “Task Center” on page 37

Related tasks:

- “Restoring a database” on page 64

Restoring a database

Use this scenario to understand how to perform a full database restore as a reaction to a storage device failure.

Business problem

As in the backup scenario, full backups of your systems are performed weekly, on Sundays. On Wednesday, a hard-disk on one of the storage devices used for your warehouse system fails. The hardware staff replaces the disk and you decide to do a full database restore from the backup that was made on Sunday. Since new data is only loaded into the warehouse database on Fridays, it will not be necessary to perform a rollforward recovery.

Before you begin: You should notify users that there is a problem with the warehouse database and that there will be no access to the database while the restore is running. Because you require an exclusive connection to the database, you should also inform them that any applications currently running against the database will be terminated.

Step 1: Creating a restore job

You should perform a restore to an existing database. This will:

- Delete table, index, and long field data from the existing database, and replace it with data from the backup image.
- Replace table entries for each table space being restored.

- Retain the recovery history file, unless it is damaged or has no entries. If the recovery history file is damaged, the database manager copies the file from the backup image.
- Retain the authentication type for the existing database.
- Retain the database directories for the existing database. The directories define where the database resides, and how it is cataloged.

You can have the database returned to the active state immediately after the restore because no rollforward activities are required.

To reduce the amount of time required to complete a restore operation, you might consider:

- Increasing the restore buffer size
- Increasing the number of buffers

See the *Data Recovery and High Availability Guide and Reference* for more information.

This step can be completed in the Control Center using the **Restore Database Wizard** (Restore is an option available for each database).

Step 2: Restoring the database

You can run the restore immediately or schedule it to run at a specific time.

The results of the restore can be checked in the Task Center. You can use this information to refine future restores.

You should test the newly restored database before informing users that the database is back online.

Related concepts:

- “Developing a Backup and Recovery Strategy” in the *Data Recovery and High Availability Guide and Reference*
- “Control Center” on page 22
- “Task Center” on page 37

Related tasks:

- “Backing up a database online for the first time” on page 63

Database health

Built into DB2 is the ability to collect data about its operation and performance. Health indicators measure changes in the state of database objects and resources. Monitoring these changes determines whether the state is healthy or unhealthy. If there is a distinct change in a state, warnings and alerts can be generated advising you of potential problems and providing recommendations for resolving them.

Reacting to alerts

Use this scenario to understand how DB2 will alert you to potential problems with the health of your database and how you can access its recommendations for remedying such situations.

Business problem

In your company, there are applications that use the FINANCE database in the DB2 instance. While working with this database you are alerted to problems with the amount of memory being used by sort operations.

Before you begin: You must ensure that the appropriate health indicators are enabled for monitoring sort activity. In this case the **Database Object Health Indicator Settings** for the FINANCE database requires that **Shared Sort Memory Utilization** indicator is enabled.

Step 1: Setting the actions for alerts

For actions to occur when an alert condition is met, you must define those actions. You decide that two actions will take place when an alert occurs:

1. E-mail notifications are sent to a group of selected contacts.
You use the Health Center Configure Notifications dialog to create a list of contacts who are notified when an alert is generated for an instance. The notification identifies the alert, provides information on how to get a description of the alert, and includes recommendations that can help resolve the problem. Each instance can have a different set of contacts.
2. A graphical indicator is displayed to indicate that an alert has been issued.
You use the Tool Settings notebook to specify that Health Center Status Beacon will appear on the status line for each DB2 Center whenever an alert occurs. Users should use this as a visual cue that they should view the alert details in the Health Center.

Step 2: Working with e-mail notifications

You are working with an application that generates employee reports from the FINANCE database in the DB2 instance. Currently, you do not have any of the DB2 GUI tools open. The employee report application generates a number of SQL queries that include ORDER BY clauses. While the reports are being generated, you receive an e-mail generated by the Health Center indicating that there is a limit on the total amount of shared memory used for sorting.

The e-mail also includes instructions on accessing the GUI tools and opening the Health Center. You locate the DB2 instance and FINANCE database in the object tree and review the details of the alert and the suggested recommendations. You then decide on a course of action based on the details and recommendations.

Step 3: Working with status beacon notifications

You are developing a report generating application which uses the FINANCE database in the DB2 instance. You are working with the Command Center, testing the application (which is running SQL queries for the reports) and trying to configure the database for optimal performance. You notice that Health Center Status Beacon has appeared on the Command Center status line, indicating that there is a warning or alert in the Health Center. You open the Health Center and locate the DB2 instance and FINANCE database in the object tree and review the details of the warning and the suggested recommendations. You then decide on a course of action based on the details and recommendations.

Related concepts:

- “Command Center” on page 35
- “Health Center” on page 40

Replication

The replication environment that you need depends on when you want data updated and how you want transactions handled.

Replicating data

Use this scenario to gain some experience using the DB2 Replication Center and its Capture and Apply programs. This scenario copies changes from a DB2 replication source to a target table in a database on DB2 for Windows Enterprise Server Edition (ESE) or Workgroup Server Edition (WSE).

Business problem

In your company, there is a group that uses an application that generates reports. This application needs information that exists in the DEPARTMENT table of the SAMPLE database. In order to protect the integrity of the data in the DEPARTMENT table (source table), you want to copy its data to a target table that can be read only by the report-generating application. For ease of administration, you want to keep the target table on the same computer as the source table.

You require a simple data-distribution configuration, in which changes from a replication source are being replicated to a single read-only copy.

Before you begin: The steps in this scenario use the data in the DEPARTMENT table from the SAMPLE database. The fully qualified name is *schema*.DEPARTMENT; where *schema* is the user ID that created the table.

If you want to follow this scenario on your computer, you must:

1. Make sure that you have created the default DB2 instance. This scenario assumes that all databases are within the same instance.
2. Make sure that you have access to the SAMPLE database. This database will be both the source server and the capture control server for this scenario.
3. Use the DB2 Control Center to create a new database called COPYDB, which you will use as the target server and the apply control server. Both the name and the alias for the database should be COPYDB.

For the remainder of this scenario, use the user ID with which you created the SAMPLE and COPYDB databases. Because this ID created the databases, you have the required authority (DBADM or SYSADM) to perform replication tasks.

Step 1: Planning your replication strategy

Before you perform any replication tasks, you must design and plan your replication strategy. You are then ready to set up your replication environment.

You already know that the replication source is the *schema*.DEPARTMENT table in the SAMPLE database. Before you set up your environment, you must decide what you want to replicate from that table; you decide to register all columns and subscribe to all columns. Registration defines what portions of a source are available for replication. Subscription identifies what parts are replicated in a target.

You decide that you want your replication target to be the COPYDB database. Currently, the table to which you want to replicate data does not exist in that database; you want the Replication Center to create the target table according to your specifications. This method of automatically generating a target table is preferred because it ensures correct mapping of the columns in the source table to the columns being replicated in the target table. You can use existing target tables instead, but this scenario assumes that the target table does not exist.

Because the columns in the target table simply reflect the data from the source table, and because there will be only one row in the target table for each row in the source table, you can use a *user copy* type of target table.

For the purpose of this scenario, you decide to store the change data (CD) table, the target table, and the replication control tables in their respective

default table spaces. Although the SAMPLE and COPYDB databases exist in the same computer, their table spaces are in separate containers.

Note: Change data tables are used to record all committed changes made to a replication source.

Step 2: Creating replication control tables for the Capture program

The Capture program reads the replication control tables for current registration information and stores this status in these tables. It also reads the DB2 log for changes to the registered columns in the registered tables. Any database that acts as a Capture control server must contain the Capture control tables.

Create Capture Control Tables is an option available from the **Capture Control Servers** folder under the **Replication Definitions** folder in the Replication Center.

Step 3: Enabling the source database for replication

The Capture program reads the DB2 log for log records that include changes to registered tables. The log must be an archive log so that the log file will not be reused by DB2 before the Capture program can read the log. For UNIX and Windows environments, the DB2 default is circular logging, so you must change this setting to archive logging.

Changing to archive logging enables rollforward database recovery. This change also requires that you back up your database before proceeding.

Enable Database for Replication is an action that can be performed on the SAMPLE database in the **Capture Control Servers** folder.

Step 4: Registering a replication source

After you create the Capture control tables and enable the database for replication, register the DEPARTMENT table as a replication source. Remember that you are registering all columns for replication.

Select one of the **Capture changes** options for the Row-capture rule. This means that only changed data is replicated to the target table (change-capture replication).

Register Tables is an option available from the **Registered Tables** folder under SAMPLE database, which now appears in the **Capture Control Servers** folder (you may need to refresh the view).

When you defined the DEPARTMENT table as a replication source, the Replication Center created the CD table and CD-table index for this replication source, and it updated the Capture control tables.

Step 5: Creating replication control tables for the Apply program

The Apply program reads the replication control tables for current subscription-set information and stores its status and the status of the subscription sets in these tables. Any database that acts as an Apply control server must contain the Apply control tables.

Create Apply Control Tables is an option available from the **Apply Control Servers** folder under the **Replication Definitions** folder.

Step 6: Creating a subscription set and a subscription-set member

After you register the source table, you need to create a subscription set. A subscription set defines a relationship between the replication source database (SAMPLE) and a target database (COPYDB). A subscription-set member defines a relationship between the replication source table (DEPARTMENT) and a target table (call it DEPTCOPY). The subscription set ensures that a group of subscription-set members is processed at the same time. If the processing of one of the members fails, the Apply program rolls back any changes made while processing the other members.

Create is an option available from the **Subscription Sets** folder under COPYDB database in the **Apply Control Servers** folder of **Replication Definitions**.

Step 7: Creating an Apply password file

Because the Apply program needs to connect to the Capture control server, the Apply control server, and the target server, you must create a password file for user authentication. Because the contents of the password file are encrypted, only the Apply program can read the file, although you can modify the file using the **anspwd** command.

You should create the password file in the directory in which you will start the Apply program. When you start the Apply program, you specify the file name for the password file and the value for the directory in which the Apply program will store its log and work files. One of the Apply program's work files is the password file.

Step 8: Replicating data

After you register the replication source and create the subscription set, start the Capture and Apply programs to perform the initial full refresh for the target table and begin change-capture replication.

You can use the Replication Center to view the results of your replication by comparing the contents of the target table with its source table.

Related concepts:

- "Instances" on page 3

- “Table spaces” on page 5
- “Authorities” on page 9

Related tasks:

- “Backing up a database online for the first time” on page 63
- “Setting up for replication” in the *Replication Guide and Reference*

Related reference:

- “asnpwd: Maintaining password files (UNIX and Windows)” in the *Replication Guide and Reference*

Appendix A. DB2 Universal Database technical information

Overview of DB2 Universal Database technical information

DB2 Universal Database technical information can be obtained in the following formats:

- Books (PDF and hard-copy formats)
- A topic tree (HTML format)
- Help for DB2 tools (HTML format)
- Sample programs (HTML format)
- Command line help
- Tutorials

This section is an overview of the technical information that is provided and how you can access it.

FixPaks for DB2 documentation

IBM may periodically make documentation FixPaks available. Documentation FixPaks allow you to update the information that you installed from the *DB2 HTML Documentation CD* as new information becomes available.

Note: If you do install documentation FixPaks, your HTML documentation will contain more recent information than either the DB2 printed or online PDF manuals.

Categories of DB2 technical information

The DB2 technical information is categorized by the following headings:

- Core DB2 information
- Administration information
- Application development information
- Business intelligence information
- DB2 Connect information
- Getting started information
- Tutorial information
- Optional component information
- Release notes

The following tables describe, for each book in the DB2 library, the information needed to order the hard copy, print or view the PDF, or locate the HTML directory for that book. A full description of each of the books in

the DB2 library is available from the IBM Publications Center at www.ibm.com/shop/publications/order

The installation directory for the HTML documentation CD differs for each category of information:

htmlcdpath/doc/htmlcd/%L/category

where:

- *htmlcdpath* is the directory where the HTML CD is installed.
- *%L* is the language identifier. For example, en_US.
- *category* is the category identifier. For example, core for the core DB2 information.

In the PDF file name column in the following tables, the character in the sixth position of the file name indicates the language version of a book. For example, the file name db2d1e80 identifies the English version of the *Administration Guide: Planning* and the file name db2d1g80 identifies the German version of the same book. The following letters are used in the sixth position of the file name to indicate the language version:

Language	Identifier
Arabic	w
Brazilian Portuguese	b
Bulgarian	u
Croatian	9
Czech	x
Danish	d
Dutch	q
English	e
Finnish	y
French	f
German	g
Greek	a
Hungarian	h
Italian	i
Japanese	j
Korean	k
Norwegian	n
Polish	p
Portuguese	v
Romanian	8
Russian	r
Simp. Chinese	c
Slovakian	7
Slovenian	l

Spanish	z
Swedish	s
Trad. Chinese	t
Turkish	m

No form number indicates that the book is only available online and does not have a printed version.

Core DB2 information

The information in this category covers DB2 topics that are fundamental to all DB2 users. You will find the information in this category useful whether you are a programmer, a database administrator, or you work with DB2 Connect, DB2 Warehouse Manager, or other DB2 products.

The installation directory for this category is `doc/htmlcd/%L/core`.

Table 1. Core DB2 information

Name	Form Number	PDF File Name
<i>IBM DB2 Universal Database Command Reference</i>	SC09-4828	db2n0x80
<i>IBM DB2 Universal Database Glossary</i>	No form number	db2t0x80
<i>IBM DB2 Universal Database Master Index</i>	SC09-4839	db2w0x80
<i>IBM DB2 Universal Database Message Reference, Volume 1</i>	GC09-4840	db2m1x80
<i>IBM DB2 Universal Database Message Reference, Volume 2</i>	GC09-4841	db2m2x80
<i>IBM DB2 Universal Database What's New</i>	SC09-4848	db2q0x80

Administration information

The information in this category covers those topics required to effectively design, implement, and maintain DB2 databases, data warehouses, and federated systems.

The installation directory for this category is doc/htmlcd/%L/admin.

Table 2. Administration information

Name	Form number	PDF file name
<i>IBM DB2 Universal Database Administration Guide: Planning</i>	SC09-4822	db2d1x80
<i>IBM DB2 Universal Database Administration Guide: Implementation</i>	SC09-4820	db2d2x80
<i>IBM DB2 Universal Database Administration Guide: Performance</i>	SC09-4821	db2d3x80
<i>IBM DB2 Universal Database Administrative API Reference</i>	SC09-4824	db2b0x80
<i>IBM DB2 Universal Database Data Movement Utilities Guide and Reference</i>	SC09-4830	db2dmx80
<i>IBM DB2 Universal Database Data Recovery and High Availability Guide and Reference</i>	SC09-4831	db2hax80
<i>IBM DB2 Universal Database Data Warehouse Center Administration Guide</i>	SC27-1123	db2ddx80
<i>IBM DB2 Universal Database Federated Systems Guide</i>	GC27-1224	db2fpx80
<i>IBM DB2 Universal Database Guide to GUI Tools for Administration and Development</i>	SC09-4851	db2atx80
<i>IBM DB2 Universal Database Replication Guide and Reference</i>	SC27-1121	db2e0x80
<i>IBM DB2 Installing and Administering a Satellite Environment</i>	GC09-4823	db2dsx80
<i>IBM DB2 Universal Database SQL Reference, Volume 1</i>	SC09-4844	db2s1x80
<i>IBM DB2 Universal Database SQL Reference, Volume 2</i>	SC09-4845	db2s2x80
<i>IBM DB2 Universal Database System Monitor Guide and Reference</i>	SC09-4847	db2f0x80

Application development information

The information in this category is of special interest to application developers or programmers working with DB2. You will find information about supported languages and compilers, as well as the documentation required to access DB2 using the various supported programming interfaces, such as embedded SQL, ODBC, JDBC, SQLj, and CLI. If you view this information online in HTML you can also access a set of DB2 sample programs in HTML.

The installation directory for this category is `doc/htmlcd/%L/ad`.

Table 3. Application development information

Name	Form number	PDF file name
<i>IBM DB2 Universal Database Application Development Guide: Building and Running Applications</i>	SC09-4825	db2axx80
<i>IBM DB2 Universal Database Application Development Guide: Programming Client Applications</i>	SC09-4826	db2a1x80
<i>IBM DB2 Universal Database Application Development Guide: Programming Server Applications</i>	SC09-4827	db2a2x80
<i>IBM DB2 Universal Database Call Level Interface Guide and Reference, Volume 1</i>	SC09-4849	db2l1x80
<i>IBM DB2 Universal Database Call Level Interface Guide and Reference, Volume 2</i>	SC09-4850	db2l2x80
<i>IBM DB2 Universal Database Data Warehouse Center Application Integration Guide</i>	SC27-1124	db2adx80
<i>IBM DB2 XML Extender Administration and Programming</i>	SC27-1234	db2sxx80

Business intelligence information

The information in this category describes how to use components that enhance the data warehousing and analytical capabilities of DB2 Universal Database.

The installation directory for this category is doc/htmlcd/%L/wareh.

Table 4. Business intelligence information

Name	Form number	PDF file name
<i>IBM DB2 Warehouse Manager Information Catalog Center Administration Guide</i>	SC27-1125	db2dix80
<i>IBM DB2 Warehouse Manager Installation Guide</i>	GC27-1122	db2idx80

DB2 Connect information

The information in this category describes how to access host or iSeries data using DB2 Connect Enterprise Edition or DB2 Connect Personal Edition.

The installation directory for this category is doc/htmlcd/%L/conn.

Table 5. DB2 Connect information

Name	Form number	PDF file name
<i>APPC, CPI-C, and SNA Sense Codes</i>	No form number	db2apx80
<i>IBM Connectivity Supplement</i>	No form number	db2h1x80
<i>IBM DB2 Connect Quick Beginnings for DB2 Connect Enterprise Edition</i>	GC09-4833	db2c6x80
<i>IBM DB2 Connect Quick Beginnings for DB2 Connect Personal Edition</i>	GC09-4834	db2c1x80
<i>IBM DB2 Connect User's Guide</i>	SC09-4835	db2c0x80

Getting started information

The information in this category is useful when you are installing and configuring servers, clients, and other DB2 products.

The installation directory for this category is doc/htmlcd/%L/start.

Table 6. Getting started information

Name	Form number	PDF file name
<i>IBM DB2 Universal Database Quick Beginnings for DB2 Clients</i>	GC09-4832	db2itx80

Table 6. Getting started information (continued)

Name	Form number	PDF file name
<i>IBM DB2 Universal Database Quick Beginnings for DB2 Servers</i>	GC09-4836	db2isx80
<i>IBM DB2 Universal Database Quick Beginnings for DB2 Personal Edition</i>	GC09-4838	db2i1x80
<i>IBM DB2 Universal Database Installation and Configuration Supplement</i>	GC09-4837	db2iyx80
<i>IBM DB2 Universal Database Quick Beginnings for DB2 Data Links Manager</i>	GC09-4829	db2z6x80

Tutorial information

Tutorial information introduces DB2 features and teaches how to perform various tasks.

The installation directory for this category is `doc/htmlcd/%L/tutr`.

Table 7. Tutorial information

Name	Form number	PDF file name
<i>Business Intelligence Tutorial: Introduction to the Data Warehouse</i>	No form number	db2tux80
<i>Business Intelligence Tutorial: Extended Lessons in Data Warehousing</i>	No form number	db2tax80
<i>Development Center Tutorial for Video Online using Microsoft Visual Basic</i>	No form number	db2tdx80
<i>Information Catalog Center Tutorial</i>	No form number	db2aix80
<i>Video Central for e-business Tutorial</i>	No form number	db2twx80
<i>Visual Explain Tutorial</i>	No form number	db2tvx80

Optional component information

The information in this category describes how to work with optional DB2 components.

The installation directory for this category is doc/htmlcd/%L/opt.

Table 8. Optional component information

Name	Form number	PDF file name
<i>IBM DB2 Life Sciences Data Connect Planning, Installation, and Configuration Guide</i>	GC27-1235	db2lsx80
<i>IBM DB2 Spatial Extender User's Guide and Reference</i>	SC27-1226	db2sbx80
<i>IBM DB2 Universal Database Data Links Manager Administration Guide and Reference</i>	SC27-1221	db2z0x80
<i>IBM DB2 Universal Database Net Search Extender Administration and Programming Guide</i>	SH12-6740	N/A

Note: HTML for this document is not installed from the HTML documentation CD.

Release notes

The release notes provide additional information specific to your product's release and FixPak level. They also provides summaries of the documentation updates incorporated in each release and FixPak.

Table 9. Release notes

Name	Form number	PDF file name
<i>DB2 Release Notes</i>	See note.	See note.
<i>DB2 Installation Notes</i>	Available on product CD-ROM only.	Available on product CD-ROM only.

Note: The HTML version of the release notes is available from the Information Center and on the product CD-ROMs. To view the ASCII file on UNIX-based platforms, see the Release.Notes file. This file is located in the DB2DIR/Readme/%L directory, where %L represents the locale name and DB2DIR represents:

- /usr/opt/db2_08_01 on AIX
- /opt/IBM/db2/V8.1 on all other UNIX operating systems

Related tasks:

- "Printing DB2 books from PDF files" on page 81

- “Ordering printed DB2 books” on page 82
- “Accessing online help” on page 82
- “Finding product information by accessing the DB2 Information Center from the administration tools” on page 86
- “Viewing technical documentation online directly from the DB2 HTML Documentation CD” on page 87

Printing DB2 books from PDF files

You can print DB2 books from the PDF files on the *DB2 PDF Documentation* CD. Using Adobe Acrobat Reader, you can print either the entire book or a specific range of pages.

Prerequisites:

Ensure that you have Adobe Acrobat Reader. It is available from the Adobe Web site at www.adobe.com

Procedure:

To print a DB2 book from a PDF file:

1. Insert the *DB2 PDF Documentation* CD. On UNIX operating systems, mount the DB2 PDF Documentation CD. Refer to your *Quick Beginnings* book for details on how to mount a CD on UNIX operating systems.
2. Start Adobe Acrobat Reader.
3. Open the PDF file from one of the following locations:
 - On Windows operating systems:
x:\doc\language directory, where *x* represents the CD-ROM drive letter and *language* represents the two-character territory code that represents your language (for example, EN for English).
 - On UNIX operating systems:
/cdrom/doc/%L directory on the CD-ROM, where */cdrom* represents the mount point of the CD-ROM and *%L* represents the name of the desired locale.

Related tasks:

- “Ordering printed DB2 books” on page 82
- “Finding product information by accessing the DB2 Information Center from the administration tools” on page 86
- “Viewing technical documentation online directly from the DB2 HTML Documentation CD” on page 87

Related reference:

- “Overview of DB2 Universal Database technical information” on page 73

Ordering printed DB2 books

Procedure:

To order printed books:

- Contact your IBM authorized dealer or marketing representative. To find a local IBM representative, check the IBM Worldwide Directory of Contacts at www.ibm.com/planetwide
- Phone 1-800-879-2755 in the United States or 1-800-IBM-4YOU in Canada.
- Visit the IBM Publications Center at www.ibm.com/shop/publications/order

You can also obtain printed DB2 manuals by ordering Doc Packs for your DB2 product from your IBM Reseller. The Doc Packs are subsets of the manuals in the DB2 library selected to help you to get started using the DB2 product that you purchased. The manuals in the Doc Packs are the same as those that are available in PDF format on the *DB2 PDF Documentation CD* and contain the same content as the documentation that is available on the *DB2 HTML Documentation CD*.

Related tasks:

- “Printing DB2 books from PDF files” on page 81
- “Finding topics by accessing the DB2 Information Center from a browser” on page 84
- “Viewing technical documentation online directly from the DB2 HTML Documentation CD” on page 87

Related reference:

- “Overview of DB2 Universal Database technical information” on page 73

Accessing online help

The online help that comes with all DB2 components is available in three types:

- Window and notebook help
- Command line help
- SQL statement help

Window and notebook help explain the tasks that you can perform in a window or notebook and describe the controls. This help has two types:

- Help accessible from the **Help** button
- Infopops

The **Help** button gives you access to overview and prerequisite information. The infopops describe the controls in the window or notebook. Window and notebook help are available from DB2 centers and components that have user interfaces.

Command line help includes Command help and Message help. Command help explains the syntax of commands in the command line processor. Message help describes the cause of an error message and describes any action you should take in response to the error.

SQL statement help includes SQL help and SQLSTATE help. DB2 returns an SQLSTATE value for conditions that could be the result of an SQL statement. SQLSTATE help explains the syntax of SQL statements (SQL states and class codes).

Note: SQL help is not available for UNIX operating systems.

Procedure:

To access online help:

- For window and notebook help, click **Help** or click that control, then click **F1**. If the **Automatically display infopops** check box on the **General** page of the **Tool Settings** notebook is selected, you can also see the infopop for a particular control by holding the mouse cursor over the control.
- For command line help, open the command line processor and enter:
 - For Command help:
 - `? command`

where *command* represents a keyword or the entire command.

For example, `? catalog` displays help for all the CATALOG commands, while `? catalog database` displays help for the CATALOG DATABASE command.

- For Message help:
 - `? XXXnnnnn`

where *XXXnnnnn* represents a valid message identifier.

For example, `? SQL30081` displays help about the SQL30081 message.

- For SQL statement help, open the command line processor and enter:
 - `? sqlstate` or `? class code`

where *sqlstate* represents a valid five-digit SQL state and *class code* represents the first two digits of the SQL state.

For example, ? 08003 displays help for the 08003 SQL state, while ? 08 displays help for the 08 class code.

Related tasks:

- “Finding topics by accessing the DB2 Information Center from a browser” on page 84
- “Viewing technical documentation online directly from the DB2 HTML Documentation CD” on page 87

Finding topics by accessing the DB2 Information Center from a browser

The DB2 Information Center accessed from a browser enables you to access the information you need to take full advantage of DB2 Universal Database and DB2 Connect. The DB2 Information Center also documents major DB2 features and components including replication, data warehousing, metadata, and DB2 extenders.

The DB2 Information Center accessed from a browser is composed of the following major elements:

Navigation tree

The navigation tree is located in the left frame of the browser window. The tree expands and collapses to show and hide topics, the glossary, and the master index in the DB2 Information Center.

Navigation toolbar

The navigation toolbar is located in the top right frame of the browser window. The navigation toolbar contains buttons that enable you to search the DB2 Information Center, hide the navigation tree, and find the currently displayed topic in the navigation tree.

Content frame

The content frame is located in the bottom right frame of the browser window. The content frame displays topics from the DB2 Information Center when you click on a link in the navigation tree, click on a search result, or follow a link from another topic or from the master index.

Prerequisites:

To access the DB2 Information Center from a browser, you must use one of the following browsers:

- Microsoft Explorer, version 5 or later

- Netscape Navigator, version 6.1 or later

Restrictions:

The DB2 Information Center contains only those sets of topics that you chose to install from the *DB2 HTML Documentation CD*. If your Web browser returns a File not found error when you try to follow a link to a topic, you must install one or more additional sets of topics from the *DB2 HTML Documentation CD*.

Procedure:

To find a topic by searching with keywords:

1. In the navigation toolbar, click **Search**.
2. In the top text entry field of the Search window, enter one or more terms related to your area of interest and click **Search**. A list of topics ranked by accuracy displays in the **Results** field. The numerical ranking beside the hit provides an indication of the strength of the match (bigger numbers indicate stronger matches).

Entering more terms increases the precision of your query while reducing the number of topics returned from your query.

3. In the **Results** field, click the title of the topic you want to read. The topic displays in the content frame.

To find a topic in the navigation tree:

1. In the navigation tree, click the book icon of the category of topics related to your area of interest. A list of subcategories displays underneath the icon.
2. Continue to click the book icons until you find the category containing the topics in which you are interested. Categories that link to topics display the category title as an underscored link when you move the cursor over the category title. The navigation tree identifies topics with a page icon.
3. Click the topic link. The topic displays in the content frame.

To find a topic or term in the master index:

1. In the navigation tree, click the "Index" category. The category expands to display a list of links arranged in alphabetical order in the navigation tree.
2. In the navigation tree, click the link corresponding to the first character of the term relating to the topic in which you are interested. A list of terms with that initial character displays in the content frame. Terms that have multiple index entries are identified by a book icon.

3. Click the book icon corresponding to the term in which you are interested. A list of subterms and topics displays below the term you clicked. Topics are identified by page icons with an underscored title.
4. Click on the title of the topic that meets your needs. The topic displays in the content frame.

Related concepts:

- “Accessibility” on page 93
- “DB2 Information Center accessed from a browser” on page 96

Related tasks:

- “Finding product information by accessing the DB2 Information Center from the administration tools” on page 86
- “Updating the HTML documentation installed on your machine” on page 88
- “Troubleshooting DB2 documentation search with Netscape 4.x” on page 91
- “Searching the DB2 documentation” on page 92

Related reference:

- “Overview of DB2 Universal Database technical information” on page 73

Finding product information by accessing the DB2 Information Center from the administration tools

The DB2 Information Center provides quick access to DB2 product information and is available on all operating systems for which the DB2 administration tools are available.

The DB2 Information Center accessed from the tools provides six types of information.

Tasks Key tasks you can perform using DB2.

Concepts

Key concepts for DB2.

Reference

DB2 reference information, such as keywords, commands, and APIs.

Troubleshooting

Error messages and information to help you with common DB2 problems.

Samples

Links to HTML listings of the sample programs provided with DB2.

Tutorials

Instructional aid designed to help you learn a DB2 feature.

Prerequisites:

Some links in the DB2 Information Center point to Web sites on the Internet. To display the content for these links, you will first have to connect to the Internet.

Procedure:

To find product information by accessing the DB2 Information Center from the tools:

1. Start the DB2 Information Center in one of the following ways:
 - From the graphical administration tools, click on the **Information Center** icon in the toolbar. You can also select it from the **Help** menu.
 - At the command line, enter **db2ic**.
2. Click the tab of the information type related to the information you are attempting to find.
3. Navigate through the tree and click on the topic in which you are interested. The Information Center will then launch a Web browser to display the information.

Related concepts:

- “Accessibility” on page 93
- “DB2 Information Center accessed from a browser” on page 96

Related tasks:

- “Finding topics by accessing the DB2 Information Center from a browser” on page 84
- “Searching the DB2 documentation” on page 92

Viewing technical documentation online directly from the DB2 HTML Documentation CD

All of the HTML topics that you can install from the *DB2 HTML Documentation CD* can also be read directly from the CD. Therefore, you can view the documentation without having to install it.

Restrictions:

As the Tools help is installed from the DB2 product CD and not from the *DB2 HTML Documentation CD*, you must install the DB2 product to view the help.

Procedure:

1. Insert the *DB2 HTML Documentation CD*. On UNIX operating systems, mount the *DB2 HTML Documentation CD*. Refer to your *Quick Beginnings* book for details on how to mount a CD on UNIX operating systems.
2. Start your HTML browser and open the appropriate file:

- For Windows operating systems:

```
e:\program files\IBM\SQLLIB\doc\htmlcd\%L\index.htm
```

where *e* represents the CD-ROM drive, and %L is the locale of the documentation that you wish to use, for example, **en_US** for English.

- For UNIX operating systems:

```
/cdrom/program files/IBM/SQLLIB/doc/htmlcd/%L/index.htm
```

where */cdrom/* represents where the CD is mounted, and %L is the locale of the documentation that you wish to use, for example, **en_US** for English.

Related tasks:

- “Finding topics by accessing the DB2 Information Center from a browser” on page 84
- “Copying files from the DB2 HTML Documentation CD to a Web server” on page 90

Related reference:

- “Overview of DB2 Universal Database technical information” on page 73

Updating the HTML documentation installed on your machine

It is now possible to update the HTML installed from the *DB2 HTML Documentation CD* when updates are made available from IBM. This can be done in one of two ways:

- Using the Information Center (if you have the DB2 administration GUI tools installed).
- By downloading and applying a DB2 HTML documentation FixPak .

Note: This will NOT update the DB2 code; it will only update the HTML documentation installed from the *DB2 HTML Documentation CD*.

Restrictions:

To run the HTML Updater for all UNIX platforms, the user must have write authority to the directory where the *DB2 HTML Documentation CD* is installed. The following should be run as root from the command prompt:

- Using the Bourne shell or any derivative:
(`~db2inst1/sqlllib/db2profile;db2ic`)

where `db2inst1` can be any owner ID.

- Using the C shell:
(`source~db2inst1/sqlllib/db2profile;db2ic`)

where `db2inst1` can be any owner ID.

Procedure:

To use the Information Center to update your local documentation:

1. Start the DB2 Information Center in one of the following ways:
 - From the graphical administration tools, click on the **Information Center** icon in the toolbar. You can also select it from the **Help** menu.
 - At the command line, enter **db2ic**.
2. Ensure your machine has access to the external Internet; the updater will download the latest documentation FixPak from the IBM server if required.
3. Select **Information Center** —> **Update Local Documentation** from the menu to start the update.
4. Supply your proxy information (if required) to connect to the external Internet.

The Information Center will download and apply the latest documentation FixPak, if one is available.

To manually download and apply the documentation FixPak :

1. Ensure your machine is connected to the Internet.
2. Open the DB2 support page in your Web browser at:
www.ibm.com/software/data/db2/udb/winos2unix/support.
3. Follow the link for Version 8 and look for the "Documentation FixPaks" link.
4. Determine if the version of your local documentation is out of date by comparing the documentation FixPak level to the documentation level you have installed. This current documentation on your machine is at the following level: **DB2 v8.1 GA**.
5. If there is a more recent version of the documentation available then download the FixPak applicable to your operating system. There is one FixPak for all Windows platforms, and one FixPak for all UNIX platforms.
6. Apply the FixPak:
 - For Windows operating systems: The documentation FixPak is a self extracting zip file. Place the downloaded documentation FixPak in an

empty directory, and run it. It will create a **setup** command which you can run to install the documentation FixPak.

- For UNIX operating systems: The documentation FixPak is a compressed tar.Z file. Uncompress and untar the file. It will create a directory named `delta_install` with a script called **installdocfix**. Run this script to install the documentation FixPak.

Related tasks:

- “Copying files from the DB2 HTML Documentation CD to a Web server” on page 90

Related reference:

- “Overview of DB2 Universal Database technical information” on page 73

Copying files from the DB2 HTML Documentation CD to a Web server

The entire DB2 information library is delivered to you on the *DB2 HTML Documentation CD* and may be installed on a Web server for easier access. Simply copy to your Web server the documentation for the languages that you want.

Note: You might encounter slow performance if you access the HTML documentation from a Web server through a low-speed connection.

Procedure:

To copy files from the *DB2 HTML Documentation CD* to a Web server, use the appropriate source path:

- For Windows operating systems:

```
E:\program files\IBM\SQLLIB\doc\htmlcd\%L\*.*
```

where *E* represents the CD-ROM drive and *%L* represents the language identifier.

- For UNIX operating systems:

```
/cdrom/program files/IBM/SQLLIB/doc/htmlcd/%L/*.*
```

where *cdrom* represents the mount point for the CD-ROM drive and *%L* represents the language identifier.

Related tasks:

- “Searching the DB2 documentation” on page 92

Related reference:

- “Supported DB2 interface languages, locales, and code pages” in the *Quick Beginnings for DB2 Servers*
- “Overview of DB2 Universal Database technical information” on page 73

Troubleshooting DB2 documentation search with Netscape 4.x

Most search problems are related to the Java support provided by web browsers. This task describes possible workarounds.

Procedure:

A common problem with Netscape 4.x involves a missing or misplaced security class. Try the following workaround, especially if you see the following line in the browser Java console:

```
Cannot find class java/security/InvalidParameterException
```

- On Windows operating systems:

From the *DB2 HTML Documentation CD*, copy the supplied `x:program files\IBM\SQLLIB\doc\htmlcd\locale\InvalidParameterException.class` file to the `java\classes\java\security\` directory relative to your Netscape browser installation, where *x* represents the CD-ROM drive letter and *locale* represents the name of the desired locale.

Note: You may have to create the `java\security\` subdirectory structure.

- On UNIX operating systems:

From the *DB2 HTML Documentation CD*, copy the supplied `/cdrom/program files/IBM/SQLLIB/doc/htmlcd/locale/InvalidParameterException.class` file to the `java/classes/java/security/` directory relative to your Netscape browser installation, where *cdrom* represents the mount point of the CD-ROM and *locale* represents the name of the desired locale.

Note: You may have to create the `java/security/` subdirectory structure.

If your Netscape browser still fails to display the search input window, try the following:

- Stop all instances of Netscape browsers to ensure that there is no Netscape code running on the machine. Then open a new instance of the Netscape browser and try to start the search again.
- Purge the browser’s cache.
- Try a different version of Netscape, or a different browser.

Related tasks:

- “Searching the DB2 documentation” on page 92

Searching the DB2 documentation

You can search the library of DB2 documentation to locate information that you need. A pop-up search window opens when you click the search icon in the navigation toolbar of the DB2 Information Center (accessed from a browser). The search can take a minute to load, depending on the speed of your computer and network.

Prerequisites:

You need Netscape 6.1 or higher, or Microsoft's Internet Explorer 5 or higher. Ensure that your browser's Java support is enabled.

Restrictions:

The following restrictions apply when you use the documentation search:

- Search is not case sensitive.
- Boolean searches are not supported.
- Wildcard and partial searches are not supported. A search on *java** (or *java*) will only look for the literal string *java** (or *java*) and would not, for example, find *javadoc*.

Procedure:

To search the DB2 documentation:

1. In the navigation toolbar, click the **Search** icon.
2. In the top text entry field of the Search window, enter one or more terms (separated by a space) related to your area of interest and click **Search**. A list of topics ranked by accuracy displays in the **Results** field. The numerical ranking beside the hit provides an indication of the strength of the match (bigger numbers indicate stronger matches).
Entering more terms increases the precision of your query.
3. In the **Results** list, click the title of the topic you want to read. The topic displays in the content frame of the DB2 Information Center.

Note: When you perform a search, the first (highest-ranking) result is automatically loaded into your browser frame. To view the contents of other search results, click on the result in the results list.

Related tasks:

- "Troubleshooting DB2 documentation search with Netscape 4.x" on page 91

Online DB2 troubleshooting information

With the release of DB2[®] UDB Version 8, there will no longer be a *Troubleshooting Guide*. The troubleshooting information once contained in this guide has been integrated into the DB2 publications. By doing this, we are able to deliver the most up-to-date information possible. To find information on the troubleshooting utilities and functions of DB2, access the DB2 Information Center from any of the tools.

Refer to the DB2 Online Support site if you are experiencing problems and want help finding possible causes and solutions. The support site contains a large, constantly updated database of DB2 publications, TechNotes, APAR (product problem) records, FixPaks, and other resources. You can use the support site to search through this knowledge base and find possible solutions to your problems.

Access the Online Support site at www.ibm.com/software/data/db2/udb/winos2unix/support, or by clicking the **Online Support** button in the DB2 Information Center. Frequently changing information, such as the listing of internal DB2 error codes, is now also available from this site.

Related concepts:

- “DB2 Information Center accessed from a browser” on page 96

Related tasks:

- “Finding product information by accessing the DB2 Information Center from the administration tools” on page 86

Accessibility

Accessibility features help users with physical disabilities, such as restricted mobility or limited vision, to use software products successfully. These are the major accessibility features in DB2[®] Universal Database Version 8:

- DB2 allows you to operate all features using the keyboard instead of the mouse. See “Keyboard Input and Navigation” on page 94.
- DB2 enables you customize the size and color of your fonts. See “Accessible Display” on page 94.
- DB2 allows you to receive either visual or audio alert cues. See “Alternative Alert Cues” on page 94.
- DB2 supports accessibility applications that use the Java[™] Accessibility API. See “Compatibility with Assistive Technologies” on page 94.

- DB2 comes with documentation that is provided in an accessible format. See “Accessible Documentation”.

Keyboard Input and Navigation

Keyboard Input

You can operate the DB2 Tools using only the keyboard. You can use keys or key combinations to perform most operations that can also be done using a mouse.

Keyboard Focus

In UNIX-based systems, the position of the keyboard focus is highlighted, indicating which area of the window is active and where your keystrokes will have an effect.

Accessible Display

The DB2 Tools have features that enhance the user interface and improve accessibility for users with low vision. These accessibility enhancements include support for customizable font properties.

Font Settings

The DB2 Tools allow you to select the color, size, and font for the text in menus and dialog windows, using the Tools Settings notebook.

Non-dependence on Color

You do not need to distinguish between colors in order to use any of the functions in this product.

Alternative Alert Cues

You can specify whether you want to receive alerts through audio or visual cues, using the Tools Settings notebook.

Compatibility with Assistive Technologies

The DB2 Tools interface supports the Java Accessibility API enabling use by screen readers and other assistive technologies used by people with disabilities.

Accessible Documentation

Documentation for the DB2 family of products is available in HTML format. This allows you to view documentation according to the display preferences set in your browser. It also allows you to use screen readers and other assistive technologies.

DB2 tutorials

The DB2[®] tutorials help you learn about various aspects of DB2 Universal Database. The tutorials provide lessons with step-by-step instructions in the areas of developing applications, tuning SQL query performance, working with data warehouses, managing metadata, and developing Web services using DB2.

Before you begin:

Before you can access these tutorials using the links below, you must install the tutorials from the *DB2 HTML Documentation* CD-ROM.

If you do not want to install the tutorials, you can view the HTML versions of the tutorials directly from the *DB2 HTML Documentation* CD. PDF versions of these tutorials are also available on the *DB2 PDF Documentation* CD.

Some tutorial lessons use sample data or code. See each individual tutorial for a description of any prerequisites for its specific tasks.

DB2 Universal Database tutorials:

If you installed the tutorials from the *DB2 HTML Documentation* CD-ROM, you can click on a tutorial title in the following list to view that tutorial.

Business Intelligence Tutorial: Introduction to the Data Warehouse Center
Perform introductory data warehousing tasks using the Data Warehouse Center.

Business Intelligence Tutorial: Extended Lessons in Data Warehousing
Perform advanced data warehousing tasks using the Data Warehouse Center. (Not provided on CD. You can download this tutorial from the Downloads section of the Business Intelligence Solutions Web site at <http://www.ibm.com/software/data/bi/>.)

Development Center Tutorial for Video Online using Microsoft[®] Visual Basic
Build various components of an application using the Development Center Add-in for Microsoft Visual Basic.

Information Catalog Center Tutorial
Create and manage an information catalog to locate and use metadata using the Information Catalog Center.

Video Central for e-business Tutorial
Develop and deploy an advanced DB2 Web Services application using WebSphere[®] products.

Visual Explain Tutorial

Analyze, optimize, and tune SQL statements for better performance using Visual Explain.

DB2 Information Center accessed from a browser

The DB2[®] Information Center gives you access to all of the information you need to take full advantage of DB2 Universal Database[™] and DB2 Connect[™] in your business. The DB2 Information Center also documents major DB2 features and components including replication, data warehousing, the Information Catalog Center, Life Sciences Data Connect, and DB2 extenders.

The DB2 Information Center accessed from a browser has the following features if you view it in Netscape Navigator 6.1 or later or Microsoft Internet Explorer 5 or later. Some features require you to enable support for Java or JavaScript:

Regularly updated documentation

Keep your topics up-to-date by downloading updated HTML.

Search

Search all of the topics installed on your workstation by clicking **Search** in the navigation toolbar.

Integrated navigation tree

Locate any topic in the DB2 library from a single navigation tree. The navigation tree is organized by information type as follows:

- Tasks provide step-by-step instructions on how to complete a goal.
- Concepts provide an overview of a subject.
- Reference topics provide detailed information about a subject, including statement and command syntax, message help, requirements.

Master index

Access the information installed from the *DB2 HTML Documentation CD* from the master index. The index is organized in alphabetical order by index term.

Master glossary

The master glossary defines terms used in the DB2 Information Center. The glossary is organized in alphabetical order by glossary term.

Related tasks:

- “Finding topics by accessing the DB2 Information Center from a browser” on page 84

- “Finding product information by accessing the DB2 Information Center from the administration tools” on page 86
- “Updating the HTML documentation installed on your machine” on page 88

Appendix B. Notices

IBM may not offer the products, services, or features discussed in this document in all countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country/region or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country/region where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions; therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make

improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product, and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licenses of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information that has been exchanged, should contact:

IBM Canada Limited
Office of the Lab Director
8200 Warden Avenue
Markham, Ontario
L6G 1C7
CANADA

Such information may be available, subject to appropriate terms and conditions, including in some cases payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems, and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements, or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information may contain examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious, and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information may contain sample application programs, in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows:

© (*your company name*) (*year*). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. *_enter the year or years_*. All rights reserved.

Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both, and have been used in at least one of the documents in the DB2 UDB documentation library.

ACF/VTAM	LAN Distance
AISPO	MVS
AIX	MVS/ESA
AIXwindows	MVS/XA
AnyNet	Net.Data
APPN	NetView
AS/400	OS/390
BookManager	OS/400
C Set++	PowerPC
C/370	pSeries
CICS	QBIC
Database 2	QMF
DataHub	RACF
DataJoiner	RISC System/6000
DataPropagator	RS/6000
DataRefresher	S/370
DB2	SP
DB2 Connect	SQL/400
DB2 Extenders	SQL/DS
DB2 OLAP Server	System/370
DB2 Universal Database	System/390
Distributed Relational Database Architecture	SystemView
DRDA	Tivoli
eServer	VisualAge
Extended Services	VM/ESA
FFST	VSE/ESA
First Failure Support Technology	VTAM
IBM	WebExplorer
IMS	WebSphere
IMS/ESA	WIN-OS/2
iSeries	z/OS
	zSeries

The following terms are trademarks or registered trademarks of other companies and have been used in at least one of the documents in the DB2 UDB documentation library:

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Index

A

- access plans
 - Visual Explain 52
- accessibility
 - features 93
- administration notification log
 - Health Center 40
 - Journal 42
- advisors 15
- AIX
 - DB2 tools vi
- Alert Center, replaced by Health Center 40
- alerts
 - tool 40
- aliases
 - description 8
- authorities
 - description 9

B

- backing up
 - databases 63
- backups
 - database 61
- beacons
 - Health Center 40
- buffer pools
 - description 10

C

- capture control tables
 - replication 67
- catalogs
 - description 7
 - Information Catalog Center 39
- clients
 - configuration 55
- cloning subsystems
 - Control Center 28
- Command Center
 - description 35
 - web version 57
- commands
 - tool 35
- configuration
 - clients 55
- Configuration Assistant (CA)
 - description 55

- configuration parameters
 - description 11
- Control Center
 - description 22
 - federated systems 26
 - OS/390 and z/OS objects 28
 - required authority vi
- control tables
 - replication 67
- customized views 15

D

- data elements
 - description 12
- Data Warehouse Center
 - description 34
- data warehousing
 - tool 34
- database aliases
 - description 8
- database manager
 - description 7
- database partition groups
 - description 5
- databases
 - description 4
- DB2 Connect
 - Control Center 22
- DB2 documentation search
 - using Netscape 4.x 91
- DB2 Information Center 96
- DB2 tutorials 95
- details view 15
- Development Center
 - description 45
- disability 93
- displaying help 15

E

- explain tools
 - Visual Explain 52

F

- federated systems
 - Control Center 26
 - description 14
- filtering 15

H

- Health Center
 - description 40
 - web version 58
- health monitor
 - tool 40

I

- Indoubt Transaction Manager
 - description 56
- indoubt transactions
 - tool 56
- infopops 15
- Information Catalog Center
 - description 39
- information catalogs
 - tool 39
- Information Center
 - description 46
- instances
 - description 3

J

- Journal
 - description 42

L

- large objects (LOBs)
 - description 6
- launchpads 15
- License Center
 - description 44
- licenses
 - tool 44
- logging
 - archive 63
 - circular 63
- logs
 - archive 61
 - circular 61
- long data
 - description 6

M

- memory
 - visualizer 48
- Memory Visualizer
 - description 48
- menu bars 15

O

- object trees 15
- offline
 - backup 63
- online
 - backup 63
 - help, accessing 82
- ordering DB2 books 82
- OS/390 and z/OS objects
 - Control Center 28

P

- packages
 - description 8
- printed books, ordering 82
- procedures
 - tool 45

R

- registry variables
 - description 12
- replication
 - description 13
 - tool 30
- Replication Center
 - description 30
 - using 67
- replication sources
 - enabling 67
- restore
 - database 62
- rollforward
 - database 62

S

- SATCTLDB database
 - Satellite Administration Center 33
- Satellite Administration Center
 - description 33
- satellites
 - administration 33
- screen resolution vi
- Script Center, replaced by Task Center 37
- scripts
 - tool 37
- sources
 - replication 67
- spatial data
 - adding 50
- Spatial Extender
 - description 50
- SQL Assist
 - description 50

- starting the GUI administration
 - tools 15
- Stored Procedure Builder, replaced by Development Center 45
- stored procedures
 - tool 45
- subsystems
 - cloning with Control Center 28
- systems
 - description 2

T

- table aliases
 - description 8
- table spaces
 - description 5
- targets
 - replication 67
- Task Center
 - description 37
- tasks
 - tool 37
- toolbars 15
- transactions
 - indoubt 56
- transparent DDL
 - Control Center 26
- troubleshooting
 - DB2 documentation search 91
 - online information 93
- tutorials 95

V

- Visual Explain
 - description 52

W

- warehousing
 - tool 34
- Web tools
 - Command Center 57
 - Health Center 58
- wizards 15

Z

- z/OS objects
 - Control Center 28

Contacting IBM

In the United States, call one of the following numbers to contact IBM:

- 1-800-237-5511 for customer service
- 1-888-426-4343 to learn about available service options
- 1-800-IBM-4YOU (426-4968) for DB2 marketing and sales

In Canada, call one of the following numbers to contact IBM:

- 1-800-IBM-SERV (1-800-426-7378) for customer service
- 1-800-465-9600 to learn about available service options
- 1-800-IBM-4YOU (1-800-426-4968) for DB2 marketing and sales

To locate an IBM office in your country or region, check IBM's Directory of Worldwide Contacts on the web at www.ibm.com/planetwide

Product information

Information regarding DB2 Universal Database products is available by telephone or by the World Wide Web at www.ibm.com/software/data/db2/udb

This site contains the latest information on the technical library, ordering books, client downloads, newsgroups, FixPaks, news, and links to web resources.

If you live in the U.S.A., then you can call one of the following numbers:

- 1-800-IBM-CALL (1-800-426-2255) to order products or to obtain general information.
- 1-800-879-2755 to order publications.

For information on how to contact IBM outside of the United States, go to the IBM Worldwide page at www.ibm.com/planetwide



Printed in U.S.A.

SC09-4851-00



Spine information:



IBM® DB2 Universal Database™ SC09-4851

Version 8