

IBM<sup>®</sup> DB2<sup>®</sup> Universal Database  
Enterprise - Extended Edition  
for Windows<sup>®</sup>



# Quick Beginnings

*Version 7*



IBM<sup>®</sup> DB2<sup>®</sup> Universal Database  
Enterprise - Extended Edition  
for Windows<sup>®</sup>



# Quick Beginnings

*Version 7*

Before using this information and the product it supports, be sure to read the general information under "Appendix F. Notices" on page 149.

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.

Order publications through your IBM representative or the IBM branch office serving your locality or by calling 1-800-879-2755 in the United States or 1-800-IBM-4YOU in Canada.

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 1993, 2000. All rights reserved.**

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

---

# Contents

<b>Welcome to DB2 Universal Database!</b> . . .	<b>vii</b>
Conventions . . . . .	vii

---

## **Part 1. Introduction to DB2 Enterprise - Extended Edition . . . . 1**

<b>Chapter 1. Introduction to DB2 Enterprise - Extended Edition . . . . .</b>	<b>3</b>
Processing in a Partitioned Database Environment . . . . .	4
Cost-Based Query Optimization . . . . .	4
Configuration . . . . .	5
Machines and Storage . . . . .	7
Nodegroups and Data Partitioning . . . . .	7
Multiple Logical Nodes . . . . .	9
Instances . . . . .	10
Fast Communications Manager . . . . .	11
High Availability . . . . .	11
Virtual Interface Architecture . . . . .	12

---

## **Part 2. Installing and Configuring DB2 Enterprise - Extended Edition 13**

<b>Chapter 2. Planning for Installation . . . . .</b>	<b>15</b>
Memory Requirements . . . . .	15
Disk Requirements . . . . .	16
DB2 Servers . . . . .	16
DB2 Clients . . . . .	16
Software Requirements . . . . .	17
Server Product Requirements . . . . .	17
Client Product Requirements . . . . .	18
Possible Client-to-Server Connectivity Scenarios . . . . .	19
Migrating from Previous Versions of DB2 . . . . .	20
Preparing Databases and Instances for Migration . . . . .	21
Verifying Databases are Ready for Migration . . . . .	22
Migrating from Visual Warehouse Version 3.1 or Visual Warehouse Version 5.2 . . . . .	24
Migrating the Visual Warehouse Server and Active Warehouse Control Database . . . . .	25
Migrating Warehouse Agents . . . . .	25

Migrating Warehouse Transformers . . . . .	25
Migrating DataGuide . . . . .	26
NetQuestion Search System . . . . .	26
Stopping the Windows Cluster Service . . . . .	26
Your next step . . . . .	27

<b>Chapter 3. Installing DB2 Universal Database Enterprise - Extended Edition on Windows. . . . .</b>	<b>29</b>
Before You Begin . . . . .	29
Performing the Installation . . . . .	32
Adding a Logical Database Partition Server to an Instance . . . . .	35
Changing the Database Partition Server Configuration . . . . .	37
Verifying the Installation . . . . .	39

<b>Chapter 4. DB2 Post-Installation Migration Tasks . . . . .</b>	<b>43</b>
Migrating Databases . . . . .	43
Updating Single-Partition Instances for use in Partitioned Database Systems . . . . .	45
Optional Post Migration Actions . . . . .	46

<b>Chapter 5. Your Next Step. . . . .</b>	<b>49</b>
DB2 Universal Database Administration Guide . . . . .	49
DB2 Universal Database Administration Guide: Planning . . . . .	49
DB2 Universal Database Administration Guide: Implementation . . . . .	49
DB2 Universal Database Administration Guide: Performance . . . . .	50
Task Information in other DB2 Documentation . . . . .	50

---

## **Part 3. Installing and Configuring DB2 Clients . . . . . 51**

<b>Chapter 6. Installing DB2 Clients . . . . .</b>	<b>53</b>
DB2 Run-Time Client . . . . .	53
DB2 Administration Client . . . . .	54
DB2 Application Development Client . . . . .	54
Distributed Installation . . . . .	54

DB2 Thin Client . . . . .	55
---------------------------	----

**Chapter 7. Installing DB2 Clients on Windows 32-Bit Operating Systems . . . . . 57**

Before You Begin Installing . . . . .	57
Installing Without Administrator Authority	57
Installation Steps . . . . .	58

**Chapter 8. Configuring Client-to-Server Communications Using the Client Configuration Assistant . . . . . 61**

LDAP Directory Support Considerations . . . . .	61
Before You Begin . . . . .	61
Configuration Steps . . . . .	62
Adding a Database Using a Profile . . . . .	63
Adding a Database Using Discovery . . . . .	64
Adding a Database Manually . . . . .	66
Creating and Using Profiles . . . . .	68
Server Profiles . . . . .	68
Client Profiles . . . . .	69

**Chapter 9. Control Center Installation and Configuration . . . . . 73**

Application versus Applet . . . . .	73
Machine Configurations . . . . .	74
Supported Java Virtual Machines for the Control Center . . . . .	75
Setting Up and Working with the Control Center . . . . .	76
Control Center Services Setup (Applet Mode only) . . . . .	76
Working with the Control Center . . . . .	78
Functional Considerations . . . . .	80
Installation Tips for Control Center Help on UNIX Operating Systems . . . . .	80
Configuring TCP/IP on OS/2 . . . . .	81
Enabling Local Loopback . . . . .	81
Enabling Localhost . . . . .	81
Verifying TCP/IP Configuration on OS/2 . . . . .	82
Considerations for Remotely Administering Clusters Using the Control Center. . . . .	82
Troubleshooting Information . . . . .	83
Administering DB2 for OS/390 and DB2 Connect Enterprise Edition servers with the Control Center . . . . .	84
Preparing DB2 for OS/390 Servers for the Control Center . . . . .	84
Working with the Control Center . . . . .	85
Other Sources of Information . . . . .	85

**Part 4. Appendixes . . . . . 87**

**Appendix A. Basic Task Knowledge . . . . . 89**

Starting First Steps . . . . .	89
Starting the Client Configuration Assistant . . . . .	89
Starting the DB2 Control Center . . . . .	90
Entering Commands Using the Command Center . . . . .	90
Entering Commands Using the Command Line Processor . . . . .	91
DB2 Command Window . . . . .	92
Interactive Input Mode . . . . .	92
Working with the System Administrative Group . . . . .	93
Granting Advanced User Rights on Windows . . . . .	94
Windows NT. . . . .	94
Windows 2000 . . . . .	95
Working with the Business Intelligence Functions . . . . .	95
Setting the Number of Licensed Processors. . . . .	95
Upgrading DB2 from Try and Buy Mode . . . . .	96
Uninstalling DB2 . . . . .	97

**Appendix B. National Language Support (NLS) . . . . . 99**

Code Page and Language Support . . . . .	99
--	----

**Appendix C. Naming Rules . . . . . 101**

General Naming Rules . . . . .	101
Database, Database Alias, and Catalog Node Name Rules. . . . .	101
Object Name Rules . . . . .	102
Username, User ID, Group Name, and Instance Name Rules. . . . .	103
Workstation Name (nname) Rules . . . . .	104
DB2SYSTEM Naming Rules . . . . .	105
Password Rules . . . . .	105

**Appendix D. About DB2 Universal Database for UNIX, Windows, and OS/2 . . . . . 107**

DB2 Products . . . . .	107
DB2 Everywhere . . . . .	107
DB2 Universal Database . . . . .	108
Host Databases . . . . .	110
DB2 Connect . . . . .	110
Related Products . . . . .	111
DB2 Relational Connect . . . . .	111
DB2 Warehouse Manager . . . . .	111
DB2 OLAP Server. . . . .	112

Intelligent Miner . . . . .	112	Managing Warehouses using the Data	
DB2 Spatial Extender. . . . .	112	Warehouse Center . . . . .	128
DB2 Net Search Extender . . . . .	112	Understanding the Administration Server	128
DB2 Data Links Manager . . . . .	113	Developing Applications using the DB2	
Tivoli Enterprise . . . . .	113	Application Development Client . . . . .	129
Working with DB2 Data. . . . .	113	Running Your Own Applications. . . . .	130
Accessing DB2 Data from Remote Clients	113	<b>Appendix E. Using the DB2 Library . . . . .</b>	<b>131</b>
Accessing Multiple DB2 Servers . . . . .	114	DB2 PDF Files and Printed Books . . . . .	131
Accessing Host or AS/400 DB2 Data from		DB2 Information . . . . .	131
the Desktop using DB2 Connect		Printing the PDF Books . . . . .	140
Enterprise Edition. . . . .	115	Ordering the Printed Books . . . . .	141
Accessing DB2 Data from the Web using		DB2 Online Documentation . . . . .	142
Java . . . . .	119	Accessing Online Help . . . . .	142
Accessing DB2 Data from the Web using		Viewing Information Online . . . . .	144
Net.Data. . . . .	120	Using DB2 Wizards . . . . .	146
Accessing DB2 Data from Host and		Setting Up a Document Server . . . . .	147
AS/400 Client Machines . . . . .	122	Searching Information Online. . . . .	148
Administering Instances and Databases with		<b>Appendix F. Notices . . . . .</b>	<b>149</b>
the DB2 Administration Tools. . . . .	122	Trademarks . . . . .	152
Managing Instances and Database Objects		<b>Index . . . . .</b>	<b>155</b>
using the Control Center . . . . .	122	<b>Contacting IBM . . . . .</b>	<b>159</b>
Managing Communications on the Server	126	Product Information . . . . .	159
Monitoring Databases using DB2			
Performance Monitor . . . . .	126		
Viewing SQL Access Plans using Visual			
Explain . . . . .	126		
Managing Connections to Databases using			
the Client Configuration Assistant . . . . .	127		





---

# Welcome to DB2 Universal Database!

The DB2 Universal Database Quick Beginnings books provide a focused introduction to the installation and configuration of DB2 products.

This *Quick Beginnings* book will guide you through the planning, installation, migration (if necessary), and setup of a partitioned database system. You will first install the instance-owning database partition server and then create (or migrate) an instance. On another machine, you will install a database partition server that will participate in the instance that you created. After you have set up and configured your partitioned database system, you will create the SAMPLE database. Finally, once the DB2 server has been installed and the SAMPLE database has been created, you will install a DB2 client and configure a connection between the client and server using the DB2 GUI tools.



---

## Conventions

This book uses the following highlighting conventions:

- **Boldface** indicates commands or graphical user interface (GUI) controls such as names of fields, folders, icons, or menu choices.
- *Italics* indicates variables that you should replace with your own value. It is also used to indicate book titles and to emphasize words.
- Monospace indicates file names, directory paths, and examples of text you enter exactly as shown.



This icon marks a fast path. A fast path guides you to information specific to your configuration where multiple options are available.



This icon marks a tip. It provides additional information that can help you complete a task.

For a complete description of the DB2 library, see “Appendix E. Using the DB2 Library” on page 131.



- If you do not follow the documented installation method with the recommended defaults, it may be necessary to refer to the *Administration Guide* and the *Command Reference* to complete the installation and configuration.
- The term *Windows 32-bit operating systems* refers to Windows 95, Windows 98, Windows NT, or Windows 2000.
- The term *Windows 9x* refers to Windows 95 or Windows 98.
- The term *DB2 client* refers to a DB2 Run-Time Client, a DB2 Administration Client, or a DB2 Application Development Client.
- The term *database partition server* is also referred to as a *node*.
- In this book, the term *DB2 Universal Database* refers to DB2 Universal Database on OS/2, UNIX, and Windows 32-bit operating systems, unless otherwise stated.

---

# **Part 1. Introduction to DB2 Enterprise - Extended Edition**



---

# Chapter 1. Introduction to DB2 Enterprise - Extended Edition

A *database* is simply a collection of data. A *database manager* is the software that allows users to store and access data in a database using system resources, including CPU, memory, disk space, and communications. In a partitioned database system, a single database manager and the collection of data and system resources that it manages are referred to collectively as a *database partition server (node)*. A *partitioned database system* is the collection of all the database partition servers that you create to handle data requests.

In DB2 Enterprise - Extended Edition (DB2 EEE), the quality, functionality, reliability, and robustness of a the database manager is extended to the entire partitioned database system.

In a partitioned database system, multiple database partition servers can be assigned to a machine (or to multiple machines), and the database manager at each machine is responsible for a portion of a database's total data (each database partition server houses a portion of the entire database). This portion of the database is known as a *database partition (node)*. The fact that databases are partitioned across database partition servers is transparent to users and applications.

A partitioned database system can maintain very large databases and open opportunities for new applications. DB2 EEE provides fast response time for both decision-support (DSS) and online transaction processing (OLTP) applications.

DB2 EEE can be configured to execute on a *shared-nothing* hardware architecture, in which machines do not compete for resources. Each machine has exclusive access to its own disks and memory, and the database partition servers that run on the machines communicate with each other through the use of messages. For a database system to exploit shared-nothing architecture, typically one database partition server is assigned to each machine. Another possible configuration is running multiple logical nodes, in which more than one database partition server runs on a machine. For more information, see "Multiple Logical Nodes" on page 9.

---

## Processing in a Partitioned Database Environment

In DB2 Enterprise - Extended Edition (DB2 EEE), a database is distributed across multiple machines, and database partition servers are installed on a set of machines. Because the database is partitioned across multiple machines, you can use multiple CPUs across multiple machines to satisfy requests for information. The retrieval and update requests are decomposed automatically into subrequests and executed in parallel on the database partition servers on each machine.

As an illustration of the power of processing in a partitioned database system, assume that you have 100 000 000 records that you want to scan in a single-partition database. This scan would require that a single database manager search 100 000 000 records. Now suppose that these records are spread evenly over 20 database partition servers; each database manager only has to scan 5 000 000 records. If each database partition server scans at the same time and with the same speed, the time required to do the scan should be approximately 5% of that of a single-partition system handling this task.

User interaction with DB2 EEE is handled through a database partition server (node). This database partition server is known as the *coordinator node* for the partitioned database system. The instance-owning machine, by default, is the coordinator node for a partitioned database system. Any database partition server can act as a coordinator node. The database partition server that a client or application connects to becomes the coordinator node. You should consider spreading out users across database partitions servers to distribute the coordinator function. For more information, refer to the *Administration Guide*.

DB2 EEE keeps communications overhead as low as possible. For example, if a row is being added to a table, the database partition server checks a *partitioning map*, which specifies the database partition server where the row is stored. The row is only sent to that database partition server, with the result that only the interested database partition servers take part in the insert.

---

## Cost-Based Query Optimization

DB2 Enterprise - Extended Edition (DB2 EEE) uses a *cost-based query optimizer*, which compares different methods for doing a unit of work, and selects the most efficient one. The optimizer provides the following features:

### Transparent parallelism

Both new and existing applications that use data-manipulating SQL statements do not have to be changed when they are migrated to DB2 EEE. You only have to rebind them so the optimizer can generate the best plans for existing SQL queries.

### **Comprehensive use of data partitioning information**

The optimizer uses information about how base tables, and the intermediate tables that result from queries, are partitioned across database partitions. This information is used to determine the best execution strategy.

### **Full-fledged cost-based SQL optimization**

The optimizer has information about how the data is partitioned. With this information, the optimizer considers different execution plans and chooses the one with the lowest cost. While comparing different strategies, it accounts for the inherent parallelism of different operations, and the costs introduced by messaging between database partition servers.

When generating plans, the optimizer considers different parallel joining methods, including *collocated*, *directed*, and *broadcast* joins. For more information on joins, refer to the *Administration Guide*.

### **Inter-partition and Intra-partition parallelism of all relational operations.**

All operations, such as index and table scans, aggregation, set operations, joins, inserts, deletes, and updates can employ both *Inter-partition parallelism* and *Intra-partition parallelism*.

Inter-partition parallelism means that the operator is executed in parallel by each database partition server. For example, assume that you issue a **SELECT** statement to fetch data that meets some condition. The coordinator node sends this request to the other database partition servers to select this data set from that data that is stored on each database partition. Each database partition server then sends this data back to the coordinator node which does the final processing and returns a resulting set.

Intra-partition parallelism means that different operators in the same query can be executed in parallel by the same database partition server. For example, if a SQL query included a scan, join, and sort, the database partition server would process these operators, to the best of its abilities, in parallel.

---

## **Configuration**

Figure 1 on page 6 shows an example of a DB2 Enterprise - Extended Edition (DB2 EEE) hardware configuration.

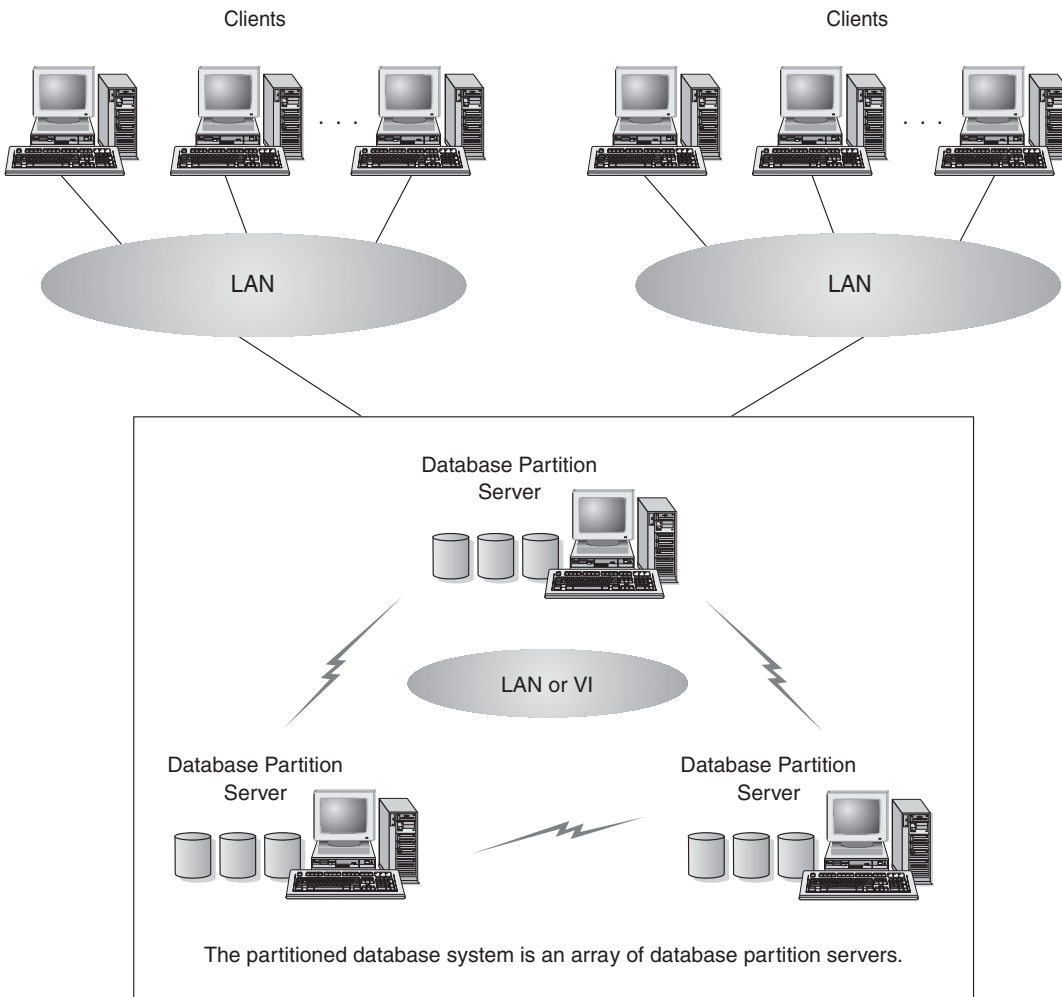


Figure 1. DB2 Enterprise - Extended Edition Hardware Configuration

DB2 EEE can run on a cluster of individual machines interconnected by a LAN. Communications between all database partition serves is handled by a component of DB2 EEE called the Fast Communications Manager (FCM).

The following sections provide information that you should be familiar with before you configure your partitioned database system. Specifically, they describe:

- Machines and Storage
- Nodegroups and Data Partitioning
- Multiple Logical Nodes
- Instances



- Fast Communication Manager (FCM)
- High Availability
- Virtual Interface (VI) Architecture

## Machines and Storage

DB2 Enterprise - Extended Edition implements a shared-nothing architecture, therefore each database partition server is the equivalent of a single-partition database system. Thus, the database storage capacity for the partitioned database system is equal to that provided by a single-partition database system multiplied by the number of database partition servers. You can store tables of up to 512 GB (gigabytes) per database partition. For example, in a database that has 128 partitions, the maximum size of one table is approximately 64 TB (terabytes).

## Nodegroups and Data Partitioning

You can define named subsets of one or more database partitions in a database. Each subset you define is known as a *nodegroup*. Each subset that contains more than one database partition is known as a *multipartition nodegroup*. Multipartition nodegroups can only be defined within database partitions that belong to the same database.

Three default nodegroups are created when you create a database: IBMDEFAULTGROUP, IBMCATGROUP, and IBMTEMPGROUP.

If you want, you can create table spaces in the default nodegroups IBMDEFAULTGROUP and IBMCATGROUP, and then create tables within those table spaces.

The IBMDEFAULTGROUP nodegroup contains all the database partitions for the database. When you create a database, a database partition is created at each database partition server (node) in the instance.

The IBMCATGROUP nodegroup for the database is created at the database partition server where you enter the **create database** command. This nodegroup only contains the database partition that is local to the database partition server where the command was entered. This database partition server is referred to as the *catalog node* of the database because the IBMCATGROUP nodegroup contains the catalog tables for the database.

You cannot directly work with the third default nodegroup, IBMTEMPGROUP. Like the IBMDEFAULTGROUP nodegroup, it also contains all the database partitions of a database. This nodegroup is used to contain all temporary table spaces.

Figure 2 on page 8 shows an example of a database in which there are three nodegroups. Nodegroup 1 is a multipartition nodegroup made of four

database partitions, nodegroup 2 is a single-partition nodegroup, and nodegroup 3 is a multipartition nodegroup.




Figure 2. Nodegroups in a Database

When you want to create table spaces for a database, you first create the nodegroup where the table spaces will be stored, then create a table space in the nodegroup. After this, you create the tables in the table space.

You can drop database partitions from a nodegroup, or if you have increased the size of the database system you can add new database partitions to existing nodegroups. For information about adding and dropping nodes in nodegroups, refer to the *Administration Guide*.

As your database increases in size, you can add database partition servers to the database system for improved performance. This is known as scaling the database system. When you add a database partition server, a database partition is created for each database that already exists in the database

system. You then add the new database partition to an existing nodegroup that belongs to that database. Finally, you redistribute data in that nodegroup to utilize the new database partition. For information about scaling databases, refer to the *Administration Guide*.

Each table defined in a multipartition nodegroup has a *partitioning key* associated with it. The partitioning key is an ordered set of columns whose values are used in conjunction with a *partitioning map* to determine the database partition on which a row of a given table resides. The partitioning map is an array of 4 096 database partition numbers.

Columns of any data type (except LONG VARCHAR, LONG VARGRAPHIC, BLOB, or CLOB) can be used as the partitioning key. A table defined in a single-partition nodegroup may or may not have a partitioning key. Tables with only long-field columns can only be defined in single-partition nodegroups, and they cannot have a partitioning key. For more information about creating tables, refer to the *SQL Reference*.

The use of nodegroups and partitioning keys means that:

- Data can be distributed across multiple database partitions to reduce I/O and processing bottlenecks
- Data can be redistributed when large volumes of system activity or an increase in table size require the addition of more machines.

For more information about creating nodegroups, refer to the *SQL Reference*. For more information about using nodegroups, refer to the *Administration Guide*.

## Multiple Logical Nodes

Typically, you configure DB2 Enterprise - Extended Edition to have one database partition server assigned to each machine. There are situations, however, in which it would be advantageous to have more than one database partition server assigned to each machine. If these database partition servers (nodes) participate in the *same* instance, this is referred to as a Multiple Logical Node (MLN) configuration.

A Multiple Logical Node (MLN) configuration is useful when the system runs queries on a machine that has symmetric multiprocessor (SMP) architecture. Another benefit is that multiple logical nodes can exploit SMP hardware configurations. In addition, because database partitions are smaller, you can obtain better performance when performing such tasks as backing up and restoring database partitions and table spaces, and creating indexes. As a general rule, we recommend that you run one MLN per 4 processors. Depending on the operating system where you are running DB2 EEE, this may vary for performance reasons.

## Instances

An instance has its own databases and instance directory. The instance directory contains the database manager configuration file, system database directories, node directories, and the node configuration file. For more information on instances in a partitioned database system, refer to the *Administration Guide*.

In DB2 Enterprise - Extended Edition, an instance contains all the database partition servers (nodes) that were defined to take part in a given partitioned database system. The instance-owning machine (known as node 0) owns the shared directory where this information is stored. Other database partition servers that are added to an instance are said to be participating in the instance.

Each instance has different security from other instances on the same machine. This is shown in Figure 3, which shows two separate instances. Instance 1 contains six database partition servers and Instance 2 contains eight database partition servers. (Multiple database partition servers are indicated when more than one line is shown between a database partition server and the instance directory.) The two instances overlap, but this is due to the assignment of two database partition servers to each of the three machines in the middle of the figure.

Database partition servers only belong to one instance. For example, Instance 1 does not contain database partition servers that belong to Instance 2.

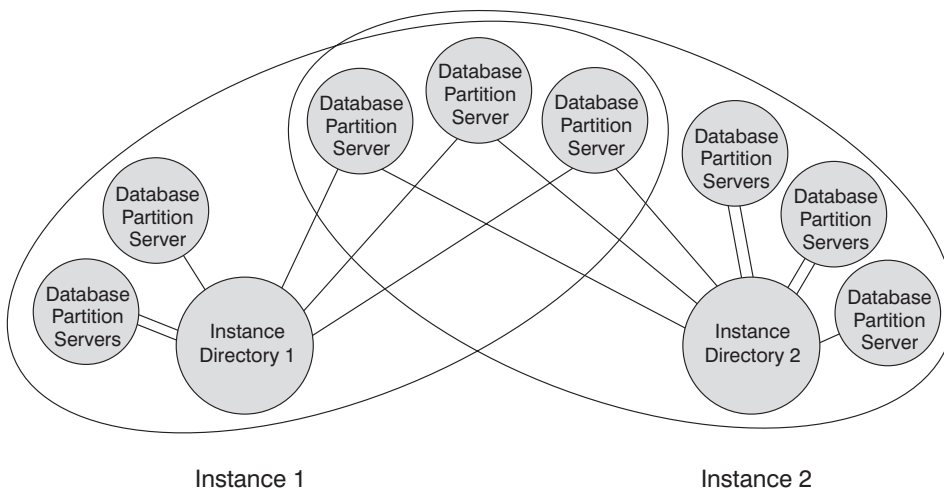


Figure 3. Two Instances

You can have multiple instances on the same machine, with each configured differently:

- To have distinct test and production environments
- To restrict access to specific databases
- To exploit different database configurations.

Each instance is owned by the machine where the instance was first created (known as the instance-owning machine, node 0). The instance-owning machine stores information that is common to all of the database partition servers. For information about creating instances, refer to the *Administration Guide*.

## Fast Communications Manager

The fast communication manager (FCM) provides communication support for DB2 Enterprise - Extended Edition. Each database partition server has one FCM thread to provide communications between database partition servers to handle agent requests, and to deliver message buffers. The FCM thread is started when you start the instance.

If communications fail between database partition servers or if they re-establish communications, the FCM thread updates information (that you can query with the database system monitor) and causes the appropriate action (such as the rollback of an affected transaction) to be performed. You can use the database system monitor to help you set the FCM configuration parameters. For information about FCM-related database system monitor output, refer to the *System Monitor Guide and Reference*.



You can specify the number of FCM message buffers with the `fc_m_num_buffers` database manager configuration parameter. For a description of this and other FCM parameters, refer to the *Administration Guide*.

## High Availability

You can set up your partitioned database system so that if a machine fails, the database server on the failed machine can run on another machine. MSCS can perform both failure detection and the restarting of resources in a clustered environment, such as failover support for physical disks and IP addresses. On Windows, you implement failover support using Microsoft Cluster Service (MSCS). To use MSCS, you require Windows NT Server 4.0 Enterprise Edition or Windows 2000 Advanced Server with the MSCS component installed. For more information, refer to the *Administration Guide*.

## Virtual Interface Architecture

On Windows, DB2 Enterprise – Extended Edition (DB2 EEE) can take advantage of Virtual Interface (VI) Architecture. VI Architecture was developed through the efforts of various companies to address the needs for a standard high-volume interconnect for data transfer between servers. VI Architecture permits high volumes of data to pass very quickly between clustered servers.

Prior to the release of VI Architecture, communications between database partition servers in a cluster was done through the network infrastructure provided by the operating system. This amounted to processing overhead on the operating system each time any communications between partitioned database servers took place. VI Architecture defines a thin, fast interface that connects software applications directly to the networking hardware, while retaining the robust security protection of the operating system. In a communications-intensive environment, implementing VI Architecture with DB2 EEE can realize significant improvements in overall system throughput of database transactions and queries. For more information, refer to the *Administration Guide*.

---

## **Part 2. Installing and Configuring DB2 Enterprise - Extended Edition**





---

## Chapter 2. Planning for Installation

Before you install DB2, you should make sure that your system meets DB2's hardware and software requirements. If you are migrating from a previous version of DB2, there are also pre-installation migration tasks you should perform to prepare your databases.

This chapter describes the following requirements you should take into consideration before installing DB2:

- "Memory Requirements".
- "Disk Requirements" on page 16.
- "Software Requirements" on page 17.
- "Possible Client-to-Server Connectivity Scenarios" on page 19.
- "Migrating from Previous Versions of DB2" on page 20.



If you know that your system meets all hardware and software requirements, and you want to begin installing your DB2 product right away, go to "Chapter 3. Installing DB2 Universal Database Enterprise - Extended Edition on Windows" on page 29.

For information on the DB2 family of products, see "Appendix D. About DB2 Universal Database for UNIX, Windows, and OS/2" on page 107.

---

---

### Memory Requirements

The amount of random-access memory (RAM) that you require to run your DB2 server depends on the size of your databases and on the administration tools you will use. For example, if you plan to use the DB2 GUI tools to administer and configure your DB2 databases, the minimum amount of RAM we recommend is 128 MB.

The DB2 GUI tools are a set of administration and configuration tools that include the Command Center, Control Center, and the Data Warehouse Center. Additional GUI tools are available on Windows 32-bit and OS/2 operating systems. These include the Client Configuration Assistant, Event Monitor and Event Analyzer. For more information about DB2 GUI tools and the platforms on which they are available, refer to the *Administration Guide*.

To run a DB2 Run-Time Client or an DB2 Application Development Client, you need a minimum of 16 MB of RAM. If you are planning to run a DB2 Administration Client, you need a minimum of 32 MB of RAM.

---

## Disk Requirements

The actual fixed disk requirements of your installation may vary depending on your file system and the components you install. Ensure that you have included a disk space allowance for your operating system, application development tools, application data, and communications products. For more information about space requirements for data, refer to the *Administration Guide*.

### DB2 Servers

A *typical* installation of DB2 requires a minimum of 245 MB of disk space. This amount includes the online product documentation, tools, and the Java Runtime Environment.

### DB2 Clients

Use Table 1 to estimate the amount of disk space you need on each of your client workstations. You may require additional amounts of disk space depending on your file system.

*Table 1. Disk Requirements for Client Components*

Client Components	Recommended Minimum Disk (MB)
<b>OS/2</b>	
DB2 Run-Time Client	30 MB
DB2 Application Development Client	125 MB, not including the Java Development Kit (JDK)
DB2 Administration Client	95 MB
<b>UNIX platforms</b>	
DB2 Run-Time Client	30 to 40 MB (70 MB for Silicon Graphics IRIX)
DB2 Application Development Client	90 to 120 MB, not including the JDK (40 MB for NUMA-Q)
DB2 Administration Client	80 to 110 MB
<b>Note:</b> PTX/NUMA-Q and Silicon Graphics IRIX operating systems do not support the DB2 Administration Client.	
<b>Windows 32-bit Operating Systems</b>	
DB2 Run-Time Client	25 MB
DB2 Application Development Client	325 MB, including the JDK
DB2 Administration Client	125 MB

The DB2 Application Development Client and DB2 Administration Client include tools and documentation, except on NUMA-Q systems.

---

## Software Requirements

This section outlines the software required to run DB2 products.

### Server Product Requirements

Table 2 lists the operating system and communications software required for DB2 Universal Database.

On all platforms you will need a Java Runtime Environment (JRE) Version 1.1.8 to run DB2's Java-based tools, such as the Control Center. If you intend to run the Control Center as an applet, you will need a Java-enabled browser. See "Chapter 9. Control Center Installation and Configuration" on page 73 for more information.

*Table 2. Software Requirements*

Hardware/Software Requirements	Communications
<ul style="list-style-type: none"><li>Windows NT Version 4.0 with Service Pack 5, or later</li><li>Windows 2000</li></ul>	DB2 Universal Database for Windows
	<p>APPC, TCP/IP, IPX/SPX, Named Pipes, and MPTN (APPC over TCP/IP)</p> <ul style="list-style-type: none"><li>The Windows NT and Windows 2000 base operating systems provide NetBIOS, IPX/SPX, Named Pipes, and TCP/IP connectivity.</li></ul> <p>For SNA (APPC) connectivity, one of the following communication products is required:</p> <ul style="list-style-type: none"><li>Windows NT: IBM Communications Server Version 5.01 or later</li><li>Windows 2000: IBM Communications Server Version 6.1 or later</li><li>Windows 9x and Windows NT: IBM Personal Communications for Windows Version 4.2 or later</li><li>Windows 2000: IBM Personal Communications for Windows Version 4.3 CSD2 or later</li><li>Microsoft SNA Server Version 3 Service Pack 3 or later</li></ul>
	<p><b>Notes:</b></p> <ol style="list-style-type: none"><li>If SNA two-phase commit is necessary, IBM Communications Server Version 5.01 or later or Microsoft SNA Server Version 4 Service Pack 3 or later is required.</li><li>If you plan to use LDAP (Lightweight Directory Access Protocol), you require either a Microsoft LDAP client or an IBM SecureWay LDAP client V3.1.1. For more information, refer to the <i>Administration Guide</i>.</li><li>If you plan to use the Tivoli Storage Manager facilities for backup and restore of your databases, you require the Tivoli Storage Manager Client Version 3 or later.</li><li>If you have the IBM Antivirus program installed on your operating system, it must be disabled or uninstalled to complete a DB2 installation.</li><li>If you plan to use the Simple Network Management Protocol (SNMP) subagent, you require DPI 2.0 provided by IBM SystemView Agent.</li><li>If you plan to set up failover support, then you require Windows NT Enterprise Edition V4.0, or Windows 2000 with Microsoft Cluster Server support.</li></ol>

---

## Client Product Requirements

Table 3 lists the software requirements needed for a DB2 Administration Client, DB2 Run-Time Client, or a DB2 Application Development Client.

On all platforms you will need a Java Runtime Environment (JRE) Version 1.1.8 to run the DB2 tools, such as the Control Center. If you intend to run the Control Center as an applet on Windows 32-bit or OS/2 systems, you will need a Java-enabled browser. See “Chapter 9. Control Center Installation and Configuration” on page 73 for more information.

Table 3. Software Requirements for Clients

Component	Hardware/Software Requirements	Communications
<ul style="list-style-type: none"><li>• DB2 Run-Time Client for Windows 9x</li><li>• DB2 Administration Client for Windows 9x</li><li>• DB2 Application Development Client for Windows 9x</li></ul>	<ul style="list-style-type: none"><li>• Windows 95 4.00.950 or later</li><li>• Windows 98</li></ul> <p><b>Note:</b> When the DB2 Application Development Client is installed, the JDK 1.1.8 is installed.</p>	<p>IPX/SPX, Named Pipes, NetBIOS, or TCP/IP</p> <ul style="list-style-type: none"><li>• The Windows 9x base operating system provides NetBIOS, IPX/SPX, TCP/IP, and Named Pipes connectivity. <b>Note:</b> IPX/SPX connectivity is only supported to Windows NT and Windows 2000 servers.</li><li>• If you plan to use LDAP (Lightweight Directory Access Protocol), you require either a Microsoft LDAP client or an IBM SecureWay LDAP client V3.1.1. For more information, refer to the <i>Administration Guide</i>.</li><li>• If you plan to use the Tivoli Storage Manager facilities for backup and restore of your databases, you require the Tivoli Storage Manager Client Version 3 or later.</li><li>• If you have the IBM Antivirus program installed on your operating system, it must be disabled or uninstalled to complete a DB2 installation.</li></ul>

Table 3. Software Requirements for Clients (continued)

Component	Hardware/Software Requirements	Communications
<ul style="list-style-type: none"> <li>• DB2 Run-Time Client for Windows</li> <li>• DB2 Administration Client for Windows</li> <li>• DB2 Application Development Client for Windows</li> </ul>	<ul style="list-style-type: none"> <li>• Windows NT Version 4.0 with Service Pack 3 or later</li> <li>• Windows Terminal Server (can only run the DB2 Run-Time Client)</li> <li>• Windows 2000</li> </ul> <p><b>Note:</b> When the DB2 Application Development Client is installed, the JDK 1.1.8 is installed.</p>	<p>APPC, IPX/SPX, Named Pipes, NetBIOS, or TCP/IP</p> <ul style="list-style-type: none"> <li>• The Windows NT and Windows 2000 base operating systems provide NetBIOS, IPX/SPX, TCP/IP, and Named Pipes connectivity.</li> <li>• For APPC connectivity, you require one of the following products:               <ul style="list-style-type: none"> <li>– IBM eNetwork Communications Server for Windows V5.01 or later.</li> <li>– Windows 2000: IBM eNetwork Personal Communications for Windows Version 4.3 CSD2 or later</li> <li>– Windows NT: IBM eNetwork Personal Communications for Windows Version 4.2 or later</li> <li>– Microsoft SNA Server Version 3 Service Pack 3 or later</li> <li>– Wall Data Rumba</li> </ul> </li> <li>• If you plan to use DCE (Distributed Computing Environment) and connect to a DB2 for OS/390 V5.1 database, that database must be enabled for DCE support using OS/390 DCE Base Services Version 3.</li> <li>• If you plan to use LDAP (Lightweight Directory Access Protocol), you require either a Microsoft LDAP client or an IBM SecureWay LDAP client V3.1.1. For more information, refer to the <i>Administration Guide</i>.</li> <li>• If you plan to use the Tivoli Storage Manager facilities for backup and restore of your databases, you require the Tivoli Storage Manager Client Version 3 or later.</li> <li>• If you have the IBM Antivirus program installed on your operating system, it must be disabled or uninstalled to complete a DB2 installation.</li> </ul>

### Possible Client-to-Server Connectivity Scenarios

The following table shows the communication protocols that can be used when connecting a specific DB2 client to a specific DB2 server. DB2 Workgroup, DB2 Enterprise, and DB2 Enterprise - Extended Editions can service requests from host or AS/400 clients (DRDA ARs).

Table 4. Possible Client-to-Server Connectivity Scenarios

Client	Server						
	AIX	HP-UX	Linux	OS/2	PTX/NUMA-Q	Solaris	Windows NT/ Windows 2000
AS/400 V4R1	APPC	N/A	N/A	APPC	N/A	APPC	APPC
AS/400 V4R2	APPC TCP/IP	TCP/IP	TCP/IP	APPC TCP/IP	TCP/IP	APPC TCP/IP	APPC TCP/IP
AIX	APPC TCP/IP	TCP/IP	TCP/IP	APPC TCP/IP	TCP/IP	APPC TCP/IP	APPC TCP/IP
HP-UX	APPC TCP/IP	TCP/IP	TCP/IP	APPC TCP/IP	TCP/IP	APPC TCP/IP	APPC TCP/IP
Linux	TCP/IP	TCP/IP	TCP/IP	TCP/IP	TCP/IP	TCP/IP	TCP/IP
MVS	APPC	N/A	N/A	APPC	N/A	APPC	APPC
OS/2	APPC IPX/SPX(1),(2) TCP/IP	TCP/IP	TCP/IP	APPC IPX/SPX(1),(2) NetBIOS TCP/IP	TCP/IP	APPC IPX/SPX(1) TCP/IP	APPC IPX/SPX(1) NetBIOS TCP/IP
OS/390	APPC TCP/IP	TCP/IP	TCP/IP	APPC TCP/IP	TCP/IP	APPC TCP/IP	APPC TCP/IP
PTX/NUMA-Q	TCP/IP	TCP/IP	TCP/IP	TCP/IP	TCP/IP	TCP/IP	TCP/IP
Silicon Graphics IRIX	TCP/IP	TCP/IP	TCP/IP	TCP/IP	TCP/IP	TCP/IP	TCP/IP
SQL/DS	APPC	N/A	N/A	APPC	N/A	APPC	APPC
Solaris	APPC TCP/IP	TCP/IP	TCP/IP	APPC TCP/IP	TCP/IP	APPC TCP/IP	APPC TCP/IP
VSE & VM V5	APPC	N/A	N/A	APPC	N/A	APPC	APPC
VSE V6	APPC	N/A	N/A	APPC	N/A	APPC	APPC
VM V6	APPC TCP/IP	TCP/IP	TCP/IP	APPC TCP/IP	TCP/IP	APPC TCP/IP	APPC TCP/IP
Windows 9x	TCP/IP	TCP/IP	TCP/IP	NetBIOS TCP/IP	TCP/IP	TCP/IP	IPX/SPX(1) NPIPE NetBIOS TCP/IP
Windows NT/ Windows 2000	APPC IPX/SPX(1) TCP/IP	TCP/IP	TCP/IP	APPC IPX/SPX(1) NetBIOS TCP/IP	TCP/IP	APPC IPX/SPX(1) TCP/IP	APPC IPX/SPX(1) NPIPE NetBIOS TCP/IP

1. Direct Addressing  
2. File Server Addressing

## Migrating from Previous Versions of DB2

If you are migrating from Version 5.0 or Version 5.2 to DB2 Version 7, you must prepare your databases and instances before installing DB2 Version 7. If you are upgrading from Version 6, there is nothing you have to do prior to installing Version 7.

The migration of pre-Version 5.0 databases and instances is not supported in DB2 Version 7.

When you migrate from a Version 5.x or Version 6 database system to a DB2 Enterprise - Extended Edition Version 7 database system, depending on the

installation option you choose, the installation program will migrate all instances to the Version 7 multipartition or single-partition format. If you are migrating a database with populated tables, refer to the *Administration Guide* for descriptions of how to add database partition servers to a system and redistribute data across database partitions.

After you install DB2 Enterprise - Extended Edition, any databases that you created with earlier versions of DB2 will not automatically work with the new version of DB2. You must follow the steps in this section to migrate your instances and databases.

Windows allows only one version of DB2 to be installed on a machine. For example, if you have DB2 Version 6 and install Version 7, Version 6 will be deleted during the installation.

## Preparing Databases and Instances for Migration

This section describes how to prepare existing DB2 Version 5.x databases and instances for migration to a format usable by DB2 Version 7. If you want to migrate more than one instance, you must repeat these steps for each instance.

To prepare your databases for migration:

1. Ensure that there are no applications using any databases owned by the DB2 instance you are preparing to migrate. To get a list of all applications owned by the instance, enter the **db2 list applications** command. If all applications are disconnected, this command will return the following message:

```
SQL1611W No data was returned by the Database System Monitor.  
SQLSTATE=00000
```

You can end a session by entering the **db2 terminate** command.

2. Ensure that all databases are cataloged. To view a list of all the cataloged databases in the current instance, enter the following command:  

```
db2 list database directory
```
3. Make a backup copy of all Version 5.x databases. You do not have to back up Version 6 databases. Refer to the *Administration Guide* for your DB2 product for information on making a backup copy of a database and to the *Command Reference* for the syntax of the backup command.
4. When all applications are complete and you have backed up your databases, stop all database server processes owned by the DB2 instance by entering the **db2stop** command.
5. Stop the DB2 license daemon by entering the **db2licd -end** command.
6. Stop all command line processor sessions by entering the **db2 terminate** command in each session that was running the command line processor.

Next, you should verify that all cataloged databases are ready for migration before installing DB2 Version 7.

## Verifying Databases are Ready for Migration

To ensure that you can migrate your databases to the DB2 Version 7 format, you should run the **db2ckmig** command before installing DB2 Version 7.

**Note:** This section applies to databases created with DB2 Version 5.x only. You do not have to run this command against databases created with DB2 V6.

To run the **db2ckmig** command:

1. Insert your DB2 Version 7 product CD-ROM into the drive.
2. Change to the `x:\db2\common` directory where `x:` represents your CD-ROM drive letter.
3. Enter the **db2ckmig** command to verify that the databases on your system can be correctly migrated. The syntax of the command is as follows:

### DB2CKMIG Command

```
► db2ckmig database_alias [/i drive:\path\filename]  
                |  
                +-----/e  
  
                |  
                +-----/u userid /p password |  
                |
```

*database\_alias* Specifies a *database\_alias* name of a database to be verified for migration. This parameter is required if the */e* parameter is not specified.

*/e* Specifies that all cataloged databases are to be verified for migration. This parameter is required if the *database\_alias* parameter is not specified.

*/l drive:\path\filename* Specifies a drive, target path and filename to keep a list of errors and warnings generated for the scanned database. The *path* variable is optional; if you do not specify a path, the path from which you execute the **db2ckmig** command will be used. You must specify a *filename*.

*/u userid* Specifies the user account used to connect to the database. This parameter must be specified if you are logged on as a user without connect authority.

*/p password* Specifies the password of the user account used to connect to the database. This parameter must be specified if you are logged on as a user without connect authority.



You can enter the **db2ckmig** command on remote systems. The database parameter must specify the *database\_alias* name of the remote database. You can run this command from any database partition server. The file will be written on your local system.

For example, to check that all databases cataloged on your system can be migrated and to log all the messages from this command to the `c:\temp\message.txt` file, enter the following command:

```
x:\db2\common\db2ckmig /e /l c:\temp\message.txt
```

where *x*: represents your CD-ROM drive.

4. If any errors are found, the **db2ckmig** command generates a log file and places it in the path and file specified by the `/l` option. If there are errors, refer to the information that follows for suggested corrective actions. Once the errors have been corrected, re-enter the **db2ckmig** command to ensure that the databases are ready to be migrated.
5. Make a backup copy of the database. For more information, refer to the *Administration Guide*.

### **Possible DB2CKMIG Error Conditions**

#### **A database is in backup pending state**

Perform a backup of the database.

#### **A database is in roll-forward pending state**

Recover the database as required. Perform or resume a roll-forward database to end of logs and stop.

#### **Table space ID is not in normal state**

Recover the database and table space as required. Perform or resume a roll-forward database to end of logs and stop.

#### **A database is in an inconsistent state**

Restart the database to return it to a consistent state.

#### **Structured type and function have the same name**

A structured type and function (with no arguments) belonging to the same schema cannot have the same name. The type or function, and objects using the type or function, must be dropped and recreated using another name. To correct this error:

1. Back up the database.
2. Export the data from any tables that are dependent on the structured types or functions.

3. Drop any tables dependent on the structured types or functions, and then drop the structured types or functions. These drops may drop other objects such as views, indexes, triggers, or functions.
4. Create structured types or functions with different type or function names and recreate the tables using the new structured type or function names. Recreate any dropped views, indexes, triggers, or functions.
5. Import or load the data into the object.

**The database contains user-defined distinct types (UDTs) that use the type name BIGINT, DATALINK, or REFERENCE.**

These data type names are reserved for the Version 7 database manager. To correct this error:

1. Back up the database.
2. Export the data from any tables that are dependent on these data types.
3. Drop any tables dependent on these data types, and then drop the data types. These drops may drop other objects such as views, indexes, triggers, or functions.
4. Create data types with different type or function names and recreate the tables using the new data type names. Recreate any dropped views, indexes, triggers, or functions.
5. Import or load the data into the object.

Refer to the *Administration Guide* for more information about the actions required to correct these conditions.

---

## **Migrating from Visual Warehouse Version 3.1 or Visual Warehouse Version 5.2**

The function provided by the Visual Warehouse product has been integrated into DB2 Universal Database. This function is accessed in DB2 Universal Database through the Data Warehouse Center.

When you install DB2 Universal Database on Windows, all previous versions of Visual Warehouse components on that system are uninstalled. To ensure that you continue to have a working environment after the installation, all warehouse components that interoperate in your environment must be at the Version 7 level. This includes:

- The server
- Warehouse control databases
- Warehouse agents (local and remote)
- Warehouse transformers
- The Information Catalog Manager (DataGuide\*)

A partial migration of Visual Warehouse is not allowed.

**Note:** Any settings that you have customized for Visual Warehouse will be lost during the migration process. For example, you might have a customized setting in the `IWH.environment` file for an AIX or Solaris Operating Environment warehouse agent. Before you begin the installation process, make a backup copy of any Visual Warehouse settings that you have customized.

## **Migrating the Visual Warehouse Server and Active Warehouse Control Database**

The Visual Warehouse server and your active warehouse control database (the most recently used database) are migrated together during the installation process. After you install DB2 Universal Database, you must manually migrate any other warehouse control databases that you want to continue to use with the Data Warehouse Center.

Refer to *Data Warehouse Center Administration Guide* for:

- Information about migrating warehouse control databases
- Information about how the structure of metadata tables has changed from Visual Warehouse to Data Warehouse Center

## **Migrating Warehouse Agents**

When a warehouse server is installed, a default warehouse agent is also installed. However, you must migrate any remote warehouse agents that you want to continue to use after the installation completes. Warehouse agents are available with the DB2 Warehouse Manager product.

To migrate a warehouse agent:

1. Remove the existing warehouse agent.
2. Install the DB2 Warehouse Manager warehouse agent.

For information on removing and installing Visual Warehouse agents, see the *DB2 Warehouse Manager Installation Guide*.

## **Migrating Warehouse Transformers**

Warehouse transformers are available with the DB2 Warehouse Manager product. After installing the warehouse transformers on the warehouse agent system, you must enable the warehouse transformers so that they can be used with new warehouse targets. This is done using the Data Warehouse Center.

For more information on enabling warehouse transformers, see the *DB2 Warehouse Manager Installation Guide*.

## Migrating DataGuide

**Note:** DataGuide is known as the Information Catalog Manager in DB2 Universal Database Version 7. It is available with the DB2 Warehouse Manager product.

When you install DB2 Universal Database Version 7, any DataGuide components that are on that system are uninstalled. If you then install the Information Catalog Manager as a component of the DB2 Warehouse Manager, the new Information Catalog Manager components will be installed and the DataGuide information catalog database can be upgraded. For information about upgrading DataGuide information catalog databases, refer to the *Information Catalog Manager Administration Guide*.

**Note:** If you do not upgrade to the Information Catalog Manager at this time, you will not be able to use DataGuide with the Data Warehouse Center after the installation. DataGuide components on other systems will continue to work with the DataGuide information catalog database if it has not been upgraded by a DB2 Warehouse Manager installation.

---

## NetQuestion Search System

If you installed online product documentation with your previous version of DB2 for Windows, OS/2, AIX, HP-UX or Solaris, or if you installed another IBM product like VisualAge C++ or VisualAge for Java, an online search system called NetQuestion was also installed automatically.

If the version of NetQuestion that comes with DB2 Version 7 is later than the version of NetQuestion currently present on your system, the current version will be upgraded and any existing document indices will be re-registered with NetQuestion. This will be done automatically for you during DB2 installation.

For more information on NetQuestion, refer to the *Installation and Configuration Supplement*.

---

## Stopping the Windows Cluster Service

If you are running in the Microsoft Cluster Server (MSCS) environment and the DB2 server was configured to run with MSCS, you need to stop the MSCS Cluster Service before installation by issuing the **net stop clussvc** command.

**Notes:**

1. If DB2 is running on the current machine, it may fail-over and run on another machine in the MSCS cluster.
2. You need to upgrade all machines in the MSCS cluster before you migrate the database(s).

---

## Your next step

After you have determined that your system meets all hardware and software requirements, and after you have prepared any existing databases and instances for migration, you can now install DB2 Version 7 using either the interactive or distributed method. For installation procedures, see the following sections:

- “Chapter 3. Installing DB2 Universal Database Enterprise - Extended Edition on Windows” on page 29 for interactive installation.
- The *Installation and Configuration Supplement* for distributed installation.



---

## Chapter 3. Installing DB2 Universal Database Enterprise - Extended Edition on Windows



If you are migrating from a single-partition database system, or a previous release of this product, you must complete certain procedures before installing DB2 Universal Database Enterprise - Extended Edition Version 7. Go to “Migrating from Previous Versions of DB2” on page 20 for more information.

---

You will perform the following steps to create a partitioned database system:

1. You will install the instance-owning database partition server on one machine (physical node) using the setup program, and create (or migrate) a partitioned database instance.
2. On another machine, you will use the setup program to install a database partition server that will participate in the instance. You may perform this step several times, depending upon the number of machines (physical nodes) participating in your partitioned database system.

If you want to install a DB2 Client, go to “Chapter 6. Installing DB2 Clients” on page 53.

For information on how to deploy this product using a distributed installation, refer to the *Installation and Configuration Supplement*.

---

### Before You Begin

Before you begin the installation, be sure that you have the following items and information:

1. Ensure that all of the machines that will participate in this partitioned database system:
  - a. Meet all the requirements outlined in “Chapter 2. Planning for Installation” on page 15.
  - b. Belong to the same Windows domain.
  - c. Have consistent time and date settings.



To be considered consistent, the difference in GMT time between all machines that participate in a partitioned database system must be within 1 hour. You can use the *max\_time\_diff* configuration parameter to change this restriction. For more information, refer to the *Administration Guide*.

- d. Have a TCP/IP port range free that spans 4 port numbers. This port range will be used by the Fast Communications Manager (FCM). FCM is a feature of DB2 that handles communications between database partition servers. The port range you select must be available on every machine that will participate in the partitioned database system. By default, the setup program will generate a range for you, use this port range or provide your own.

To review the TCP/IP ports that are already in use, open the services file located in the `x:\winnt\system32\drivers\etc` directory, where `x`: represents the drive where you installed Windows.

Make note of this port range here: \_\_\_\_\_ to \_\_\_\_\_.



If you plan to run multiple logical nodes (MLNs) on any of the machines that will participate in your partitioned database system, then you must specify a port range that spans the number of MLNs on the machine running the most MLNs.

For example, if you had a database partition server that was running 6 database managers (by installing a database partition server and then adding 5 logical nodes) you would need to specify 6 consecutive ports as your port range for the installation. For more information, refer to the *Administration Guide*.

- e. Have a TCP/IP port free for use by the DB2 Performance Monitor. The port you select must be available on every machine that will participate in the partitioned database system. By default, the setup program will generate a value for you. You can use this port or provide your own.

To review the TCP/IP ports that are already in use on a machine, open the services file located in the `x:\winnt\system32\drivers\etc` directory, where `x`: represents the drive where you installed Windows.

Make note of this port here: \_\_\_\_\_.

- f. Can communicate with each other using TCP/IP. To ensure that two machines can communicate with each other using TCP/IP, perform the following steps:

Step 1) On a machine that will participate in this partitioned database system, enter the following command:

```
hostname
```

This command will return the *hostname* of this machine

Step 2) On another machine that will participate in this partitioned database system, enter the following command:

```
ping hostname
```

You will receive output similar to the following:



Pinging myserver.com [9.21.27.230] with 32 bytes of data:

```
Reply from 9.21.27.230: bytes=32 time<10ms TTL=128
Reply from 9.21.27.230: bytes=32 time<10ms TTL=128
```

This output verifies that both machines can communicate with each other.



If you are planning to use multiple network adapters, you must use the `db2nchg` command to specify the TCP/IP address for the netname field in the `db2nodes.cfg` file. For more information, you can refer to the *Administration Guide*, or to “Changing the Database Partition Server Configuration” on page 37.

2. A domain user account that belongs to the *Local Administrators* group on the machine where you are going to perform the installation. The domain user account you specify must be defined as a Local Administrator on *every* machine that will participate in the partitioned database system.  
This user account must also have the “*Act as part of the operating system*” advanced user right. For more information on granting advanced user rights, refer to “Granting Advanced User Rights on Windows” on page 94.
3. A domain user account to be associated with the DB2 instance. This domain user account does not need to be given Administrator or Domain Administration authority.
4. During installation, you will be asked to provide a user account that will be used by the Administration Server to log on to the system and to start itself as a service.  
By default, the setup program will create a user account with the user name `db2admin` and the password you supply. You can accept the default user account, create your own user account by modifying the default value, or provide your own user account. If this account already exists on your system, you must use the password that was previously set for this user account. If you provide your own user account, you must ensure that it conforms to DB2’s naming rules. For more information, see “Appendix C. Naming Rules” on page 101.
5. To verify that DB2 installed correctly, you will need to have a user account that belongs to the DB2 System Administrative (SYSADM) group. The account name must comply with DB2’s naming rules, as described in “Username, User ID, Group Name, and Instance Name Rules” on page 103.  
By default, any local user account that belongs to the *Local Administrators* group, on the local machine where the account is defined, has SYSADM authority on the instance. For more information, see “Working with the System Administrative Group” on page 93. For more information on valid DB2 usernames, see “Appendix C. Naming Rules” on page 101.

6. During the installation, you may also be prompted for user names and passwords for the products and services that you install.
7. If you want to use LDAP with Windows 2000, you must extend the directory schema to contain DB2 object classes and attribute definitions. You must do this once, before you install any DB2 product.

To extend the directory schema, execute the **db2schex.exe** program from the installation CD with Schema Admins authority. You can execute this program with Schema Admins authority, without logging off and logging on again, as follows:

```
runas /user:MyDomain\Administrator x:\db2\common\db2schex.exe
```

where **x:** represents the CD-ROM letter. When **db2schex.exe** completes, you can continue with the installation.

---

## Performing the Installation

You will first install the instance-owning database partition server on one machine, and then install a database partition server on each of the other machines that will participate in your partitioned database system.

Use the following steps to install an instance-owning database partition server, or a database partition server that will participate in your partitioned database system:

- Step 1. Log on to the system with the domain user account that you will use to perform the installation. This is the domain user account that has Local Administrator authority on each machine participating in the partitioned database system.
- Step 2. Shut down any other programs so that the setup program can update files as required.
- Step 3. Insert the CD-ROM into the drive. The auto-run feature automatically starts the setup program. The setup program will determine the system language, and launch the setup program for that language. If you want to run the setup program in a different language, or the setup program failed to start, see the following tip:



To manually invoke the setup program, perform the following steps:

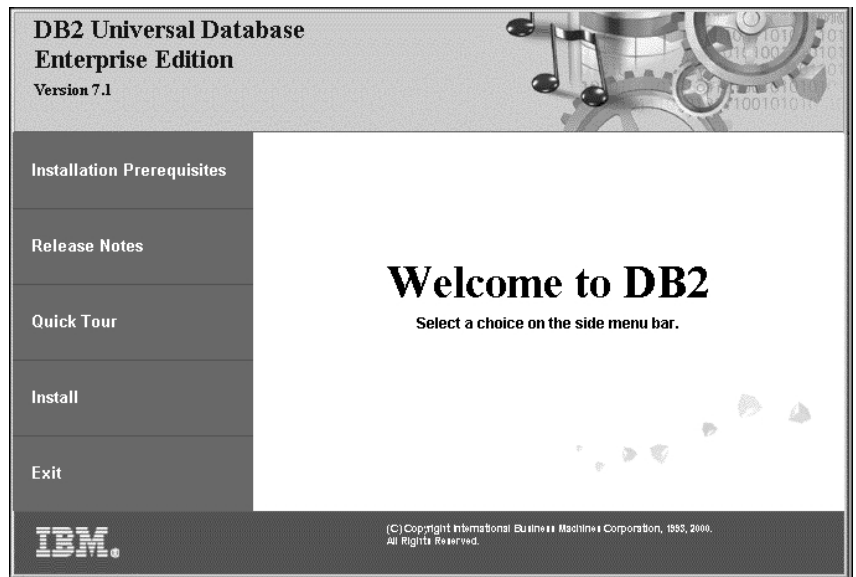
- a. Click **Start** and select the **Run** option.
- b. In the **Open** field, enter the following command:  
`x:\setup /i language`

where:

- *x*: represents your CD-ROM drive
- *language* represents the country code for your language (for example, EN for English). Table 8 on page 99 lists the code for each available language.

- c. Click **OK**.

Step 4. The DB2 Launchpad opens. It looks similar to the following:



Step 5. From the DB2 Launchpad, you can view the Installation Prerequisites and the Release Notes, you can take a Quick Tour to explore the features, capabilities, and benefits of DB2 Universal Database Version 7, or you can proceed directly to the installation.

Step 6. Once you have initiated the installation from the DB2 Launchpad, proceed by following the setup program's prompts. Online help is available to guide you through the remaining steps. Invoke online help by clicking **Help** or pressing **F1** at any time. You can click **Cancel** at any time to end the installation.



To install the *instance-owning database partition server*, you must ensure that you select the **This machine will be the instance-owning database partition server** option in the Select Installation Option window.

If you are using the setup program to install a *database partition server* that will participate in an instance, you must ensure that you select the **This machine will be a new node on an existing partitioned database system** option in the Select Installation Option window.



If you have just finished installing the instance-owning database partition server, return to Step 1 to install database partition servers on all other machines that will participate in the instance you just created.

Remember to wait for the installation to complete on one machine, and to restart the machine before installing on the next machine in your cluster.

If you have finished running the setup program on two or more machines, creating both the instance-owning database partition server and at least one database partition server to participate in the instance, you are finished creating a partitioned database system.

The installation program has:

- Created DB2 program groups and items (or shortcuts).
- Registered a security service.
- Reserved the port that you specified for the DB2 Performance Monitor in the services file with the service name `db2ccmsv`.
- Reserved the port range that you specified for FCM in the services file with the service names `DB2_DB2MPP` for the first port number and `DB2_DB2MPP_END` for the last port number.
- Updated the Windows registry.
- Created a default instance named `DB2MPP` and added it as a service (if you selected to create a default instance).

The instance directory is located in the `\sql11ib` directory. During the installation, the DB2 setup program gives the instance directory the shared name `DB2-instance_name`. Read and write permissions are automatically granted to everyone in the domain. After completing the installation, you can change the permissions to restrict access to the directory.

If the setup program detected any pre-Version 7 single-partition instances, they were migrated to the Version 7 single-partition format. If the setup program detected any pre-Version 7 multi-partition database instances, they were migrated to the Version 7 multi-partition format.

- Created the DB2 Administration Server.



For information on errors encountered during installation, see the `db2.log` file. The `db2.log` file stores general information and error messages resulting from the install and uninstall activities. By default, the `db2.log` file is located in the `x:\db2log` directory, where `x`: represents the drive on which your operating system is installed.

For more information, refer to the *Troubleshooting Guide*.

### Your Next Step:

- You can use the **db2ncrt** command to add a logical database partition server to your instance, thus creating a multiple logical node (MLN) configuration. To add a logical database partition server to your instance, refer to “Adding a Logical Database Partition Server to an Instance”.
- If you do not want to add a logical database partition server to your instance, you can verify your installation by following the steps outlined in “Verifying the Installation” on page 39.

---

## Adding a Logical Database Partition Server to an Instance

Once you have created your partitioned database system by installing the instance-owning database partition server on one machine, and installing a database partition server on the other machines participating in the partitioned database system, you have the option of adding additional database partition server nodes (logical nodes) to your instance to create a multiple logical node (MLN) configuration.

Use the **db2ncrt** command to add a database partition server (logical node) to an instance. Adding a database partition server to a machine where an instance already exists using the **db2ncrt** command adds a database partition server as a logical node to the machine.



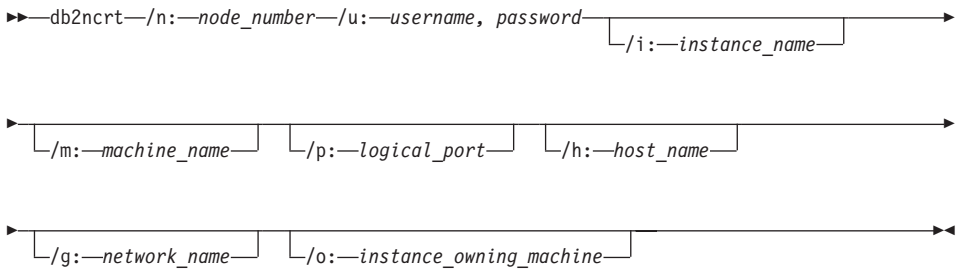
Do not use the **db2ncrt** command if there are databases in an instance, or if you are migrating databases from a previous version of DB2.

Instead, you should use the **db2start addnode nodenum** command to scale your system after you have completed the post-installation migration tasks. This ensures that the database is correctly added to the new database partition server. If you want to add a node to an instance in which a database has been created, refer to the Administration Guide.

Do not edit the `db2nodes.cfg` file, since changing the file may cause inconsistencies in the partitioned database system.

The syntax of the **db2ncrt** command is as follows:

### **db2ncrt command**



where:

#### */n:node\_number*

Is a unique node number to identify the database partition server. This number can be from 1 to 999. This parameter is required.

#### */u:domain\_name\username, password*

Is the domain, logon account name, and password for the DB2 service that represents the database partition server. This parameter is required.

#### */i:instance\_name*

Is the instance name. This parameter is optional. If you do not specify this parameter, the default is the current instance.

#### */m:machine\_name*

Is the computer name of the Windows workstation on which the node resides. This parameter is required if you add the database partition server on a remote machine.

#### */p:logical\_port*

Specifies the logical port number used for the database partition server. This parameter is optional. If you do not specify this parameter, the logical port number assigned will be 0.

If you are creating a logical node, you must specify this parameter and select a logical port number that is not in use. Note the following restrictions:

- On every machine, there must be a database partition server that has a logical port 0.
- The port number cannot exceed the port range reserved for FCM communications in the services file in `x:\winnt\system32\drivers\etc\` directory. For example, if you

reserve a range of 4 ports for the current instance, then the maximum port number would be 3 (ports 1, 2, and 3; port 0 is for the default logical node).

**/h:host\_name**

Specifies the TCP/IP host name that is used by FCM for internal communications (the port is not a node). This parameter is required if you add the database partition server on a remote machine.

**/g:network\_name**

Specifies the network name for the database partition server. This parameter is optional.

Use this parameter if you have multiple IP addresses on a machine and you want to use a specific IP address for the database partition server. You can enter the *network\_name* parameter using the network name or IP address.

**/o:instance\_owning\_machine**

Is the computer name of the instance-owning machine; the default is the local machine. This parameter is required when the **db2nrcr** command is invoked on any machine that is not the instance-owning machine.

For example, if you want to add a new database partition server to the instance TESTMPP (so that you are running multiple logical nodes) on the instance-owning machine SHAYER, and you want this new node to be known as node 2 using logical port 1, enter the following command:

```
db2nrcr /n:2 /u:QBPAULZ\paulz,pas1wrđ /i:TESTMPP /m:TEST /p:1 /o:SHAYER
```

## Changing the Database Partition Server Configuration

To re-configure a database partition server, you can use the **db2nchg** command. With the **db2nchg** command you can do the following:

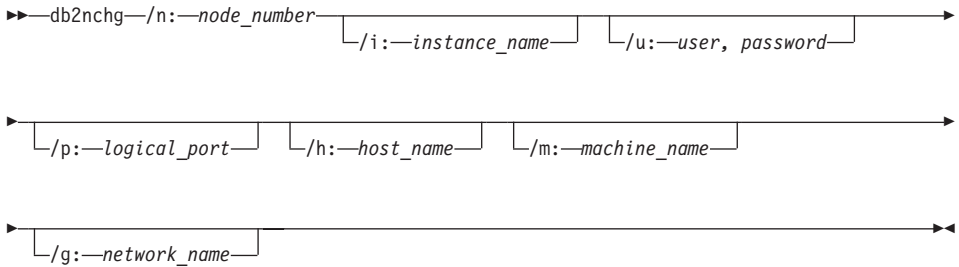
- Move the database partition server (node) from one machine to another.
- Change the TCP/IP host name of the machine.
- Use a different logical port number.
- Use a different network name for the database partition server (node).



You can only use this command if the database partition server is stopped.

The syntax of the **db2nchg** command is as follows:

### **db2nchg** command



#### */n:node\_number*

Specifies the node number of the database partition server's configuration you want to change. This parameter is required.

#### */i :instance\_name*

Specifies the instance that this database partition server participates in. This parameter is optional. If you do not specify this parameter, the default is the current instance.

#### */u:domain\_name\username, password*

Specify this option to change the domain, logon account name, and password for the DB2 service that represents the database partition server. This parameter is optional. If you do not specify this parameter, the domain, logon account, and password will remain the same.

#### */p:logical\_port*

Specify this option to change the logical port for the database partition server. This parameter must be specified if you move the database partition server to a different machine. If you do not specify this parameter, the logical port number will remain unchanged.

#### */h:host\_name*

Specify this option to change the TCP/IP hostname used by FCM for internal communications. This parameter is optional. If you did not specify this parameter, the host name is unchanged.

#### */m:machine\_name*

Specify this option to move the database partition server to another machine. The database partition server can only be moved if there are no existing databases in the instance. This parameter is optional.

#### */g:network\_name*

Specify this option to change the network name for the database partition server. This parameter is optional.



Use this parameter if you have multiple IP addresses on a machine and you want to use a specific IP address for the database partition server. You can enter the *network\_name* using the network name or IP address.

For example, to change the logical port assigned to node 2, which participates in the instance TESTMPP, to use the logical port 3, enter the following command: `db2nchg /n:2 /i:TESTMPP /p:3`

---

## Verifying the Installation



If you are verifying the installation of a DB2 Enterprise - Extended Edition server, ensure that you perform these steps on the *instance-owning database partition server*

You can verify that DB2 is installed correctly by creating the DB2 UDB SAMPLE database on your system, and accessing data from the database.

If you have installed First Steps, you can verify the installation by performing the following steps.

- Step 1. Log on to the instance-owning database partition server with the user account that you created to verify the installation.
- Step 2. Start First Steps. For more information, see “Starting First Steps” on page 89.
- Step 3. Select **Create SAMPLE databases** on the First Steps Launchpad to open the Create SAMPLE databases window.
- Step 4. In the Create SAMPLE databases window, select all of the following databases that you want to create:
  - **DB2 UDB Sample database**
  - **Data Warehousing Sample database**
  - **OLAP Sample database**

You use the DB2 UDB sample database to verify the installation. You need the other two databases if you want to run the Business Intelligence tutorial. For more information about these databases, refer to *Data Warehouse Center Administration Guide* and *OLAP Setup and User's Guide*.

**Note:** The Data Warehousing and OLAP sample databases only apply if you installed the Base Warehouse and OLAP Starter Kit components.

- Step 5. Click **OK**.

You must ensure that the drive where the SAMPLE database will be created exists on every machine that participates in your partitioned database system. By default, the SAMPLE database is created on the drive where DB2 is installed.

This command may take a few minutes to process. For a detailed description of the contents of the SAMPLE database, refer to the *Administration Guide*. When the SAMPLE database has been created, you will receive a completion message. Click **OK**.

- Step 6. Once the database is created, select **Work with the SAMPLE databases** on the First Steps Launchpad to start the Control Center. The Control Center allows you to perform administration tasks on different instance and database objects.

In the left pane of the Control Center screen, expand the object tree to view the SAMPLE database and SAMPLE database objects. Select the Tables object to view the SAMPLE database tables in the right pane of the Control Center screen. For more information, see “Administering Instances and Databases with the DB2 Administration Tools” on page 122.

To verify that the Data Warehouse component installed correctly, select the Data Warehouse Center icon.

- Step 7. To verify that the OLAP Starter Kit component installed correctly, click **Start** → **Programs** → **IBM DB2** → **OLAP** → **OLAP Desktop**.

If you did not install the DB2 tools, you can verify the installation by creating the SAMPLE database and connecting to it using the command line processor as follows:

- Step 1. Log on to the instance-owning database partition server with the DB2 user account that you created to verify the installation.
- Step 2. Enter the **db2samp1** command to create the SAMPLE database.

By default, the SAMPLE database will be created on the drive where DB2 was installed; however, you can specify a drive on which to create this database. For example, to create the SAMPLE database on the F: drive, enter the following command:

```
db2samp1 F:
```

This command may take a few minutes to process. For a detailed description of the contents of the SAMPLE database, refer to the *Administration Guide*. There is no completion message; when the command prompt returns, the process is complete.

You must ensure that the drive where the SAMPLE database will be created exists on every machine that participates in your partitioned database system.

The `SAMPLE` database is automatically cataloged with the database alias `SAMPLE` when it is created.

Enter the **`dwcsmp`** and **`olapsmp`** commands to create data warehousing and OLAP sample databases. When prompted, provide a user ID and password. You need these databases if you want to run the Business Intelligence tutorial. For more information about these databases, see the Data Warehouse Center and OLAP manuals.

**Step 3.** Start the database manager by entering the **`db2start`** command.

**Step 4.** Enter the following DB2 commands from a DB2 Command Window to connect to the `SAMPLE` database, retrieve a list of all the employees that work in department 20, and reset the database connection:

```
db2 connect to sample
db2 "select * from staff where dept = 20"
db2 connect reset
```

For information about entering DB2 commands, see “Appendix A. Basic Task Knowledge” on page 89.



After you have verified the installation, you can remove the `SAMPLE` database to free up disk space. Enter the **`db2 drop database sample`** command to drop the `SAMPLE` database.



---

## Chapter 4. DB2 Post-Installation Migration Tasks

After installing DB2 Version 7, you can now migrate databases and complete other migration activities. Do not scale the database system before migrating all the databases. If you scale the database system, migration will fail. You should perform the following steps after DB2 is installed to complete the migration process:

If you are migrating from . . .	. . . perform these actions:
DB2 Universal Database Version 5.x Enterprise - Extended Edition	<ol style="list-style-type: none"><li data-bbox="779 529 1235 624">1. Migrate your databases. See "Migrating Databases" for more information.</li><li data-bbox="779 633 1235 815">2. Update all instances that you want to use in a partitioned database system. See "Updating Single-Partition Instances for use in Partitioned Database Systems" on page 45 for details.</li></ol>
DB2 Universal Database Version 6	Update all instances that you want to use as partitioned database systems.  See "Updating Single-Partition Instances for use in Partitioned Database Systems" on page 45 for details.

---

### Migrating Databases

This section applies only if you are migrating databases from Version 5.x to the Version 7.1 format. If you are migrating from Version 6.1 to Version 7.1, you can skip this section.

If you are updating a single-partition Version 7 instance to the Version 7 multipartition format, you do not need to migrate the databases in that instance.

When pre-Version 7 databases are migrated for use in a Version 7 partitioned database system, default nodegroups are created for them. For more information about default nodegroups and how they are used, see "Nodegroups and Data Partitioning" on page 7.

To migrate Version 5.x databases owned by an instance:



## Updating Single-Partition Instances for use in Partitioned Database Systems

You can update any single-partition instances that you want to use in a partitioned database system. To do this, use the **db2iupdt** command. The syntax of the **db2iupdt** command is as follows:

```
▶▶ db2iupdt -instance_name -/u:domain_name\username,password [ /p:instance_profile_path ]
▶ [ /r:baseport,endport ] [ /h:hostname ] ▶▶
```

where:

*instance\_name*

Specifies the name of the instance to be updated. This parameter is required.

*/u:domain\_name\username,password*

Specifies the domain user account for the DB2 service. This parameter is required.

*/p:instance\_profile\_path*

Specifies the instance profile path. If you do not specify the instance profile path, the instance directory is created in the \sql11ib directory, and given the shared name DB2-*instance\_name*. Read and write permissions will automatically be granted to everyone in the domain. You can change the permissions to restrict access to that directory. This parameter is optional.

*/r:baseport,endport*

Specifies the range of TCP/IP ports to be used for FCM communications among database partition servers. This parameter is optional.

When this option is specified, the services file of the local machine is updated with entries similar to the following:

```
DB2_instance_name      12000/tcp
DB2_instance_name_END  12004/tcp
```

If you do not specify a valid port range, DB2 will select one for you.

*/h:hostname*

Overrides the default TCP/IP hostname if there is more than one TCP/IP hostname for the current machine. This parameter is optional.

---

## Optional Post Migration Actions

There are optional activities you may want to undertake following database migration. You can also apply these optional activities to a down-level database backup which is restored to Version 7, because at the end of the restore, the database is migrated to Version 7.

- **Migration of Unique Indexes**

**Note:** This section applies only to migration scenarios involving Version 5.x databases, or to Version 6 databases that were previously migrated from the Version 5.x format but whose unique indexes were *not* migrated using the **db2uiddl** command.

Version 5.x unique indexes are not automatically migrated to Version 7 semantics for several reasons:

- Converting unique indexes is a very time-consuming operation.
- You may have applications that depend on the previous version’s unique index semantics.
- You may want to manage the staged conversion of unique indexes on your own schedule, when needed, using the **db2uiddl** command.

All existing applications will continue to work even if the unique indexes are not converted to Version 7 semantics. You have to convert unique indexes to Version 7 semantics only if support for deferred uniqueness checking is required.

To convert unique indexes:

1. Log in with a user account that has SYSADM authority. See “Working with the System Administrative Group” on page 93 for more information.
2. Start the database manager by entering the **db2start** command.
3. Run the **db2uiddl** command against your migrated database. Refer to the *Command Reference* for the syntax of this command.

The **db2uiddl** command searches the database catalog tables and generates all the CREATE UNIQUE INDEX statements for user tables in an output file.

4. Review the output generated from the **db2uiddl** command. We recommend that you remove any unwanted indexes from the output file to reduce the time needed to execute it. Comments in the output will flag other situations that require your attention.
5. Connect to the database by entering the **db2 connect to** *database\_alias* command, where *database\_alias* is the alias of the database you are migrating.



6. Execute the commands in the output file, generated by the **db2uiddl** command, using a command similar to the following:

```
db2 -tvf filename
```

where *filename* is the name of the file generated by the **db2uiddl** command.

- **Update Statistics**

When database migration is completed, the old statistics that are used to optimize query performance are retained in the catalogs. However, DB2 Version 7 has statistics that are modified or do not exist in Version 5.x. To take advantage of these statistics, you may want to execute the **runstats** command on tables, particularly those tables that are critical to the performance of your SQL queries.

Refer to the *Command Reference* for the syntax of the **runstats** command. For details on the statistics, refer to the *Administration Guide*.

- **Rebind Packages**

During database migration, all existing packages are invalidated. After the migration process, each package is rebuilt when it is used for the first time by the Version 7 database manager.

For better performance we recommend that you run the **db2rbind** command to rebuild all packages stored in the database. Refer to the *Command Reference* for the syntax of this command.

- **Update database and database manager configuration**

Some of the database configuration parameters are changed to Version 7 defaults or to other values during database migration. Refer to the *Administration Guide* for more information about configuration parameters.

We recommend that you run the DB2 Performance Monitor for suggestions in choosing appropriate configuration parameters. For more information, refer to the *Administration Guide*.

- **Migrate Explain Tables**

To migrate the explain tables in a database that has been migrated to Version 7, run the following command:

```
db2exmig -d dbname -e explain_schema [-u userid password]
```

where:

- *dbname* represents the database name. This parameter is required.
- *explain\_schema* represents the schema name of the explain tables to be migrated. This parameter is required.
- *userid* and *password* represent the current user's ID and password. These parameters are optional.

The explain tables belonging to the user ID that is running **db2exmig**, or that is used to connect to the database, will be migrated. The explain tables

migration tool will rename the Version 5.x and Version 6 tables, create a new set of tables using the **EXPLAIN.DDL**, and copy the contents of the old tables to the new tables. Finally, it will drop the old tables. The **db2exmig** migration utility will preserve any user added columns on the explain tables.

---

## Chapter 5. Your Next Step

This section provides a road map to information that you may need after you have finished installing DB2 Enterprise Extended – Edition. It includes a list of tasks that are covered in this book, a detailed listing of the information provided in the DB2 Universal Database Administration Guide, and a table outlining task information found in other DB2 documentation.

The remaining chapters of this book provide information on the following tasks:

- “Chapter 6. Installing DB2 Clients” on page 53
- “Chapter 7. Installing DB2 Clients on Windows 32-Bit Operating Systems” on page 57
- “Chapter 8. Configuring Client-to-Server Communications Using the Client Configuration Assistant” on page 61
- “Chapter 9. Control Center Installation and Configuration” on page 73

---

### DB2 Universal Database Administration Guide

The DB2 Universal Database Administration Guide is divided into three volumes. Refer to the appropriate volume for information on the task you want to perform.

#### DB2 Universal Database Administration Guide: Planning

- Logical and physical database design
- Designing distributed databases
- Designing for transaction managers
- Designing for high availability
- High availability in the Windows environment

#### DB2 Universal Database Administration Guide: Implementation

- Administering DB2 using GUI tools
- Creating a database
- Altering a database
- Controlling database access
- Auditing DB2 activities
- Utilities for moving data
- Using distributed computing environment (DCE) directory services
- Issuing commands to multiple database partition servers

- How DB2 for Windows works with Windows security
- Configuring multiple logical nodes
- Lightweight Directory Access Protocol (LDAP) directory services

## **DB2 Universal Database Administration Guide: Performance**

- Tuning and application performance
- Tuning and configuring your system
- Scaling your configuration
- Redistributing data across database partitions
- Configuring DB2
- SQL Explain tools
- Configuring XA Transaction Managers to use DB2 UDB

---

### **Task Information in other DB2 Documentation**

*Table 5. Task Information in other DB2 Documentation*

<b>If you want to ...</b>	<b>Refer to the ...</b>
Install multiple DB2 clients without having to run the setup program at each DB2 client workstation	<i>DB2 Installation and Configuration Supplement</i>
Monitor a system's performance	<i>DB2 System Monitor Guide and Reference</i>
Load or export data from an existing DB2 database.	<i>DB2 Data Movement Utilities Guide and Reference</i>
Connect to Oracle data sources	<i>DB2 Installation and Configuration Supplement</i>
Perform replication activities	<i>DB2 Replication Guide and Reference</i>
Perform trouble shooting activities	<i>DB2 Troubleshooting Guide</i>
Use text, image, audio, video, or spatial extenders	<i>DB2 Text Extender Administration and Programming, DB2 Image, Audio, and Video Extenders Administration and Programming, DB2 Spatial Extenders User's Guide and Reference</i>
Access data that resides on a host or AS/400 DB2 database	<i>DB2 Connect Enterprise Edition for Windows Quick Beginnings</i>
Learn more about DB2 SQL	<i>DB2 SQL Reference, DB2 SQL Getting Started</i>

Refer to "Appendix E. Using the DB2 Library" on page 131 for more information about the DB2 product library.

---

## **Part 3. Installing and Configuring DB2 Clients**



---

## Chapter 6. Installing DB2 Clients

This section describes the different DB2 clients, and provides information on distributed installation, and thin client configuration.

**Note:** DB2 clients can connect to DB2 servers *two* releases later or *one* release earlier than the client's release level, as well as to servers at the same release level. For example, a DB2 Version 5.2 client can connect to DB2 Version 5.0, 5.2, 6.1, and 7.1 servers, while a DB2 Version 7.1 client can connect to DB2 Version 6.1 and 7.1 servers.

You may install a DB2 client on any number of workstations. For information about licensing, refer to your *License Information Booklet*.

You cannot create a database on a DB2 client, only connect to databases that reside on a DB2 server.



To go to the section that provides client installation instructions:

- “Chapter 7. Installing DB2 Clients on Windows 32-Bit Operating Systems” on page 57

For information on how to install Version 7 DB2 clients on other platforms, refer to the *Installation and Configuration Supplement*.

To download installation packages for DB2 clients supported on other platforms, and pre-Version 7 clients, connect to the IBM DB2 Client Application Enabler Web site at <http://www.ibm.com/software/data/db2/db2tech/clientpak.html>

---

---

### DB2 Run-Time Client

A DB2 Run-Time Client provides the ability for workstations from a variety of platforms to access DB2 databases.

DB2 Run-Time Clients are available for the following platforms: AIX, HP-UX, Linux, OS/2, NUMA-Q, Silicon Graphics IRIX, the Solaris Operating Environment, and Windows 32-bit operating systems.

---

## DB2 Administration Client

A DB2 Administration Client provides the ability for workstations from a variety of platforms to access and administer DB2 databases. The DB2 Administration Client has all the features of the DB2 Run-Time Client and also includes all the DB2 Administration tools, documentation, and support for Thin Clients.

The DB2 Administration Client also includes the client components for DB2 Query Patroller, a sophisticated query management and workload distribution tool. In order to use Query Patroller, you must have a Query Patroller server installed. For more information refer to the *DB2 Query Patroller Installation Guide*.

DB2 Administration Clients are available for the following platforms: AIX, HP-UX, Linux, OS/2, Solaris, Windows 32-bit operating systems.

---

## DB2 Application Development Client

The DB2 Application Development Client was known as the DB2 Software Development Kit (DB2 SDK) in previous versions of DB2. DB2 Application Development Client provides the tools and environment you need to develop applications that access DB2 servers and application servers that implement Distributed Relational Database Architecture (DRDA). You can build and run DB2 applications with a DB2 Application Development Client installed. You can also run DB2 applications on a DB2 Administration Client and a DB2 Run-Time Client.

DB2 Application Development Clients are available for the following platforms: AIX, HP-UX, Linux, OS/2, NUMA-Q, Silicon Graphics IRIX, the Solaris Operating Environment, and Windows 32-bit operating systems.

The applicable DB2 Application Development Client can be found on the server product CD-ROM. The DB2 Application Development Client for all platforms can be found on the set of DB2 Application Development Client CD-ROMs.

---

## Distributed Installation

If you are planning to install DB2 products across your network, you may consider using a distributed installation. With a network-based installation, you can roll out multiple identical copies of DB2 products. For more information on performing a distributed installation, refer to the *Installation and Configuration Supplement*.



---

## DB2 Thin Client

You can install a DB2 client for Windows 9x, Windows NT, or Windows 2000 on a code server, and have Thin Client workstations access the code across a LAN connection. Thin Client workstations function like any other DB2 client. In the setup, the main difference is that the DB2 client code is installed on a code server, and not individually on each workstation. Thin Client workstations need only a minimal configuration to set parameters and establish links to a code server. For more information on installing DB2 Thin Clients, refer to the *Installation and Configuration Supplement*.



---

## Chapter 7. Installing DB2 Clients on Windows 32-Bit Operating Systems

This section contains the information you need to install a DB2 client on Windows 32-bit operating systems.

---

### Before You Begin Installing

1. Ensure that your system meets all of the memory, hardware, and software requirements to install your DB2 client. For more information, see “Chapter 2. Planning for Installation” on page 15.
2. You require a user account to perform the installation.

#### **Windows 9x**

Any valid Windows 9x user.

#### **Windows NT or Windows 2000**

A user account that belongs to a group with more authority than the Guests group. For example, the Users group or the Power Users group. For information about installing on Windows NT or Windows 2000 without administrator authority, see “Installing Without Administrator Authority”.

### Installing Without Administrator Authority

When installing a DB2 client without administrator authority on Windows NT and Windows 2000, you will not be able to install these components:

- Control Center
- NetQuestion
- Integrated SNA Support

The following are some specific installation scenarios:

- *A user has installed a DB2 product without having administrator authority, and then an administrator installs a DB2 product on the same machine. In this scenario, the installation performed by the administrator will remove the previous installation performed by the user without administrator authority, resulting in a clean installation of the DB2 product. The installation performed by the administrator will overwrite all of the user’s services, shortcuts, and environment variables from the previous installation of DB2.*

- A user without administrator authority has installed a DB2 product, and then a second user without administrator authority attempts to install a DB2 product on the same machine. In this scenario, the installation by the second user will fail, and return an error message that the user must be an administrator to install the product.
- An administrator has installed a DB2 product, and then a user without administrator authority attempts to install a DB2 single-user product on the same machine. In this scenario, the installation being attempted by the user without administrator authority will fail, and return an error message saying that the user must be an administrator to install the product.

---

## Installation Steps

To install a DB2 client, perform the following steps:

- Step 1. Log on to the system with the user account that you want to use to perform the installation.
- Step 2. Shut down any other programs so that the setup program can update files as required.
- Step 3. Insert the appropriate CD-ROM into the drive. The auto-run feature automatically starts the setup program. The setup program will determine the system language, and launch the setup program for that language. If you want to run the setup program in a different language, or the setup program failed to start, see the following tip:



To manually invoke the setup program, perform the following steps:

- a. Click **Start** and select the **Run** option.
- b. In the **Open** field, enter the following command:


```
x:\setup /i language
```

where:

- *x*: represents your CD-ROM drive
- *language* represents the country code for your language (for example, EN for English). Table 8 on page 99 lists the code for each available language.

- c. Click **OK**.

Step 4. The DB2 Launchpad opens. It looks similar to the following:



Step 5. From this window, you can view the Installation Prerequisites and the Release Notes, you can take a Quick Tour to explore the features, capabilities, and benefits of DB2 Universal Database Version 7, or you can proceed directly to the installation.

Once you have initiated the installation, proceed by following the setup program's prompts. Online help is available to guide you through the remaining steps. Invoke the online help by clicking **Help**, or by pressing **F1** at any time. You can click **Cancel** at any time to end the installation.



For information on errors encountered during installation, see the db2.log file. The db2.log file stores general information and error messages resulting from the install and uninstall activities. By default, the db2.log file is located in the x:\db2log directory, where x: represents the drive on which your operating system is installed.

For more information, refer to the *Troubleshooting Guide*.

The setup program performs the following actions:

- Create DB2 program groups and items (or shortcuts).
- Update the Windows registry.
- Create a default client instance called DB2.



To configure your client to access remote servers, go to “Chapter 8. Configuring Client-to-Server Communications Using the Client Configuration Assistant” on page 61.

---

---

## Chapter 8. Configuring Client-to-Server Communications Using the Client Configuration Assistant

This chapter describes how to configure client-to-server communications using the Client Configuration Assistant (CCA). In an LDAP-enabled environment, you may not need to perform the tasks described in this chapter.

### Notes:

1. The CCA is available for DB2 clients running on OS/2 and Windows 32-bit systems.
2. LDAP support is available for Windows, AIX, and the Solaris operating environment.

---

### LDAP Directory Support Considerations

In an LDAP-enabled environment, the directory information about DB2 servers and databases is kept in the LDAP directory. When a new database is created, the database is automatically registered in the LDAP directory. During a database connection, the DB2 client goes to the LDAP directory to retrieve the required database and protocol information and uses this information to connect to the database. There is no need to run the CCA to configure LDAP protocol information.

You may still want to use the CCA in the LDAP environment to:

- Manually catalog a database in the LDAP directory
- Register a database as an ODBC data source
- Configure CLI/ODBC information
- Remove a database cataloged in the LDAP directory

For more information about the LDAP Directory Support, refer to the *Installation and Configuration Supplement*.

---

### Before You Begin

When you add a database using this configuration method, the CCA will generate a default node name for the server where the database resides.

To complete the steps in this section, you should be familiar with how to start the CCA. For more information, see “Starting the Client Configuration Assistant” on page 89.

**Note:** To configure communications from a DB2 client to a server, the remote server must be configured to accept inbound client requests. By default, the server installation program automatically detects and configures most protocols on the server for inbound client connections. It is recommended that you install and configure the communications protocols desired on the server before installing DB2.

If you have added a new protocol to your network that is not detectable, or want to modify any of the default settings, refer to the *Installation and Configuration Supplement*.

If you are adding a host or AS/400 database, refer to the "Configuring DB2 Connect to Host or AS/400 Communications Using the Client Configuration Assistant" section in your *DB2 Connect Quick Beginnings* manual.

---

## Configuration Steps

To configure your workstation to access a database on a remote server, perform the following steps:

**Step 1.** Log on to the system with a valid DB2 user ID. For more information, see "Appendix C. Naming Rules" on page 101.



If you are adding a database to a system that has a DB2 server or DB2 Connect server product installed, log on to this system as a user with System Administrative (SYSADM) or System Controller (SYSCTRL) authority on the instance. For more information, see "Working with the System Administrative Group" on page 93.

This restriction is controlled by the *catalog\_noauth* database manager configuration parameter. For more information, refer to the *Administration Guide*.

**Step 2.** Start the CCA. For more information, see "Starting the Client Configuration Assistant" on page 89.

The Welcome window opens each time you start the CCA, until you add at least one database to your client.

**Step 3.** Click on the **Add** push button to configure a connection.

You can use one of the following configuration methods:

- "Adding a Database Using a Profile" on page 63.
- "Adding a Database Using Discovery" on page 64.
- "Adding a Database Manually" on page 66.



## Adding a Database Using a Profile

A server profile contains information about server instances on a system, and databases within each server instance. For information on profiles, see “Creating and Using Profiles” on page 68.

If your administrator provided you with a profile, perform the following steps:

- Step 1. Select the **Use a profile** radio button and click the **Next** push button.
- Step 2. Click the ... push button and select a profile. Select a remote database from the object tree that is displayed from the profile, and if the database selected is a gateway connection, select a connection route to the database. Click the **Next** push button.
- Step 3. Enter a local database alias name in the **Database alias** field and optionally enter a comment that describes this database in the **Comment** field. Click **Next**.
- Step 4. If you are planning to use ODBC, register this database as an ODBC data source.

**Note:** ODBC must be installed to perform this operation.

- a. Ensure that the **Register this database for ODBC** check box is selected.
  - b. Select the radio button that describes how you would like to register this database:
    - If you would like all users on your system to have access to this data source, select the **As a system data source** radio button.
    - If you would like only the current user to have access to this data source, select the **As a user data source** radio button.
    - If you would like to create an ODBC data source file to share database access, select the **As a file data source** radio button and enter the path and file name for this file in the **File data source name** field.
  - c. Click the **Optimize for application** drop down box and select the application for which you want to tune the ODBC settings.
  - d. Click **Finish** to add the database that you selected. The Confirmation window opens.
- Step 5. Click the **Test Connection** push button to test the connection. The Connect to DB2 Database window opens.
  - Step 6. In the Connect To DB2 Database window, enter a valid user ID and password for the remote database and click **OK**. If the connection is successful, a message confirming the connection appears.

If the connection test failed, you will receive a help message. To change any settings you may have incorrectly specified, click the **Change** push button in the Confirmation window to return to the Add Database Wizard. If problems persist, refer to the *Troubleshooting Guide* for more information.

- Step 7. You are now able to use this database. Click **Add** to add more databases, or click **Close** to exit the Add Database Wizard. Click **Close** again to exit the CCA.

## Adding a Database Using Discovery



This option cannot return information about DB2 systems earlier than Version 5 or any systems where an Administration Server is not running. For more information, refer to the *Administration Guide*.

You can use the Discovery feature to search a network for databases. To add a database to your system using Discovery, perform the following steps:

- Step 1. Select the **Search the network** radio button and click the **Next** push button.
- Step 2. Click the [+] sign beside the **Known Systems** icon to list all the systems known to your client.
- Step 3. Click the [+] sign beside a system to get a list of the instances and databases on it. Select the database that you want to add, click the **Next** push button, and proceed to Step 4.

If the system that contains the database that you want to add is not listed, perform the following steps:

- a. Click on the [+] sign beside the **Other Systems (Search the network)** icon to search the network for additional systems.
- b. Click on the [+] sign beside a system to get a list of the instances and databases on it.
- c. Select the database that you want to add, click **Next**, and proceed to Step 4.



The Client Configuration Assistant may be unable to detect a remote system if:

- The Administration Server is not running on the remote system.
- The Discovery function times out. By default, the Discovery function will search the network for 40 seconds; this may not be long enough to detect the remote system. You can set the *DB2DISCOVERYTIME* registry variable to specify a longer period of time.
- The network that the Discovery request is running on is configured so that the Discovery request does not reach the remote system desired.
- You are using NetBIOS as the Discovery protocol. You may need to set the *DB2NBDISCOVERRCVBUFS* registry variable to a larger value to enable the client to receive more concurrent Discovery replies.

For more information, refer to the *Administration Guide*.

If the system that you want to add is still not listed, it can be added to the list of systems by performing the following steps:

- a. Click **Add System**. The Add System window opens.
- b. Enter the required communication protocol parameters for the remote Administration Server and click **OK**. A new system is added. For more information, click **Help**.
- c. Select the database that you want to add and click **Next**.

**Step 4.** Enter a local database alias name in the **Database alias** field and optionally enter a comment that describes this database in the **Comment** field. Click **Next**.

**Step 5.** If you are planning to use ODBC, register this database as an ODBC data source.

**Note:** ODBC must be installed to perform this operation.

- a. Ensure that the **Register this database for ODBC** check box is selected.
- b. Select the radio button that describes how you would like to register this database:
  - If you would like all users on your system to have access to this data source, select the **As a system data source** radio button.
  - If you would like only the current user to have access to this data source, select the **As a user data source** radio button.

- If you would like to create an ODBC data source file to share database access, select the **As a file data source** radio button and enter the path and file name for this file in the **File data source name** field.
  - c. Click the **Optimize for application** drop down box and select the application for which you want to tune the ODBC settings.
  - d. Click **Finish** to add the database that you selected. The Confirmation window opens.
- Step 6.** Click the **Test Connection** push button to test the connection. The Connect to DB2 Database window opens.
- Step 7.** In the Connect To DB2 Database window, enter a valid user ID and password for the remote database and click **OK**. If the connection is successful, a message confirming the connection appears.
- If the connection test failed, you will receive a help message. To change any settings you may have incorrectly specified, click the **Change** push button in the Confirmation window to return to the Add Database Wizard. If problems persist, refer to the *Troubleshooting Guide* for more information.
- Step 8.** You are now able to use this database. Click **Add** to add more databases, or click **Close** to exit the Add Database Wizard. Click **Close** again to exit the CCA.

## Adding a Database Manually

If you have the information for the database you want to connect to and the server upon which it resides, you can manually enter all of the configuration information. This method is analogous to entering commands via the command line processor, however, the parameters are presented graphically for you.

To add a database to your system manually, perform the following steps:

- Step 1.** Select the **Manually configure a connection to a database** radio button and click **Next**.
- Step 2.** If you are using Lightweight Directory Access Protocol (LDAP), select the radio button that corresponds to the location where you would like your DB2 directories to be maintained:
- If you would like to maintain the DB2 directories locally, select the **Add database to your local machine** radio button and click **Next**.
  - If you would like to maintain the DB2 directories globally at an LDAP server, select the **Add database using LDAP** radio button and click **Next**.
- Step 3.** Select the radio button that corresponds to the protocol that you want to use from the **Protocol** list.

If DB2 Connect (or the DB2 Connect Support Feature) is installed on your machine and you select TCP/IP or APPC, you can select **The database physically resides on a host or AS/400 system**. If you select this check box, you will have the option of selecting the type of connection that you want to make to the host or AS/400 database:

- To make a connection through a DB2 Connect gateway, select the **Connect to the server via the gateway** radio button.
- To make a direct connection, select the **Connect directly to the server** radio button.

Click **Next**.

- Step 4. Enter the required communication protocol parameters and click **Next**. For more information, click **Help**.
- Step 5. Enter the database alias name of the remote database that you want to add in the **Database name** field and a local database alias name in the **Database alias** field.

If this is a host or AS/400 database, type the Location name for an OS/390 database, the RDB name for an AS/400 database, or the DBNAME for a VSE or VM database in the **Database name** field, and optionally add a comment that describes this database in the **Comment** field.

Click **Next**.

- Step 6. Register this database as an ODBC data source.

**Note:** ODBC must be installed to perform this operation.

- a. Ensure that the **Register this database for ODBC** check box is selected.
- b. Select the radio button that describes how you would like to register this database:
  - If you would like all users on your system to have access to this data source, select the **As a system data source** radio button.
  - If you would like only the current user to have access to this data source, select the **As a user data source** radio button.
  - If you would like to create an ODBC data source file to share database access, select the **As a file data source** radio button and enter the path and file name for this file in the **File data source name** field.
- c. Click the **Optimize for application** drop down box and select the application for which you want to tune the ODBC settings.
- d. Click **Finish** to add the database that you selected. The Confirmation window opens.

- Step 7. Click the **Test Connection** push button to test the connection. The Connect to DB2 Database window opens.
- Step 8. In the Connect To DB2 Database window, enter a valid user ID and password for the remote database and click **OK**. If the connection is successful, a message confirming the connection appears.
- If the connection test failed, you will receive a help message. To change any settings you may have incorrectly specified, click the **Change** push button in the Confirmation window to return to the Add Database Wizard. If problems persist, refer to the *Troubleshooting Guide* for more information.
- Step 9. You are now able to use this database. Click **Add** to add more databases, or click **Close** to exit the Add Database Wizard. Click **Close** again to exit the CCA.

You can use the Export function of the CCA to create a client profile for an existing client configuration and use it to create identical target clients across your network. A client profile contains database connection, ODBC/CLI, and configuration information for an existing client. Use the CCA Import function to set up multiple clients across your network. Each target client will have the same configuration and settings as the existing client. For more information on creating and using client profiles, see “Creating and Using Profiles”.



You have now completed all the tasks that are involved in *Quick Beginnings* and are ready to start using DB2 Universal Database.

If you want to deploy this product using a distributed installation, refer to the *Installation and Configuration Supplement*.

---

## Creating and Using Profiles

The information in this section describes how to create and use profiles to set up connections between DB2 clients and servers. To configure database connections on a DB2 client, you can use either a server profile or client profile.

### Server Profiles

A server profile contains information about instances on a server system, and databases within each instance. The information for each instance includes the protocol information required to set up a client to connect to databases in that instance.



We recommend that you create a server profile only after you have created the DB2 databases that you want your remote clients to access.

To create a server profile, perform the following steps:

- Step 1. Start the Control Center. For more information, refer to “Starting the DB2 Control Center” on page 90.
- Step 2. Select the system that you want to create a profile for and right click. If the system that you want to create a profile for is not shown, select the **Systems** icon, click the right mouse button, and select the **Add** option. Click the **Help** push button and follow the online help.
- Step 3. Select the **Export Server Profile** option.
- Step 4. Enter a path and filename for this profile and select **OK**.



You are ready to use this profile on your system. For more information on how to add a database to your system using a server profile, go to “Configuration Steps” on page 62.

---

## Client Profiles

Information in a client profile can be used to configure clients using the Import function in the Client Configuration Assistant (CCA). Clients can import all or a subset of the configuration information in a profile. The following scenario assumes that the database connections configured on one client will be exported and used to set up one or more clients.

**Note:** Configuration profiles can also be imported using the **db2cfimp** command. Refer to the *Command Reference* for more information.

A client profile is generated from a client using the Export function of the CCA. The information contained in a client profile is determined during the export process. Depending on the settings chosen, it can contain the existing client’s:

- Database connection information (including CLI or ODBC settings).
- Client settings (including database manager configuration parameters and DB2 registry variables).
- CLI or ODBC common parameters.
- Configuration data for the local APPC or NetBIOS communications subsystem.

To create a client profile, perform the following steps:

- Step 1. Start the CCA. For more information, see “Starting the Client Configuration Assistant” on page 89.
- Step 2. Click **Export**. The Select Export Option window opens.
- Step 3. Select one of the following export options:

- If you want to create a profile that contains all of the databases cataloged on your system, and all of the configuration information for this client, select the **All** radio button, click **OK**, and go to Step 8.
- If you want to create a profile that contains all of the databases cataloged on your system *without* any of the configuration information for this client, select the **Database connection information** radio button, click **OK**, and go to Step 8.
- If you want to select a subset of the databases that are cataloged on your system, or a subset of the configuration information for this client, select the **Customize** radio button, click **OK**, and go to the next step.

**Step 4.** Select the databases to be exported from the **Available databases** box and add them to the **Selected databases** box by clicking on the push button.



To add all of the available databases to the **Databases to be exported** box, click the >> button.

**Step 5.** Select the check boxes from the **Select custom export option** box that correspond to the options that you want to set up for the target client.

To customize settings, click on the appropriate **Customize** push button. The settings that you customize will only affect the profile to be exported, no changes will be made to your workstation. For more information, click **Help**.

**Step 6.** Click **OK**. The Export Client Profile window opens.

**Step 7.** Enter a path and file name for this client profile and click **OK**. The DB2 Message window appears.

**Step 8.** Click **OK**.

To import a client profile, perform the following steps:

**Step 1.** Start the CCA. For more information, see “Starting the Client Configuration Assistant” on page 89.

**Step 2.** Click **Import**. The Select Profile window opens.

**Step 3.** Select a client profile to import and click **OK**. The Import Profile window opens.

**Step 4.** You can select to import all or a subset of the information in a Client Profile. Select one of the following import options:

- To import everything in a client profile, select the **All** radio button.



- To import a specific database, or settings, that are defined in a Client Profile, select the **Customize** radio button. Select the check boxes that correspond to the options that you want to customize.

Step 5. Click **OK**.



If you selected the **All** radio button, you are now ready to start using your DB2 product. For more advanced topics, refer to the *Administration Guide* and the *Installation and Configuration Supplement*.

---

- Step 6. You are presented with a list of systems, instances, and databases. Select the database that you want to add and click **Next**.
- Step 7. Enter a local database alias name in the **Database alias** field and optionally enter a comment that describes this database in the **Comment** field. Click **Next**.
- Step 8. If you are planning to use ODBC, register this database as an ODBC data source.

**Note:** ODBC must be installed to perform this operation.

- a. Ensure that the **Register this database for ODBC** check box is selected.
  - b. Select the radio button that describes how you would like to register this database:
    - If you would like all users on your system to have access to this data source, select the **As a system data source** radio button.
    - If you would like only the current user to have access to this data source, select the **As a user data source** radio button.
    - If you would like to create an ODBC data source file to share database access, select the **As a file data source** radio button and enter the path and file name for this file in the **File data source name** field.
  - c. Click the **Optimize for application** drop down box and select the application for which you want to tune the ODBC settings.
  - d. Click **Finish** to add the database that you selected. The Confirmation window opens.
- Step 9. Click the **Test Connection** push button to test the connection. The Connect to DB2 Database window opens.
- Step 10. In the Connect To DB2 Database window, enter a valid user ID and password for the remote database and click **OK**. If the connection is successful, a message confirming the connection appears.

If the connection test failed, you will receive a help message. To change any settings you may have incorrectly specified, click the **Change** push button in the Confirmation window to return to the

Add Database Wizard. If problems persist, refer to the *Troubleshooting Guide* for more information.

- Step 11.** You are now able to use this database. Click **Add** to add more databases, or click **Close** to exit the Add Database Wizard. Click **Close** again to exit the CCA.

---

## Chapter 9. Control Center Installation and Configuration

This chapter describes how to install and configure the DB2 Control Center.

The Control Center is the main DB2 graphical tool for administering your database. It is available on Windows 32-bit, OS/2, and UNIX operating systems.

The Control Center provides a clear overview of all the systems and database objects being managed. You can also access other administration tools from the Control Center by selecting icons on the Control Center toolbar or from the Tools pop-up menu.

---

### Application versus Applet

You can run the Control Center either as a Java application or as a Java applet through a web server. In both cases you need a supported Java Virtual Machine (JVM) installed on your machine to run the Control Center. A JVM can be a Java Runtime Environment (JRE) for running applications, or a Java-enabled browser for running applets.

- Java *applications* run just like other applications on your machine, provided you have the correct JRE installed.

On Windows 32-bit operating systems, the correct JRE level was installed or upgraded for you during DB2 installation.

On AIX systems, the correct JRE was installed for you during DB2 installation only if another JRE was not detected on your system. If another JRE was detected on your AIX system during DB2 installation, the JRE that comes with DB2 was not installed. In this case, you must install the correct JRE level before running the Control Center.

On all other operating systems you must install the correct JRE level before running the Control Center. See Table 7 on page 75 for a list of correct JRE levels.

**Note:** Some operating systems, including OS/2 Warp Server for e-business and AIX 4.3, have built-in Java support. For more information, check with your administrator.

- Java *applets* are programs that run within Java-enabled browsers. The Control Center applet code can reside on a remote machine and is served to the client's browser through a web server. This type of client is often called a *thin client* because a minimal amount of resources (a Java-enabled browser) is required to run the Java applet.

You must use a supported Java-enabled browser to run the Control Center as a Java applet. See Table 7 on page 75 for a list of supported browsers.

## Machine Configurations

You can set up your Control Center in a number of different ways. The following table identifies four scenarios, each showing a different way of installing the required components. These scenarios are referenced throughout the Control Center Services Setup (Applet Mode only) section that follows the table.

*Table 6. Control Center Machine Configuration Scenarios*

Scenario	Machine A	Machine B	Machine C
1 - Stand-alone, Application	JRE Control Center application DB2 server		
2 - Two Tier, Application	JRE Control Center application DB2 client		DB2 server
3 - Two Tier, Browser	Supported Browser (Windows and OS/2 only) Control Center applet	Web server JDBC Applet Server DB2 server	
4 - Three Tier, Browser	Supported Browser (Windows and OS/2 only) Control Center applet	JDBC Applet Server DB2 client	DB2 server

Figure 4 on page 75 summarizes the four basic Control Center machine configurations:




Figure 4. DB2 Control Center machine configurations

## Supported Java Virtual Machines for the Control Center

The following table lists supported Java Virtual Machines (JREs and browsers) required to run the Control Center as an application or applet:

Table 7. Supported Java Virtual Machines (JVMs) for the Control Center

Operating System	Correct Java Runtime Environments	Supported Browsers
Windows 32-bit	JRE 1.1.8 (automatically installed or updated by DB2, if necessary)	Netscape 4.5 or higher (shipped), or IE 4.0 Service Pack 1 or higher
AIX	JRE 1.1.8.4 (automatically installed if no other JREs detected)	None
OS/2	JRE 1.1.8	Netscape 4.6 (shipped)
Linux	JRE 1.1.8	None

Table 7. Supported Java Virtual Machines (JVMs) for the Control Center (continued)

Operating System	Correct Java Runtime Environments	Supported Browsers
Solaris	JRE 1.1.8	None
HP-UX 11	JRE 1.1.8	None
IRIX	JRE 1.1.8 (3.1.1 SGI) + Cosmo code 2.3.1	None
PTX	JRE 1.1.8	None

For the latest information on supported JREs and browsers, go to <http://www.ibm.com/software/data/db2/udb/db2cc>

---

## Setting Up and Working with the Control Center

This section describes how to set up and customize the Control Center for your environment.

### Control Center Services Setup (Applet Mode only)

If you will run the Control Center as an application, skip this section and go to “Running the Control Center as a Java Application” on page 78.

To set up the Control Center to run as an applet:

1. Start the Control Center JDBC Applet Server.
2. On Windows NT or Windows 2000, start the security server.

#### 1. Start the Control Center JDBC Applet Server

To start the Control Center JDBC Applet Server, enter the `db2jstrt 6790` command, where 6790 represents any 4-digit port number that is not already in use.

We recommend that you start the Control Center JDBC Applet Server with a user account that has SYSADM authority.

The first time you start the Control Center JDBC Applet Server, it will create several node directory entries, together with various files for administration purposes. In Scenarios 1 and 3 in “Machine Configurations” on page 74, all of these administration files and directory entries will be created in the current DB2 instance.

Most DB2 resources are accessed over **database connect** or **instance attach**. In both cases the user must supply valid user ID and password combinations to gain access. However, some resources are accessed directly by the Control Center JDBC Applet Server, including the database and node directories (catalogs), and the Command Line Processor. Access to these resources is

performed by the Control Center JDBC Applet Server on behalf of the logged in Control Center user. Both the user and the server are required to have the proper authorization before access will be granted. For example, to update the database directory, a minimum of SYSCTRL authority is required.

It is possible to run an instance of the Control Center JDBC Applet Server with any level of security, but you will be unable to update certain resources such as database and node directories. In particular, you may encounter a **SQL1092N** message informing you of the lack of authorization for a request. The user specified in the message can be either the user signed into Control Center, or the user account running the Control Center JDBC Applet Server.

On Windows NT, you can start the Control Center JDBC Applet Server by clicking **Start** and selecting **Settings** → **Control Panel** → **Services**. Select the **DB2 JDBC Applet Server - Control Center** service and click **Start**.

On Windows 2000, you can start the Control Center JDBC Applet Server by clicking **Start** and selecting **Settings** → **Control Panel** → **Administrative Tools** → **Services**. Select the **DB2 JDBC Applet Server - Control Center** service, click the **Action** menu and then select **Start**.

On any system, you can start the Control Center JDBC Applet Server by entering:

```
net start DB2ControlCenterServer
```

This step is not required if your Control Center JDBC Applet Server is autostarted.

If you start the Control Center JDBC Applet Server as a Windows NT or Windows 2000 service, you must configure the startup in the services dialog to change the account information.

## **2. Start the Windows NT or Windows 2000 Security Server**

To work with the Control Center on Windows NT or Windows 2000, the security server must be running. During DB2 installs, the Security Server is usually set up to autostart.

You can check if the Security Server is running on Windows NT by clicking **Start** and selecting **Settings** → **Control Panel** → **Services**.

On Windows 2000, click **Start** and select **Settings** → **Control Panel** → **Administrative Tools** → **Services**.

If the **DB2 Security Server** is not started on Windows NT, select it and click **Start**. On Windows 2000, select the **Action** menu and click **Start**.

After you have started the Control Center JDBC Applet Server and started the Windows NT or Windows 2000 security server (if necessary), go to “Running the Control Center as a Java Applet”.

## Working with the Control Center

You can run the Control Center as a Java application or as a Java applet. If your environment is configured similar to Scenarios 1 or 2 in Table 6 on page 74, you must run the Control Center as an application. If your environment is configured like Scenarios 3 or 4, you must run it as an applet.

### Running the Control Center as a Java Application

To run the Control Center as a Java application, you must have the correct Java Runtime Environment (JRE) installed. See Table 7 on page 75 for the correct JRE level for your operating system.

1. To start the Control Center as an application:

#### On Windows 32-bit operating systems:

Click **Start** and select **Programs** → **IBM DB2** → **Control Center**.

#### On OS/2:

Open the **IBM DB2** folder and double-click on the **Control Center** icon.

#### On all supported platforms:

Start the Control Center from a command prompt by entering the **db2cc** command.

2. The DB2 Control Center window opens.
3. You can start working with the Control Center without an existing database by creating a sample database. Enter the **db2sampl** command on the DB2 Universal Database server. On UNIX operating systems, ensure that you are logged in to the DB2 instance before you enter the **db2sampl** command.

### Running the Control Center as a Java Applet

To run the Control Center as a Java applet, you must have a Web server set up on the machine that contains the Control Center applet code and the Control Center JDBC Applet Server. The Web server must allow access to the `sqllib` directory.

If you choose to use a virtual directory, substitute this directory for the home directory. For example, if you map `sqllib` to a virtual directory called `temp` on a server named `yourserver`, a client would use the URL:

`http://yourserver/temp`

If you do not have the DB2 documentation installed and you would like to configure your web server to work with DB2's online documentation, refer to the *Installation and Configuration Supplement*.



To run Control Center as an applet on Windows 32-bit or OS/2 operating systems, you must run **db2classes.exe** on the machine where the DB2 JDBC Applet Server resides to unzip the required Java class files.

To load the Control Center HTML page, perform the following steps:

1. Start the **Control Center Launch** page through your web server. In your browser, select **File -> Open Page**. The **Open Page** dialog box appears. Enter the URL of your Web server and the main Control Center page and click on the **Open** push button. For example, if your server is named **yourserver**, you would open `http://yourserver/cc/prime/db2cc.htm`
2. In the **Server port** field, enter a value for the Control Center JDBC Applet Server port. The default server port value is 6790.
3. Click on the **Start Control Center** push button.
4. The Control Center Sign On window opens. Enter your user ID and password. This user ID must have an account on the machine that is running the Control Center JDBC Applet Server. Your initial logon will be used for all database connections. It can be changed from the Control Center pull-down menu. A unique user profile will be assigned to each user ID. Click **OK**.
5. The DB2 Control Center window opens.
6. You can start working with the Control Center without an existing database by creating a sample database. Enter the **db2sampl** command on the DB2 Universal Database server. On UNIX operating systems, ensure that you are logged in to the DB2 instance before you enter the **db2sampl** command.

### Customizing Your Control Center HTML File

To automatically start the Control Center the next time you open `db2cc.htm`, perform the following steps:

- For Scenarios 1 or 2, modify the `autoStartCC` parameter tag in `db2cc.htm` from

```
param name="autoStartCC" value="false"
```

to

```
param name="autoStartCC" value="true"
```

- For Scenarios 3 or 4, modify the `autoStartCC`, `hostNameText`, and `portNumberText` parameter tags in `db2cc.htm` to

```
param name="autoStartCC" value="true"  
param name="hostNameText" value="yourserver"  
param name="portNumberText" value="6790"
```

where `yourserver` represents the Server name or IP address and `6790` represents the server port value of the machine to which you want to connect.

## **Configuring Your Web Server to Work with the Control Center**

For general web server configuration information, consult the setup documentation that came with your web server.

For more information on serving DB2 online documentation through a web server, refer to the *Installation and Configuration Supplement*.

---

## **Functional Considerations**

If you are using the Control Center over the Internet, be aware that there is no encryption of the data flow between the Control Center JDBC Applet Server and the browser.

To use the color options of Visual Explain on Netscape, you must set your operating system to support more than 256 colors.

On OS/2 systems, you must install the Control Center on an HPFS-formatted drive. DB2 does not support the installation of the Control Center on an OS/2 FAT drive, because an OS/2 FAT drive does not support long filenames required by Java.

Every activity will be associated with an explicit DB2 connection or attachment. For security purposes, every DB2 activity will be validated.

When you are using the Control Center under Scenarios 3 or 4, the local system is Machine B. The local system is the system name as it appears in the DB2 Control Center window.

---

## **Installation Tips for Control Center Help on UNIX Operating Systems**

When installing the Control Center online help on UNIX operating systems you should keep the following in mind:

- You should install the Control Center help and the product documentation at the same time. If you install the Control Center help and the DB2 online product documentation separately, you can expect the second installation to take some time. This is true regardless of which package is installed first.
- You must select the Control Center help for any non-English language explicitly. Installing the product messages for a particular language does not mean that the Control Center help for that language is automatically installed. However, if you install the Control Center help for a particular language, the product messages for that language are installed automatically.

- If you manually install the Control Center on UNIX-based workstations rather than using the `db2setup` utility, you must run the `db2insthtml` command to install the online documentation. For more information, refer to the *DB2 for UNIX Quick Beginnings*.

---

## Configuring TCP/IP on OS/2

To run the Control Center on OS/2 Warp 4 while disconnected from a LAN, you will need to configure TCP/IP to enable local loopback and localhost. If you are running OS/2 Warp Server for e-business, local loopback is enabled by default.

### Enabling Local Loopback

To enable local loopback on your system:

1. Open the **System Setup** folder.
2. Open the **TCP/IP Configuration** notebook.
3. View the **Network** page.
4. In the **Interface to Configure** list box, highlight **loopback interface**.
5. If the **Enable interface** check box is not selected, select it now.
6. Verify that the **IP address** is 127.0.0.1 and **Subnet Mask** is empty.

### Enabling Localhost

To enable localhost on your system:

1. To check if localhost is enabled, enter the **ping localhost** command.
  - If data is returned and localhost is enabled, you can skip steps 2 and 3 below and go directly to step 4.
  - If localhost unknown is returned, or if the command hangs, localhost is not enabled. Go to step 2.
2. If you are on a network, make sure that loopback is enabled. To enable local loopback see “Enabling Local Loopback”.
3. If you are *not* on a network, enable localhost by performing these steps:
  - a. Add the following line after other `ifconfig` lines in the `MPTN\BIN\SETUP.CMD` command file:

```
ifconfig lo 127.0.0.1
```
  - b. In the TCP/IP configuration folder, perform the following steps:
    - 1) Go to the **Configure Name Resolution Services** page.
    - 2) In the **Hostname configuration without a Nameserver** table, add an entry with *IP Address* set to 127.0.0.1 and *Hostname* set to localhost.

**Note:** If you have a hostname for your machine on the **Configure LAN Name Resolution Services** page, you must add this name as an alias when you set the *IP Address* 127.0.0.1 to localhost.

- c. Select the **Look through HOSTS list before going to the nameserver** list box. This step tells your OS/2 system that when it is looking for a host, such as localhost, it should use the host address found on your machine rather than checking the nameserver. If the host is not defined on your machine, OS/2 continues looking for the host by using the nameserver you configured.
  - d. Close **TCP/IP Configuration** and reboot the system.
  - e. You should be able to ping localhost without being connected to any network.
4. Verify that your hostname is correct. On an OS/2 command line, enter the **hostname** command. The hostname returned should match the one listed in the **TCP/IP Configuration** notebook on the **Hostnames** page and it must be less than 32 characters. If the hostname deviates from these conditions, correct it on the **Hostnames** page.
  5. Verify that your hostname is set properly in CONFIG.SYS. You should see a line similar to the following:

```
SET HOSTNAME=<correct_name>
```

where *<correct\_name>* represents the value returned by the **hostname** command. If this is not the case, make the necessary changes and reboot your system when you are finished.

## Verifying TCP/IP Configuration on OS/2

If you are having trouble running the Control Center on OS/2 while disconnected from a LAN, try running the **sniffle /P** command to diagnose the problem.

---

## Considerations for Remotely Administering Clusters Using the Control Center

After you install DB2 Universal Database Enterprise – Extended Edition for Windows on multiple nodes, the setup process will only configure the instance owning machine to be the coordinator node for communications from client machines. To remotely administer nodes using the Control Center, you will need to enable communications on each of the nodes.

You need to update the services file and specify the port that you want the server to listen on for incoming client requests from the Control Center. The services file is located in the `\winnt\system32\drivers\etc` directory.

**Note:** The location of the services file can depend on the products you have installed on your system. See your TCP/IP documentation for more information.

To enable communications on one of the nodes, reserve a TCP/IP port on each of the machines. Using a local test editor, add the port entry to the services file for TCP/IP support. For example:

```
db2cntl 4000/tcp # Reserved for Control Center - Coordinator function
```

where:

**db2cntl**

is the Service name that has been set by the instance owner,

**4000** is the port number, and

**tcp** is the communication protocol that you are using.

---

## Troubleshooting Information

For the latest service information on the Control Center, point your browser to <http://www.ibm.com/software/data/db2/udb/db2cc>

If you are having problems running the Control Center, check the following:

- Ensure the Control Center JDBC Applet Server (db2jd) is running.
- Verify that the server port number is correct.
- Check that the Control Center JDBC Applet Server is running under a user account that has SYSADM authority.
- Ensure that the Database Administration Server (DAS) is running on any DB2 Universal Database systems you are trying to administer by entering the **db2admin start** command. On UNIX-based systems, ensure that you are logged in as the DAS instance owner when you issue this command.

If you are having problems running the Control Center as an *application*, also check the following:

- Verify that the correct JRE is installed. See Table 7 on page 75 for more information.

If you are having problems running the Control Center as an *applet*, also check the following:

- Verify that you are running a supported browser. See Table 7 on page 75 for more information.
- Check your browser's Java console window for diagnostic and trace information for the Control Center.
- Make sure that the client browser does not have CLASSPATH set. To make sure CLASSPATH is not set, open a command window and enter **SET CLASSPATH=** then start your client browser from this command window. Also, note that if CLASSPATH is not set in a Windows NT or Windows 2000 environment, it may still get picked up from autoexec.bat from a Windows 9x installation on the same machine.
- Ensure that you are using the db2cc.htm file from the machine running the Control Center JDBC Applet Server.

- Remember that the Control Center works within the DB2 client's locale, and that the DB2 client is located at the Control Center JDBC Applet Server's location.

---

## **Administering DB2 for OS/390 and DB2 Connect Enterprise Edition servers with the Control Center**

The Control Center has been greatly enhanced to deliver new management functions to database administrators who need to manage DB2 for OS/390 V5.1 and later database servers.

The Control Center has also been enhanced to manage operational and performance characteristics of DB2 Connect Enterprise Edition connectivity servers. The combination of DB2 for OS/390 server management and new DB2 Connect monitoring support provide complete end-to-end administration and monitoring for desktop and web applications that work with DB2 for OS/390 servers.

The DB2 Control Center uses the familiar "explorer" interface to allow database administrators to easily navigate between different database servers and the database objects they manage. Context sensitive right-mouse activated menus provide administrators with the ability to change attributes of database objects and to launch commands and utilities.

Database objects are presented in a consistent fashion for all DB2 family servers. This greatly reduces the amount of learning that is required for administrators who need to manage both DB2 for OS/390 and DB2 Universal Databases on Windows NT, Windows 2000, UNIX and OS/2 servers. While the Control Center preserves consistency across servers it does not hide capabilities that are unique to each DB2 server. This gives database administrators the power to perform all aspects of their tasks.

The ability to manage DB2 Connect connectivity servers is delivered through management of user connections and by keeping vital statistics on various performance aspects of the connectivity server. For example, database administrators can easily view all of the users connected through a particular DB2 Connect server, and their connection characteristics.

Administrators can also collect load and performance information such as the number of SQL statements and transactions executed, number of bytes sent and received, statement and transaction execution times and much more. Collected data can be displayed using easy to understand live graphs.

### **Preparing DB2 for OS/390 Servers for the Control Center**

The DB2 Control Center uses stored procedures to deliver many of its management functions. Therefore, for the Control Center to function properly

each DB2 for OS/390 server that will be managed from the Control Center needs to have stored procedures enabled and the proper stored procedures installed on that server.

For more information on applying service and required Function Modification Identifiers, refer to the *DB2 for OS/390 Program Directory*.

## **Working with the Control Center**

Before you can work with a server and its databases you will need to catalog information about the server on the Control Center workstation. The DB2 Control Center only works with servers and databases that are cataloged on the workstation where the Control Center is running. The easiest way to accomplish this on Windows and OS/2 workstations is to use the DB2 Client Configuration Assistant (CCA).

Once the Control Center is running, start by clicking on the plus sign next to the server you want to administer. Select the database or connectivity server objects that you want to administer and right click on the object to work with object properties or to execute actions on the object. You can invoke the online help by clicking **Help**, or by pressing **F1** at any time.

## **Other Sources of Information**

For more information about using the Control Center to administer DB2 for OS/390, refer to the following online resource:

<http://www.ibm.com/software/data/db2/os390/v6facts/db2cc.html>

For complete information about DB2 for OS/390 Version 6, refer to the online library: <http://www.ibm.com/software/data/db2/os390/v6books.html>

For more information about stored procedures and the Control Center for OS/390, go to: <http://www.ibm.com/software/data/db2/os390/cc390/>





---

## Part 4. Appendixes



---

## Appendix A. Basic Task Knowledge

This section describes the basic tasks that you will need to know to use this product effectively.



Go to the task that you want to perform:

- “Starting First Steps”.
  - “Starting the Client Configuration Assistant”.
  - “Starting the DB2 Control Center” on page 90.
  - “Entering Commands Using the Command Center” on page 90.
  - “Entering Commands Using the Command Line Processor” on page 91.
  - “Working with the System Administrative Group” on page 93.
  - “Granting Advanced User Rights on Windows” on page 94.
  - “Working with the Business Intelligence Functions” on page 95.
  - “Setting the Number of Licensed Processors” on page 95.
  - “Upgrading DB2 from Try and Buy Mode” on page 96.
  - “Uninstalling DB2” on page 97.
- 

---

### Starting First Steps

Start First Steps as follows:

#### Windows 32-bit operating systems

Click on **Start** and select **Programs**—>**IBM DB2**—>**First Steps**

You can also start First Steps by entering the **db2fs** command at a command prompt.

---

### Starting the Client Configuration Assistant

Start the Client Configuration Assistant (CCA) as follows:

#### Windows 32-bit operating systems

Click on **Start** and select **Programs**—>**IBM DB2**—>**Client Configuration Assistant**

You can also start the CCA by entering the **db2cca** command at a command prompt.

---

## Starting the DB2 Control Center

You can run the DB2 Control Center as a Java *application* or as a Java *applet*.

### To run the Control Center as an application

Enter the **db2cc** command. Your system must have the correct Java Runtime Environment installed in order to run the Control Center as an application.

On Windows 32-bit and OS/2 systems, you can also start the Control Center as an application by invoking the **Control Center** icon in your **IBM DB2** program group.

### To run the Control Center as an applet

You must have a Java-enabled browser and you must perform some additional configuration steps to run the Control Center as an applet. For detailed instructions on running the Control Center as an applet or application, see “Chapter 9. Control Center Installation and Configuration” on page 73.

---

## Entering Commands Using the Command Center

This section describes how to enter commands using the Command Center. There are two versions of the Command Center. This section documents the Command Center that is accessible from the DB2 Control Center.

**Note:** If you do not have the Control Center installed, a Command Center with limited functionality is available through the IBM DB2 program group, or by entering the **db2cctr** command.

From the Command Center you can:

- Run SQL statements, DB2 commands, and operating system commands.
- See the execution result of SQL statements and DB2 commands in a results window. You can scroll through the results and save the output to a file.
- Save a sequence of SQL statements and DB2 commands to a script file. You can then schedule the script to run as a job. When a saved script is modified, all jobs dependent on the saved script inherit the new modified behavior.
- Recall and run a script file.
- See the execution plan and statistics associated with a SQL statement before execution.
- Get quick access to database administrative tools from the main tool bar.
- Display all the command scripts known to the system through the Script Center, with summary information listed for each.
- Use the SQLAssist tool to build complex queries.
- Display results in an table you can edit.

To start the Command Center, click on the **Command Center** icon in the Control Center.

The Command Center contains a large input area for entering commands. To run the commands you have entered, click on the **Execute** icon (the gears icon).



In the Command Center, you do not have to enter a command with the db2 prefix; instead you just enter the DB2 command. For example:

```
list database directory
```

To enter operating system commands, precede the operating-system command with an exclamation mark (!). For example:

```
!dir
```

If you want to enter multiple commands, you must end each command with the termination character, then press **Enter** to start the next command on a new line. The default termination character is a semicolon (;).

For example, you could connect to a database called SAMPLE and list all the system tables by entering the following command:

```
connect to sample;  
list tables for system
```

After you have clicked on the **Execute** icon, the results are displayed.

To recall commands that you have entered during your session, select the **Command history** drop down box, and select a command.

To save commands, select **Interactive** —> **Save Command As** from the menu bar. For more information, click on the **Help** push button or press the **F1** key.



You can use the **Append to Script** button and the Script page of the Command Center to store commonly used SQL statements or DB2 commands as scripts. For more information, click on the **Help** push button or press the **F1** key.

---

## Entering Commands Using the Command Line Processor

You can use the command line processor to enter DB2 commands, SQL statements, and operating system commands. It operates in the following modes:

## DB2 Command Window

The DB2 command line processor behaves like a command window from your operating system. You can enter operating system commands, DB2 commands, or SQL statements and view their output.

## Interactive Input Mode

The `db2` prefix that you use for DB2 commands (in the DB2 Command Window) is pre-entered for you. You can enter operating systems commands, DB2 commands, or SQL statements and view their output.

## File Input Mode

Processes commands that are stored in a file. For information on the file input mode, refer to the *Command Reference*.

## DB2 Command Window

To invoke a DB2 Command Window, do the following:

### Windows 32-bit operating systems

Click on **Start** and select **Programs** → **IBM DB2** → **Command Window**

You can also invoke the DB2 command window by entering the **db2cmd** command at your operating system's prompt.

If you are entering commands via the Command Window, you must include the `db2` prefix. For example:

```
db2 list database directory
```



If the DB2 command contains characters that have special meaning on the operating system you are using, you will need to enter the command in quotation marks to ensure that it is run properly.

For example, the following command would retrieve all the information from the *employee* table, even if the `*` character has a special meaning on the operating system:

```
db2 "select * from employee"
```

To enter a long command that does not fit on a single line, you must use a space followed by the line continuation character `"\"` at the end of one line, then press the **Enter** key to continue the command on to the next. For example:

```
db2 select empno, function, firstname, lastname, birthdate, from \  
db2 (cont.) => employee where function='service' and \  
db2 (cont.) => firstname='Lily' order by empno desc
```

## Interactive Input Mode

To invoke the command line processor in interactive input mode, do the following:

## Windows 32-bit operating systems

Click on **Start** and select **Programs** → **IBM DB2** → **Command Line Processor**.

You can also invoke the command line processor in interactive input mode by entering the **db2cmd** command followed by the **db2** command at your operating system's prompt.

In interactive input mode, the prompt looks like this:

```
db2 =>
```

In interactive input mode, you do not have to enter DB2 commands with a db2 prefix; instead, you just enter the DB2 command. For example:

```
db2 => list database directory
```

To enter operating-system commands in interactive mode, precede the operating system command with an exclamation mark (!). For example:

```
db2 => !dir
```

To enter a long command that does not fit on a single line, you must use a space followed by the line continuation character "\" at the end of one line, then press the **Enter** key to continue the command on to the next. For example:

```
db2 select empno, function, firstname, lastname, birthdate, from \
db2 (cont.) => employee where function='service' and \
db2 (cont.) => firstname='Lily' order by empno desc
```

To end interactive input mode, enter the **quit** command.

For more information on advanced topics using the CLP, refer to the *Command Reference*.

---

## Working with the System Administrative Group

By default, System Administrative (SYSADM) authority is granted to the following:

### Windows NT and Windows 2000

Any valid DB2 user account which belongs to the local Administrators group on the machine where the account is defined.

For example, if a user logs on to a domain account and tries to access a DB2 database, DB2 will go to a Domain Controller to enumerate groups (including the Administrator's group). You can change this behavior in either of two ways:

1. Set the registry variable `DB2_GRP_LOOKUP=local` and add the domain accounts (or global groups) to the local Administrators group.
2. Update the database manager configuration parameter `SYSADM_GROUP` to specify a new group. If you want that group enumerated on the local machine, you must also set the `DB2_GRP_LOOKUP` registry variable.

For a domain user to have SYSADM authority, it must belong to the Administrators group on the Domain Controller. Since DB2 always performs authorization at the machine where the account is defined, adding a domain user to the local Administrators group on the server does not grant the domain user SYSADM authority to this group.

To avoid adding a domain user to the Administrators group at the Domain Controller, we suggest that you create a global group and add the domain users that you want to grant SYSADM authority, and then update the DB2 configuration parameter `SYSADM_GROUP` with the name of the global group. To do so, enter the following commands:

```
db2stop
db2 update dbm cfg using sysadm_group global_group
db2start
```

For information on how to change the default SYSADM settings and how to assign this authority to a different user or set of users, refer to the *Administration Guide*.

---

## Granting Advanced User Rights on Windows

### Windows NT

To grant advanced user rights on Windows NT you must be logged on as a Local Administrator. Perform the following steps to grant user rights:

1. Click **Start** and select **Programs** → **Administrative Tools (Common)** → **User Manager for Domains**.
2. In the User Manager window, select **Policies** → **User Rights** from the menu bar.
3. In the User Rights Policy window, select the **Show Advanced User Rights** check box then in the **Right** drop down box, select the user right you want to grant. Click **Add**.
4. In the Add Users and Groups window select the user or a group you want to grant the right to and click **OK**.
5. In the User Rights Policy window, select the user or group you have added from the **Grant To** list box and click **OK**.



## Windows 2000

To grant advanced user rights on Windows 2000 you must be logged on as a Local Administrator. Perform the following steps to grant user rights:

1. Click **Start** and select **Settings** → **Control Panel** → **Administrative Tools**.
2. Select **Local Security Policy**.
3. In the left window pane, expand the **Local Policies** object, then select **User Rights Assignment**.
4. In the right window pane, select the user right that you want to assign.
5. From the menu, select **Action** → **Security...**
6. Click **Add**, then select a user or group to assign the right to, and click **Add**.
7. Click **OK**.

---

## Working with the Business Intelligence Functions

The Business Intelligence Tutorial walks you through several basic and advanced tasks using the Data Warehouse Center and the OLAP Starter Kit. You can start the Tutorial from the **Help** menu in the Data Warehouse Center or from the **Help** menu in the OLAP Starter Kit desktop. You can also start the Tutorial from the Getting Started item in the Information Center.

---

## Setting the Number of Licensed Processors

If you are using an SMP machine and you have purchased additional processor entitlements, you can update this information using the **db2licm** command.

To update the number of License Processors, perform the following steps:

1. Log in as a user with SYSADM, SYSCTRL or SYSMAINT authority.
2. The **db2licm** utility can be found in the following locations:
  - For Unix-based operating systems, if `INSTHOME/sql11ib/adm` is not in your `PATH`, change your directory.
  - For Windows 32-bit operating systems and OS/2, go to `x:\DB2DIR\bin`, where `x:\DB2DIR\` is your DB2 installation drive and path.
3. Obtain the product password by issuing the **db2licm -l** command. The DB2 products are listed as follows:

### Enterprise-Extended Edition

DB2UDBEEE DB

### Enterprise Edition

DB2UDBEE DB2

### Warehouse Manager

DB2UDBWM DB2

### Relational Connect

DB2RELC DB2

### Spatial Extenders

DB2UDBGSE

4. Update the number of processors using:  
`db2licm -n [product password] [number of processors]`

---

## Upgrading DB2 from Try and Buy Mode

You can upgrade a DB2 product from the Try and Buy mode to a licensed version in two ways. You can use the command line utility called **db2licm** or use the License Center. The License Center is a built-in feature of the Control Center.

### On OS/2 and Windows 32-bit Operating Systems

To add a license using the command line:

1. Go to the `<install directory>\bin` directory, where `<install directory>` represents the directory where you installed the product.
2. To add a license, issue the following command:

```
db2licm path/filename.lic
```

License files are contained in the `db2/license` directory of the installation CD.

For more information about the **db2licm** command, refer to the *Command Reference*.

### Adding a license using the License Center

To add a license using the License Center:

1. Launch the Control Center.
2. Select **License Center** from the **Tools** menu list.
3. Refer to the online help for the License Center available from the Control Center.

---

## Uninstalling DB2

You can uninstall DB2 as follows:

1. Drop all databases.
2. Do one of the following:
  - Windows: go to the Add/Remove Products window, and select DB2.
  - OS/2: launch the Installation Utility in the IBM DB2 folder, select the product, and then select **Action** → **Delete**.
  - Go to the `sql1lib\bin` directory, and invoke the `db2unins` command.

If you cannot run the uninstall program, for example because the installation was unsuccessful, you can remove DB2 as follows:

1. Remove any partially installed files. For example, delete the files in the `c:\sql1lib` directory.
2. Under Windows, clean up the registry, using a registry editor such as **regedt32**. Delete the following entries:

```
HKEY_LOCAL_MACHINE\SOFTWARE\IBM\DB2
HKEY_CURRENT_USER\SOFTWARE\IBM\DB2
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\services\name
```

where *name* is one of the following:

- The instance name
- The instance name followed by `-N`
- **DB2REMOTECD**
- **DB2DAS00**
- **DB2GOVERNOR**
- **DB2NTSECSEVER**
- **DB2JDS**
- **DB2\_NT\_Performance**
- **DB2LICD**
- **DB2ControlCenterServer**
- **vwd**
- **vwkernel**
- **vwlogger**
- **EssbaseService**
- **OLAPIntegrationService**
- **DlfmService**



---

## Appendix B. National Language Support (NLS)

This section contains information about the National Language Support (NLS) provided by DB2, including information about supported languages and code pages. For information on developing applications that use NLS, refer to the *Application Development Guide*.

---

### Code Page and Language Support

During installation of DB2, the country, codepage, and regional settings are established. However, you can change these settings after installing DB2: including regional settings such as code page, country language (for monetary, date, and numeric formatting), and time zone. When a new connection to a database is made, the database manager uses these new values.

You must ensure that your regional settings are set correctly. DB2 may not produce the expected results if the country, code page, or regional settings are incorrect for the intended language. Table 8 shows the languages into which the DB2 messages are translated. If the setup is run on a machine that is set up using a non-supported language, then English is the default unless the user has specified otherwise.

*Table 8. Languages and Code Pages*

Country Code	Language
bg	Bulgarian
br	Brazilian Portuguese
cn	Simplified Chinese (PRC)
cz	Czech
de	German
dk	Danish
en	English
es	Spanish
fi	Finnish
fr	French
gr	Greek
hu	Hungarian
il	Hebrew
it	Italian

Table 8. Languages and Code Pages (continued)

Country Code	Language
jp	Japanese
kr	Korean
nl	Dutch
no	Norwegian
pl	Polish
pt	Portuguese
ru	Russian
se	Swedish
si	Slovenian
tr	Turkish
tw	Traditional Chinese (Taiwan)

---

## Appendix C. Naming Rules



Go to the section that describes the naming rules that you require information on:

- “General Naming Rules”
  - “Database, Database Alias, and Catalog Node Name Rules”
  - “Object Name Rules” on page 102
  - “Username, User ID, Group Name, and Instance Name Rules” on page 103
  - “Workstation Name (rname) Rules” on page 104
  - “DB2SYSTEM Naming Rules” on page 105
  - “Password Rules” on page 105
- 

---

### General Naming Rules

Unless otherwise specified, all names can include the following characters:

- A through Z. When used in most names, characters A through Z are converted from lowercase to uppercase.
- 0 through 9
- @, #, \$, and \_ (underscore)

Unless otherwise specified, all names must begin with one of the following characters:

- A through Z
- @, #, and \$

Do not use SQL reserved words to name tables, views, columns, indexes, or authorization IDs. For a list of SQL reserved words, refer to *SQL Reference*.

---

### Database, Database Alias, and Catalog Node Name Rules

*Database names* are the identifying names assigned to databases in the database manager. *Database alias names* are synonyms given to remote databases. Database aliases must be unique within the System Database Directory in which all aliases are stored. *Catalog node names* are the identifying names that are assigned to entries in the node directory. Each entry in the node directory is an alias name for a computer on your network. To avoid confusions that could arise from multiple names for the same server, we recommend that you use the same catalog node name as the network name for the server.

When naming a database, database alias, or catalog node name, see “General Naming Rules” on page 101. In addition, the name you specify can *only* contain 1 to 8 characters.



To avoid potential problems, do not use the special characters @, #, and \$ in a database name if you intend to have a client remotely connect to a host database. Also, because these characters are not common to all keyboards, do not use them if you plan to use the database in another country.

On Windows NT and Windows 2000 systems, ensure that no instance name is the same as a service name.

---

## Object Name Rules

Database objects include:

- Tables
- Views
- Columns
- Indexes
- User-defined functions (UDFs)
- User-defined types (UDTs)
- Triggers
- Aliases
- Table spaces
- Schemas

When naming database objects, see “General Naming Rules” on page 101.

In addition, the name you specify:

- Can contain 1 to 18 characters *except* for the following:
  - Table names (including view names, summary table names, alias names, and correlation names), which can contain up to 128 characters
  - column names, which can contain up to 30 characters
  - schema names, which can contain up to 30 characters
- Cannot be any of the SQL reserved words that are listed in the *SQL Reference*.

Using delimited identifiers, it is possible to create an object that violates these naming rules; however, subsequent use of the object could result in errors.



For example, if you create a column with a + or – sign included in the name and you subsequently use that column in an index, you will experience problems when you attempt to reorganize the table. To avoid potential problems with the use and operation of your database, *do not* violate these rules.

---

## Username, User ID, Group Name, and Instance Name Rules

*Usernames* or *User IDs* are the identifiers assigned to individual users. When naming users, groups, or instances, see “General Naming Rules” on page 101.

In addition to the general naming rules:

- User IDs on OS/2 can contain 1 to 8 characters. They cannot start with a numeric digit or end with \$.
- Usernames on UNIX can contain 1 to 8 characters.
- Usernames on Windows can contain 1 to 30 characters. The Windows NT and Windows 2000 operating systems currently have a limit of 20 characters.
- Group and instance names can contain 1 to 8 characters.
- Names cannot be any of the following:
  - USERS
  - ADMINS
  - GUESTS
  - PUBLIC
  - LOCAL
- Names cannot begin with:
  - IBM
  - SQL
  - SYS
- Names cannot include accented characters.
- In general, when naming users, groups, or instances:
  - OS/2** Use uppercase names.
  - UNIX** Use lowercase names.
  - Windows 32-bit operating systems**  
Use any case.

---

## Workstation Name (nname) Rules

A *workstation* name specifies the NetBIOS name for a database server or client that resides on the local workstation. This name is stored in the database manager configuration file. The workstation name is known as the *workstation nname*. When naming workstations, see “General Naming Rules” on page 101.

In addition, the name you specify:

- Can contain 1 to 8 characters
- Cannot include &, #, and @
- Must be unique within the network

In an partitioned database system, there is still only one workstation *nname* that represents the entire partitioned database system, but each node has its own derived unique NetBIOS *nname*.

The workstation *nname* that represents the partitioned database system is stored in the instance-owning database partition server’s database manager configuration file.

Each node’s unique *nname* is a derived combination of the workstation *nname* and the node number.

For a node that does not own an instance, its NetBIOS *nname* is derived as follows:

1. The first character of the instance-owning machine’s workstation *nname* is used as the first character of the node’s NetBIOS *nname*.
2. The next 1 to 3 characters represent the node number. The range is from 1 to 999.
3. The remaining characters are taken from instance-owning machine’s workstation *nname*. The number of remaining characters depend on the length of the instance-owning machine’s workstation *nname*. This number can be from 0 to 4.

For example:

Instance-Owning Machine’s <i>Workstation nname</i>	Node Number	Derived Node NetBIOS <i>nname</i>
GEORGE	3	G3ORGE
A	7	A7
B2	94	B942
N0076543	21	N216543
GEORGE5	1	G1RGE5



If you have changed the default workstation *nname* during the installation, the workstation *nname*'s last 4 characters should be unique across the NetBIOS network to minimize the chance of deriving a conflicting NetBIOS *nname*.

---

## DB2SYSTEM Naming Rules

DB2 uses the *DB2SYSTEM* name to identify a physical DB2 machine, system, or workstation within a network. On UNIX, the *DB2SYSTEM* name defaults to the TCP/IP hostname. On OS/2, you must specify the *DB2SYSTEM* name during install. On Windows 32-bit operating systems, you do not need to specify a *DB2SYSTEM* name; the DB2 setup program detects the Windows Computer name and assigns it to *DB2SYSTEM*.

When creating a *DB2SYSTEM* name, see "General Naming Rules" on page 101.

In addition, the name you specify:

- Must be unique within a network
- Can contain a maximum of 21 characters

---

## Password Rules

When determining passwords, consider the following rules:

**OS/2**                    A maximum of 14 characters.


**UNIX**                    A maximum of 8 characters.

**Windows 32-bit operating systems**  
A maximum of 14 characters.



# Appendix D. About DB2 Universal Database for UNIX, Windows, and OS/2

The DB2 Product Family provides relational database solutions for a wide range of computing devices, from small handheld devices up to the largest IBM mainframe.



## DB2 Products

The name DB2 is used for relational database products that run on a variety of platforms.

### DB2 Everywhere

DB2 Everywhere is a very small footprint database that runs on mobile devices, such as personal digital assistants (PDAs), smart phones, and

handheld personal computers (HPCs). Each mobile device stores a subset of the data from an enterprise database, which it can use without the need for a continuous database connection.

DB2 Everywhere Sync Server, running on a mid-tier server, replicates the data in both directions between mobile devices and the enterprise database. For example, each worker in a loading dock might carry a PDA that is periodically synchronized with an inventory database on OS/390.

## DB2 Universal Database

The following table shows which products are available on each platform:

Table 9. DB2 Universal Database Platforms

Edition	Windows 95/98	Windows NT/Windows 2000	OS/2	Linux	AIX	HP-UX	Solaris	PTX/NUMA-Q
Satellite	✓	✓						
Personal	✓	✓	✓	✓				
Workgroup		✓	✓	✓	✓	✓	✓	
Enterprise		✓	✓	✓	✓	✓	✓	✓
Enterprise - Extended		✓			✓	✓	✓	✓

**Note:** DB2 UDB Workgroup Edition, DB2 UDB Enterprise Edition, and DB2 UDB Enterprise - Extended Edition are commonly referred to as "servers" or "DB2 servers". Various clients are provided with each server product.

### Satellite Edition

DB2 UDB Satellite Edition is a single-user, small footprint version of DB2 available for Windows 32-bit operating systems. It is designed for occasionally connected remote systems, such as laptop computers.

Typically, many instances of DB2 UDB Satellite Edition are managed centrally by the same server.

### Personal Edition

DB2 UDB Personal Edition is a single-user version of the full DB2 product. It contains:

- An object-relational database engine
- Business intelligence support, through the OLAP Starter Kit
- Data warehouse support, through the Data Warehouse Center
- Multimedia support, through DB2 Extenders
- Access to a variety of IBM data sources, through DB2 DataJoiner
- Replication support, through DataPropagator
- Extended GUI administration tools, through DB2 Control Center

- An application development client
- An administration client

### **Workgroup Edition**

DB2 UDB Workgroup Edition is a multi-user version of the DB2 product, designed for a small business or departmental environment. It contains all the functionality of the Personal Edition, plus:

- The ability for remote clients to access data and perform administration on a DB2 workgroup server
- Web access, through Net.Data
- IBM WebSphere Application Server

### **Enterprise Edition**

DB2 UDB Enterprise Edition is designed for large databases with many users. It contains all the functionality of the Workgroup Edition, plus:

- A license for an unlimited number of client connections
- A license for an unlimited number of web client connections
- DB2 Connect support, with DRDA access to host DB2 systems

### **Enterprise - Extended Edition**

DB2 UDB Enterprise - Extended Edition is designed for the largest databases. It is ideal for scaling to very large databases for warehousing, data mining, and large-scale OLTP applications. It contains all the functionality of the Enterprise Edition, plus:

- Support for clusters of servers

### **Developer Editions**

Two special editions of DB2 are available for people who develop DB2 applications:

- DB2 Personal Developer's Edition
- DB2 Universal Developer's Edition

DB2 Personal Developer's Edition provides tools to help a software developer create applications for a DB2 Personal Edition database. DB2 Personal Developer's Edition contains all the functionality of DB2 UDB Personal Edition, plus:

- The functionality of DB2 Connect Personal Edition
- VisualAge for Java, Entry Edition

DB2 Universal Developer's Edition provides the tools for developing client/server applications. DB2 Universal Developer's Edition contains all the functionality of DB2 UDB Enterprise Edition, plus:

- VisualAge for Java, Professional Edition

## Host Databases


The following DB2 products are available for IBM midrange and mainframe systems:

- DB2 UDB for AS/400
- DB2 Server for VSE & VM
- DB2 UDB for OS/390

---

## DB2 Connect

DB2 Connect provides connectivity to mainframe and midrange databases from Windows, OS/2, and UNIX-based platforms. You can connect to DB2 databases on OS/400, VSE, VM, MVS, and OS/390. You can also connect to non-IBM databases that comply with the Distributed Relational Database Architecture (DRDA).



The following DB2 Connect products are available:

- Personal Edition
- Enterprise Edition
- Unlimited Edition



*DB2 Connect Personal Edition* provides a direct connection from one Windows, OS/2, or Linux operating system to mainframe and midrange databases. It is designed for a two-tier environment, where each client connects directly to the host. DB2 Connect Personal Edition does not accept inbound client requests for data.

*DB2 Connect Enterprise Edition*, installed on a gateway server, connects an entire LAN to mainframe and midrange databases. It is designed for a three-tier environment, where clients connect to a host through a gateway server.

*DB2 Connect Unlimited Edition* provides an unlimited number of DB2 Connect Personal Edition and DB2 Connect Enterprise Edition licenses. You get all of these licenses for one price, based on the size of the OS/390 system that is being accessed.

---

## Related Products

The following products work together with DB2 Universal Database.

### **DB2 Relational Connect**

DB2 Relational Connect gives you the ability to access federated data by joining DB2 data with Oracle data within a query. It works under DB2 UDB Enterprise Edition or Enterprise - Extended Edition on Windows NT, Windows 2000, and AIX.

### **DB2 Warehouse Manager**

DB2 Warehouse Manager provides a warehouse or data mart administrator with enhanced management capabilities. It provides the ability to manage the process of moving data as well as manage the ad hoc query workload going against the warehouse or mart. This product includes:

- Warehouse agents, which manage the flow of data between sources and warehouse targets.
- Warehouse transformers, which clean and transform data moving into the warehouse.
- An integrated business information catalog, which guides users to the data that they want.
- Metadata interchange with end-user repositories and CASE tools.
- A sophisticated query management and workload distribution tool, DB2 Query Patroller, which was previously a separate product. In order to use Query Patroller, you must have a Query Patroller server installed. Query Patroller client support is only available with a DB2 Administration Client. For more information, refer to *DB2 Query Patroller Administration Guide*.

This product is available for DB2 UDB Enterprise Edition and Enterprise - Extended Edition systems.

## **DB2 OLAP Server**

DB2 OLAP Server provides fast, intuitive online analytic processing of data to afford users quick access to information. OLAP servers process multidimensional requests that calculate, consolidate, and retrieve information from multidimensional and relational databases.

With DB2 OLAP Server you can:

- Build analytical applications using built-in math, financial and statistical functions to define data in multidimensional cubes.
- View multidimensional data from various perspectives.
- Dynamically add new data dimensions, modify dimensional hierarchies, and change calculations.
- Use spreadsheets and web browsers as clients to analyze data without using separate SQL queries.
- Ensure data security by defining levels of access for individuals.

The OLAP Starter Kit, which supports only three concurrent users, is installed with DB2 Universal Database. You can later upgrade the OLAP Starter Kit to the full DB2 OLAP Server product.

## **Intelligent Miner**

The Intelligent Miner family consists of two products:

- DB2 Intelligent Miner for Data extracts valuable information from a large quantity of data in a relational database.
- Intelligent Miner for Text works with unstructured information, such as text files, email, and web pages.

## **DB2 Spatial Extender**

DB2 Spatial Extender lets you integrate geographical data with your existing business data. It includes:

- Data types such as points, lines, and polygons
- Functions such as area, endpoint, and intersect
- An indexing scheme for spatial data

This product is available for DB2 UDB Enterprise Edition and Enterprise - Extended Edition systems.

## **DB2 Net Search Extender**

DB2 Net Search Extender contains a DB2 stored procedure that adds the power of fast full-text retrieval to Net.Data, Java, or DB2 CLI applications. It offers application programmers a variety of search functions, such as fuzzy

search, stemming, boolean operators, and section search. Searching using DB2 Net Search Extender can be particularly advantageous in the Internet, when search performance on large indexes and scalability according to concurrent queries are important factors.

### **DB2 Data Links Manager**

DB2 Data Links Manager provides referential integrity, access control, and recovery capabilities for files that physically reside on file systems external to a DB2 Universal Database. Data Links technology includes the DATALINK data type, implemented as an SQL data type in DB2 Universal Database, which references an object stored external to a database. Data Links Manager is available on Windows NT and AIX systems. On AIX, it can be used with native (or JFS) filesystems, or in Transarc DCE-DFS file server environments.

### **Tivoli Enterprise**

Tivoli Enterprise is a suite of management applications that let you manage an entire enterprise environment, including the data center, distributed systems, and mobile laptops as a single business unit. DB2 has been certified as Tivoli Ready.

---

## **Working with DB2 Data**

DB2 is a relational database system rich in features, many of which can be accessed remotely. In addition to allowing you to store your data, DB2 lets you issue requests to administer, query, update, insert, or delete data using local or remote client applications.

### **Accessing DB2 Data from Remote Clients**

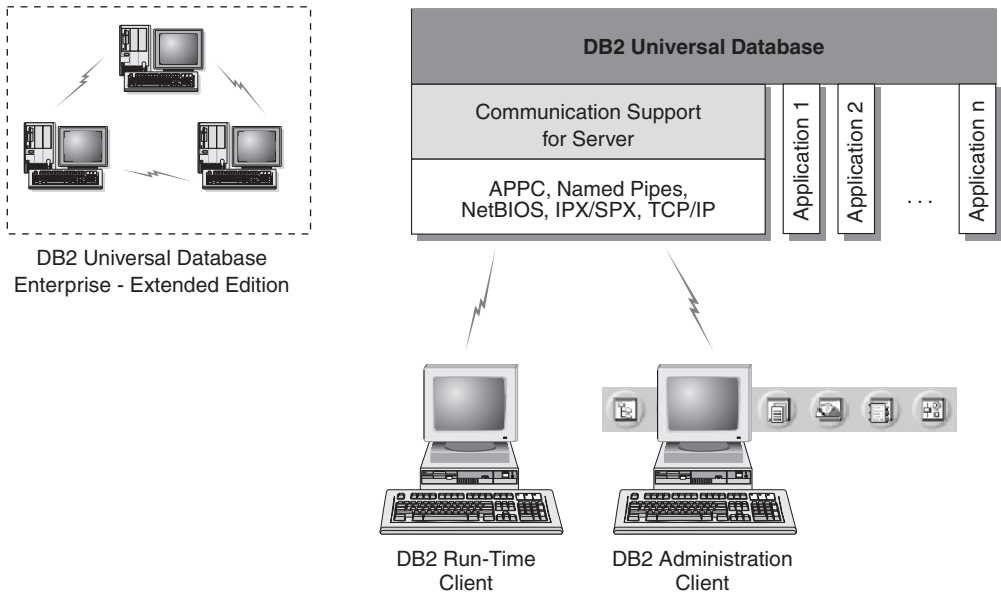
DB2 clients provide a run-time environment that enables client applications to access one or more remote databases. With a DB2 Administration Client, you can remotely administer DB2 or DB2 Connect servers. All applications must access a database through a DB2 client. A Java applet can access a remote database through a Java-enabled browser.

DB2 Version 7 clients are supported on:

- OS/2
- UNIX (AIX, HP-UX, Linux, NUMA-Q, SGI IRIX, and the Solaris Operating Environment)
- Windows 9x, Windows NT, or Windows 2000

Figure 5 on page 114 shows a server that is being accessed by local and remote applications. Remote applications must have the appropriate DB2 client installed to enable applications to access data on the remote server.

## DB2 Universal Database Enterprise - Extended Edition - Remote Client Support



Not all protocols are supported for all platforms.

Figure 5. DB2 Server with Local Applications and Remote Clients

### Accessing Multiple DB2 Servers

Once a network is operational and protocols are functional on all workstations, LAN-to-LAN connections between DB2 servers and clients require no additional software.

For example, you can have a server on a Windows NT workstation that is connected to a LAN, and another server on a UNIX workstation connected to a LAN. As long as there is a connection between the two LANs, clients on either network can access either server. See Figure 6 on page 115.

## Communicating Between DB2 Universal Database Servers




Figure 6. Accessing Data on Multiple Servers

Within a single transaction, databases on both servers are accessed and updated, and the integrity of the data on both servers is maintained. This is commonly known as two-phase commit, or distributed-unit-of-work access. Refer to the *Administration Guide* for more information.

### Accessing Host or AS/400 DB2 Data from the Desktop using DB2 Connect Enterprise Edition

A DB2 server with the DB2 Connect Server Support feature installed, or a DB2 Connect server, enables DB2 clients on a LAN access to data that is stored on host or AS/400 systems.

A great deal of the data in many large organizations is managed by DB2 for AS/400, DB2 for MVS/ESA, DB2 for OS/390, or DB2 for VSE & VM. Applications that run on any of the supported platforms can work with this data transparently, as if a local database server managed it. DB2 Connect Enterprise Edition is required for supporting applications which access host or AS/400 data and exploit transaction monitors (for example, IBM TxSeries CICS and Encina Monitor, Microsoft Transaction Server, BEA Tuxedo) as well as applications that are implemented as Java applets.

In addition, you can use a wide range of off-the-shelf or custom-developed database applications with DB2 Connect and its associated tools. For example, you can use DB2 Connect products with:

- *Spreadsheets*, such as Lotus 1-2-3 and Microsoft Excel, to analyze real-time data without having the cost and complexity of data extract and import procedures.
- *Decision support tools*, such as BusinessObjects, Brio and Impromptu, and Crystal Reports, to provide real-time information.
- *Database products*, such as Lotus Approach and Microsoft Access.
- *Development tools*, such as PowerSoft PowerBuilder, Microsoft Visual Basic, and Borland Delphi, to create client/server solutions.

DB2 Connect Enterprise Edition is most appropriate for environments where:

- Host and AS/400 database servers do not support native TCP/IP connectivity and direct connectivity from desktop workstations via SNA is not desirable.
- Application is implemented using data-aware Java applets.
- Web servers are used to implement web-based applications.
- Middle-tier application server is employed.
- Transaction monitors such as IBM TxSeries CICS and Encina Monitor, IBM Component Broker, IBM MQSeries, Microsoft Transaction Server (MTS), and BEA Tuxedo are used.

DB2 Connect provides transparent access to host or AS/400 data through a standard architecture for managing distributed data. This standard is known as Distributed Relational Database Architecture (DRDA). DRDA allows your applications to establish a fast connection to host and AS/400 databases without expensive host components or proprietary gateways.

Although DB2 Connect is often installed on an intermediate server machine to connect DB2 clients to a host or AS/400 database, it is also installed on machines where multiple local users want to access the host or AS/400 servers directly. For example, DB2 Connect may be installed on a large machine with many local users.

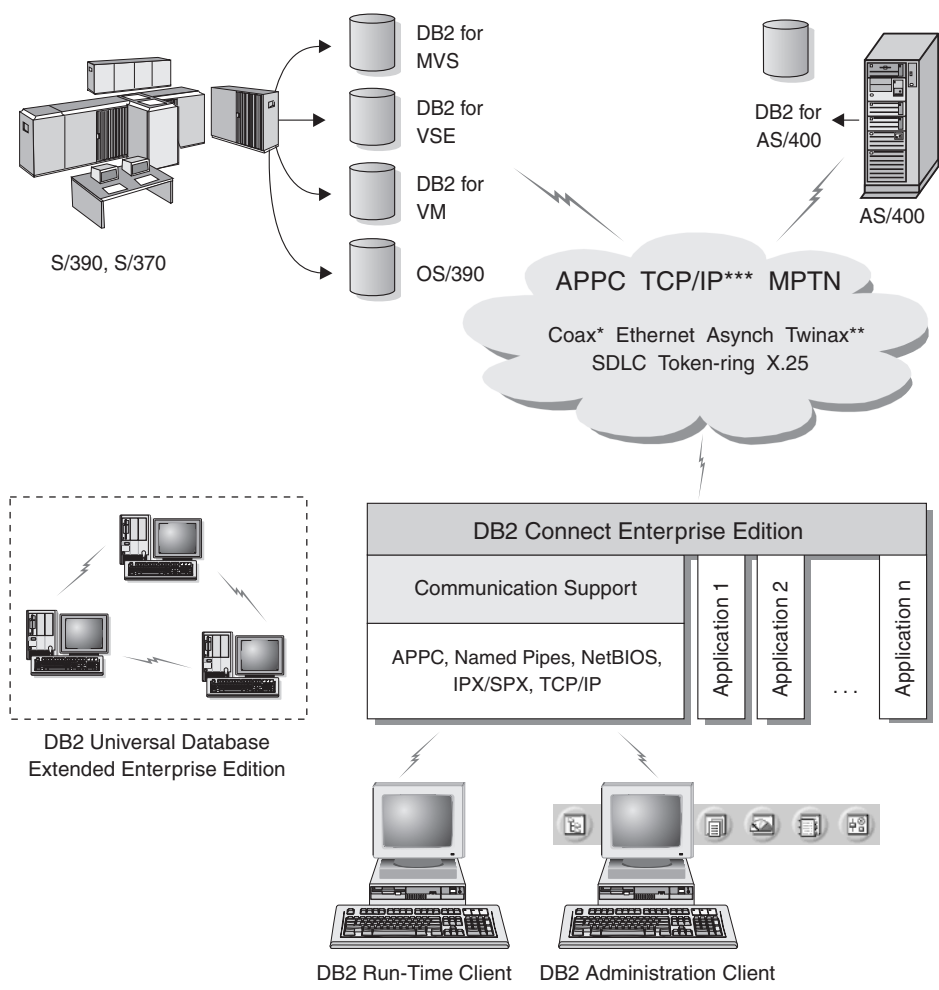
DB2 Connect may also be installed on a Web server, Transaction Processor (TP) monitor, or other 3-tier application server machines with multiple local SQL application processes and threads. In these cases, you can choose to install DB2 Connect on the same machine for simplicity, or on a separate machine to off-load CPU cycles.

A DB2 server with the DB2 Connect functionality installed, or a DB2 Connect server, enables multiple clients to connect to host or AS/400 data and can significantly reduce the effort that is required to establish and maintain access

to enterprise data. Figure 7 on page 118 illustrates IBM's solution for environments in which you want to use a DB2 client making an indirect connection to a host or AS/400 database server through DB2 Connect Enterprise Edition.

In the example, you could replace the DB2 Connect server with a DB2 server that has the DB2 Connect Server Support component installed.

## DB2 Connect Enterprise Edition



Not all protocols are supported for all platforms.

- \* For Host connections only
- \*\* For AS/400
- \*\*\* TCP/IP connectivity requires DB2 for OS/390 V5R1, DB2 for AS/400 V4R2, or DB2 for VM V6.1

Figure 7. DB2 Connect Enterprise Edition



## Accessing DB2 Data from the Web using Java

Java Database Connectivity (JDBC) and Embedded SQL for Java (SQLJ) are provided with DB2 to allow you to create applications that access data in DB2 databases from the Web.

Programming languages containing embedded SQL are called host languages. Java differs from the traditional host languages C, COBOL, and FORTRAN, in ways that significantly affect how it embeds SQL:

- SQLJ and JDBC are open standards, enabling you to easily port SQLJ or JDBC applications from other standards-compliant database systems to DB2 Universal Database.
- All Java types representing composite data, and data of varying sizes, have a distinguished value, `null`, which can be used to represent the SQL NULL state, giving Java programs an alternative to NULL indicators that are a fixture of other host languages.
- Java is designed to support programs that, by nature, are heterogeneously portable (also called "super portable" or simply "downloadable"). Along with Java's type system of classes and interfaces, this feature enables component software. In particular, an SQLJ translator written in Java can call components that are specialized by database vendors in order to leverage existing database functions such as authorization, schema checking, type checking, transactional, and recovery capabilities, and to generate code optimized for specific databases.
- Java is designed for binary portability in heterogeneous networks, which promises to enable binary portability for database applications that use static SQL.
- You can run JDBC applets inside a web page on any system with a Java-enabled browser, regardless of the platform of your client. Your client system requires no additional software beyond this browser. The client and the server share the processing of JDBC and SQLJ applets and applications.

The DB2 JDBC Applet server and the DB2 client must reside on the same machine as the Web server. The DB2 JDBC Applet server calls the DB2 client to connect to local, remote, host, and AS/400 databases. When the applet requests a connection to a DB2 database, the JDBC client opens a TCP/IP connection to the DB2 JDBC Applet on the machine where the Web server is running. See Figure 8 on page 120 for an example of a Java-enabled browser accessing data from remote DB2 databases.

## Accessing DB2 Data Using JDBC




Figure 8. Accessing Data Stored on DB2 using JDBC

JDBC and SQLJ applications can be run from any system that has a DB2 client installed; a Web browser and a Web server are not required.

For more information on Java enablement, refer to the DB2 Java Enablement web page at <http://www.ibm.com/software/data/db2/java/>

For more information on the JDBC API, point your browser to <http://splash.javasoft.com/>

### Accessing DB2 Data from the Web using Net.Data

Net.Data is provided with DB2 to allow you to create applications that access data in DB2 databases from the Web. You can use Net.Data to create applications that are stored on a Web server and viewable from any Web browser. While viewing these documents, users can either select automated queries or define new ones that retrieve the specified information directly from a DB2 database.

Automated queries do not require user input; they are links in an HTML document and, when selected, they trigger existing SQL queries and return

the results from a DB2 database. These links can be triggered repeatedly to access current DB2 data. Customized queries require user input. Users define the search characteristics on the Web page by selecting options from a list or by entering values in fields. They submit the search by clicking on a push button. Net.Data uses the information that is supplied by the user to dynamically build a complete SQL statement, and it sends the query to the DB2 database.

A demonstration of Net.Data applications is available from the IBM Software Net.Data page at <http://www.ibm.com/software/data/net.data>

Net.Data can be installed with a:

- DB2 server to allow local access to databases.
- DB2 client to allow remote access to databases.

In both cases, Net.Data and the Web server must be installed on the same system. See Figure 9 for an example of a workstation with Net.Data that is being used to access data from a remote DB2 database.

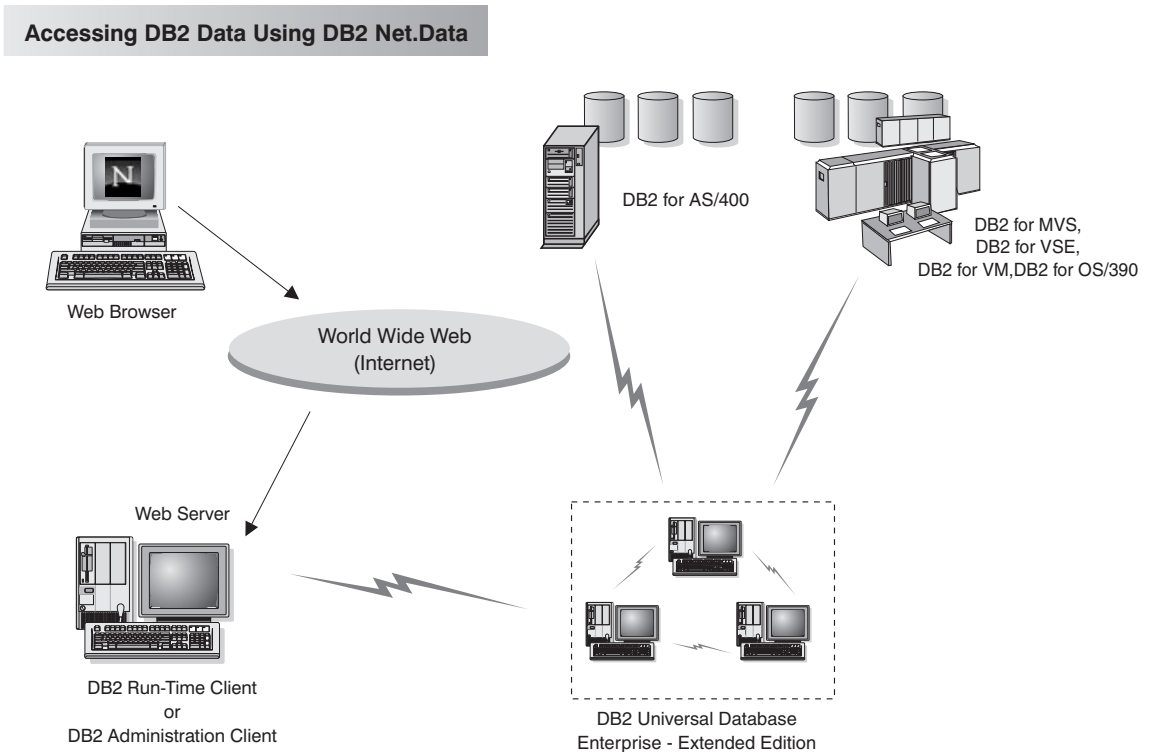


Figure 9. Accessing Internet Data Stored on DB2 using Net.Data

## Accessing DB2 Data from Host and AS/400 Client Machines

The DRDA Application Server (DRDA AS) functionality gives clients or applications on host and AS/400 machines transparent access to data stored on a LAN-based DB2 Universal Database server. This access is provided through Distributed Relational Database Architecture (DRDA), a standard architecture for managing data. You can configure your server to act as a DRDA AS for host and AS/400 clients or applications; these clients or applications are known as DRDA Application Requesters (DRDA AR).



The DRDA AS feature is available for:

- DB2 Universal Database Enterprise - Extended Edition
- DB2 Universal Database Enterprise Edition
- DB2 Universal Database Workgroup Edition

For information on how to set up your DB2 Universal Database server as a DRDA AS, refer to the *Installation and Configuration Supplement*.

---

## Administering Instances and Databases with the DB2 Administration Tools

You can administer local or remote servers using the DB2 Administration Tools. Use the Control Center to perform administration tasks such as configuring DB2 instances and databases, backing up and recovering data, scheduling jobs, and managing media, all from a graphical interface.

The Control Center uses TCP/IP to retrieve status, connection, and snapshot information from each database partition server (node). You assign a TCP/IP port number for the Performance Monitor component of the Control Center during installation.

### Managing Instances and Database Objects using the Control Center

The Control Center displays instances and database objects (such as table spaces, tables, and packages) and their relationships to each other. Using the Control Center, you can manage local and remote servers from a single point of control. See Figure 10 on page 123 for an example of the main Control Center window.




Figure 10. Control Center Main Window

The Control Center distinguishes between single-partition and multipartition database systems via *Discovery*. Discovery uses the *DB2SYSTEM*, *DB2ADMINSERVER*, and *DB2COMM* registry values. For more information on these registry values, refer to the *Administration Guide*.

From the Control Center, you can perform operations on database objects. These operations include:

- Create and drop a database
- Create, alter, and drop a table space or table
- Create, alter, and drop an index
- Create nodegroups
- Backup and recover a database partition or table space partition.
- Define the replication sources and subscriptions to replicate data between systems
- Monitor resources and events on a server.


For more information on objects in a partitioned database system, see “Chapter 1. Introduction to DB2 Enterprise - Extended Edition” on page 3.

You can also control DB2 instances by:

- Maintaining communication protocols
- Setting database manager and database configuration values that affect performance.

Wizards are provided to help you perform complex tasks. For example, a wizard is available to tune the performance of your system. See “Using DB2 Wizards” on page 146 for descriptions of the various wizards and how to start them.

The Control Center provides additional functionality to assist you in managing your servers:



### Control Center

Use the Control Center to start another session of the Control Center to administer a server.

### Satellite Center

Use the Satellite Center to manage the Satellites that are served by a particular DB2 Control Server. It provides create, remove, modify, and manage functions for Satellites and Groups. You can also create and manage scripts to administer the Satellites.

### Data Warehouse Center

Use the Data Warehouse Center to manage warehouses; create and manage warehouse objects, such as sources and targets; define, extract, transform, manipulate, and load steps and processes; and schedule and automate steps.

### Command Center

Use the Command Center to enter DB2 commands and SQL statements in an interactive window and see the execution result in a result window. You can scroll through the results and save the output to a file.

### Script Center

Use the Script Center to create scripts, which you can store and invoke at a later time. These scripts can contain DB2 commands, SQL statements, as well as operating system commands. Scripts can be scheduled to run unattended. These jobs can be run once or set up to run on a repeating schedule; a repeating schedule is particularly useful for tasks like backup.

## Alert Center

Use the Alert Center to monitor your system for early warnings of potential problems or to automate actions to correct problems discovered.

## Journal

Use the Journal to view all available information about jobs that are pending execution, executing, or that have completed execution. You can also view the recovery history log, the alerts log, and the messages log; and review the results of jobs that are run unattended.

## License Center

Use the License Center to manage licenses and display license status and usage of any DB2 products installed on your system. You can also use the License Center to configure your system for proper license monitoring.

## Stored Procedure Builder

Use Stored Procedure Builder to create stored procedures, build stored procedures on local and remote DB2 servers, modify and rebuild existing stored procedures, and run stored procedures for testing and debugging the execution of installed stored procedures.

## Tools Setting

Use the Tools Setting to change the settings for the DB2 Administration Tools.

## Information Center

The Information Center provides quick access to DB2 product information. This product information includes such items as: database tasks, reference material, DB2 documentation, warehouse administration information, troubleshooting aids, sample programs for application development, and DB2 web-related URLs.

You can also analyze performance using the DB2 Performance Monitor and Visual Explain. These tools are available from the Control Center.



Use the **DB2 Performance Monitor** to monitor the performance of your system. You can monitor activity by sampling data over a period of time or using data for a particular event. See “Monitoring Databases using DB2 Performance Monitor” on page 126 for more information.



Use **Visual Explain** to view the access plan for explained SQL statements as a graph. You can use the information available from the graph to tune your SQL queries for better performance. See “Viewing SQL Access Plans using Visual Explain” on page 126 for more information.

You can find additional information in the *Administration Guide* or in the online help.

## Managing Communications on the Server

The Control Center allows you to maintain or edit server protocol settings in the database manager configuration file.

- Configure database manager parameters by clicking with the right mouse button on an instance and selecting the **Configure** option from the pop-up menu. By default, the setup program automatically detects and configures most communication protocols that it detects on your system.
- Export database information in a profile that can be used to configure clients by clicking with the right mouse button on a system and selecting the **Export Server Profile** option from the pop-up menu.

For information on how to configure server communications, refer to the *Installation and Configuration Supplement*.

## Monitoring Databases using DB2 Performance Monitor

With the DB2 Performance Monitor, you can:

- Identify and analyze performance problems in database applications or the database manager.
- Use the early warning system to detect potential problems.
- Automate actions to correct problems that are discovered.
- Define your own statistics, in addition to the default set that is provided.

You can choose to monitor the current state of database activity or collect information when specific events occur. The Performance Monitor allows you to capture point-in-time information at specified intervals. The Event Analyzer allows you to view information about the occurrence of events such as deadlocks and transaction completions.

For additional information, refer to the *Administration Guide* or the online help. You are also able to use the Windows Performance Monitor (supported on Windows NT and Windows 2000) to monitor both database and system performance. For information on how to register DB2 resources and to use the Windows Performance Monitor, refer to the *Administration Guide*.

## Viewing SQL Access Plans using Visual Explain

Visual Explain helps database administrators and application developers to:

- View the access plan chosen by the database manager's optimizer for a given SQL statement.
- Tune SQL statements for better performance.
- Design application programs and databases.
- View all the details of an access plan, including the statistics in the system catalogs.



- Decide whether or not to add an index to a table.
- Identify the source of problems by analyzing the access plan or performance of SQL statements.
- Use the portable snapshot function to view snapshots from any remote DB2 server.
- Display access plans for queries on all supported DB2 configurations.

For additional information, refer to the *Administration Guide* or the online help.

---

## Managing Connections to Databases using the Client Configuration Assistant

The Client Configuration Assistant (CCA) helps you manage your database connections to remote servers. The CCA is available on OS/2 and Windows 32-bit operating systems and this is the preferred method to set up any OS/2, Windows 9x, Windows NT, or Windows 2000 client to communicate with a server.

You can use the command line processor to set up DB2 clients on any platform. For more information, refer to the *Installation and Configuration Supplement*.

With the CCA, you can:

- Catalog databases so that they can be used by applications. Three methods are available:
  - Use a profile provided by a database administrator to automatically define your connections. Client access is automatically set up for that database.
  - Search the network for available databases and select one. Client access is automatically set up for that database.
  - Manually configure a connection to a database by entering the required connection parameters.
- Remove cataloged databases, or change the properties of a cataloged database.
- Export and import client profiles that contain database and configuration information for a client.
- Test connections to local or remote databases identified on your system.
- Bind applications to a database by selecting utilities or bind files from a list.
- Tune the client configuration parameters on your system. Parameters are logically grouped and suggested settings are provided on the interface as parameters are selected.
- Export client configuration information to a profile.
- Import configuration information from a profile.

- Update the server password.

---

## Managing Warehouses using the Data Warehouse Center

DB2 Universal Database offers the Data Warehouse Center, a component that automates data warehouse processing. You can use the Data Warehouse Center to define the data to include in the warehouse. Then, you can use the Data Warehouse Center to schedule automatic refreshes of the data in the warehouse.

From the Data Warehouse Center, you can manage specific warehousing objects, including subject areas, warehouse sources, warehouse targets, agents, agent sites, steps, and processes.

You can also perform the following tasks from the Data Warehouse Center:

- Define a subject area. You use a subject area to logically group the processes that are related to a particular topic or function.
- Explore the source data and define warehouse sources.
- Create database tables and define warehouse targets.
- Define a process that specifies how to move and transform the source data into the appropriate format for the warehouse.
- Test and schedule steps.
- Define security and monitor database currency.
- Define a star schema model.

---

## Understanding the Administration Server

The Administration Server responds to requests from the DB2 Administration Tools and the Client Configuration Assistant (CCA). The DB2 Administration Tools allow you to start, stop, and set database manager configuration parameters for servers. The Administration Server is used by the CCA to catalog databases for a client.

The Administration Server (DAS) must reside on every server that you want to administer and detect. The Administration Server is automatically created and started for you; its default name is DB2DAS00.

The setup program creates the Administration Server on the instance-owning machine and automatically starts it at boot time. The default name of the Administration Server is DB2DAS00. For more information, see *Administration Guide*.

---

## Developing Applications using the DB2 Application Development Client

The DB2 Application Development Client is a collection of tools that are designed to meet the needs of database application developers. It includes libraries, header files, documented APIs, and sample programs to build character-based, multimedia, or object-oriented applications.

A platform-specific version of the DB2 Application Development Client is available on each server CD-ROM. In addition the Developer Edition boxes contain the Application Development Clients for multiple supported operating systems. The Personal Developer's Edition box contains the Application Development CD-ROMs for OS/2, Windows and Linux. The Universal Developer's Edition box contains the Application Development CD-ROMs for all supported operating systems.

Through a DB2 client, these applications can access all servers and, by using the DB2 Connect product (or the DB2 Connect functionality supplied with DB2 Enterprise - Extended or DB2 Enterprise Edition), they can also access DB2 Universal Database for AS/400, DB2 Universal Database for OS/390, and DB2 for VSE & VM database servers.

The DB2 Application Development Client allows you to develop applications that use the following interfaces:

- Embedded SQL
- Call Level Interface (CLI) development environment (which is compatible with ODBC from Microsoft)
- Java Database Connectivity (JDBC)
- Embedded SQL for Java (SQLJ)
- DB2 Application Programming Interfaces (APIs) that use administrative functions to manage a DB2 database.

The DB2 Application Development Client includes:

- Precompilers for Java, C, C++, COBOL, and FORTRAN.
- Libraries, include files, and code samples to develop applications that use SQLJ and DB2 CLI.
- A single control point for metadata management through the use of templates and tokens.
- JDBC and SQLJ support to develop Java applications and applets.
- Interactive SQL, through the CLP, to prototype SQL statements and perform ad-hoc database queries.
- An API to enable other application development tools to implement precompiler support for DB2 directly with their products.

- An SQL92 and MVS Conformance Flagger to identify embedded SQL statements in applications not conforming to the ISO/ANSO SQL92 Entry Level standard, or which are not supported by DB2 for OS/390.

For complete information on the functionality of the DB2 Application Development Client, and instructions on how to use them, as well as a full list of supported compilers for your platform, refer to the *Application Building Guide*.

---

## Running Your Own Applications

Various types of applications can access DB2 databases:

- Applications developed using a DB2 Application Development Client that include embedded SQL (including Java SQLJ applications and applets), APIs, stored procedures, user-defined functions, calls to DB2 CLI, or calls to JDBC applications and applets.
- ODBC applications such as Lotus Approach.
- Net.Data macros containing HTML and SQL.

The DB2 CLI/ODBC driver is an optional component during a DB2 client install. It is required to run CLI, ODBC, JDBC, and some SQLJ applications.

For more information on running your own applications, refer to the *Installation and Configuration Supplement*.

---

## Appendix E. Using the DB2 Library

The DB2 Universal Database library consists of online help, books (PDF and HTML), and sample programs in HTML format. This section describes the information that is provided, and how you can access it.

To access product information online, you can use the Information Center. For more information, see “Accessing Information with the Information Center” on page 145. You can view task information, DB2 books, troubleshooting information, sample programs, and DB2 information on the Web.

---

### DB2 PDF Files and Printed Books

#### DB2 Information

The following table divides the DB2 books into four categories:

##### **DB2 Guide and Reference Information**

These books contain the common DB2 information for all platforms.

##### **DB2 Installation and Configuration Information**

These books are for DB2 on a specific platform. For example, there are separate *Quick Beginnings* books for DB2 on OS/2, Windows, and UNIX-based platforms.

##### **Cross-platform sample programs in HTML**

These samples are the HTML version of the sample programs that are installed with the Application Development Client. The samples are for informational purposes and do not replace the actual programs.

##### **Release notes**

These files contain late-breaking information that could not be included in the DB2 books.

The installation manuals, release notes, and tutorials are viewable in HTML directly from the product CD-ROM. Most books are available in HTML on the product CD-ROM for viewing and in Adobe Acrobat (PDF) format on the DB2 publications CD-ROM for viewing and printing. You can also order a printed copy from IBM; see “Ordering the Printed Books” on page 141. The following table lists books that can be ordered.

On OS/2 and Windows platforms, you can install the HTML files under the `sql1lib\doc\html` directory. DB2 information is translated into different

languages; however, all the information is not translated into every language. Whenever information is not available in a specific language, the English information is provided

On UNIX platforms, you can install multiple language versions of the HTML files under the `doc/%L/html` directories, where `%L` represents the locale. For more information, refer to the appropriate *Quick Beginnings* book.

You can obtain DB2 books and access information in a variety of ways:

- “Viewing Information Online” on page 144
- “Searching Information Online” on page 148
- “Ordering the Printed Books” on page 141
- “Printing the PDF Books” on page 140

Table 10. DB2 Information

Name	Description	Form Number PDF File Name	HTML Directory
<b>DB2 Guide and Reference Information</b>			
<i>Administration Guide</i>	<i>Administration Guide: Planning</i> provides an overview of database concepts, information about design issues (such as logical and physical database design), and a discussion of high availability.	SC09-2946 db2d1x70	db2d0
	<i>Administration Guide: Implementation</i> provides information on implementation issues such as implementing your design, accessing databases, auditing, backup and recovery.	SC09-2944 db2d2x70	
	<i>Administration Guide: Performance</i> provides information on database environment and application performance evaluation and tuning.	SC09-2945 db2d3x70	
	You can order the three volumes of the <i>Administration Guide</i> in the English language in North America using the form number SBOF-8934.		
<i>Administrative API Reference</i>	Describes the DB2 application programming interfaces (APIs) and data structures that you can use to manage your databases. This book also explains how to call APIs from your applications.	SC09-2947 db2b0x70	db2b0

Table 10. DB2 Information (continued)

Name	Description	Form Number PDF File Name	HTML Directory
<i>Application Building Guide</i>	Provides environment setup information and step-by-step instructions about how to compile, link, and run DB2 applications on Windows, OS/2, and UNIX-based platforms.	SC09-2948 db2axx70	db2ax
<i>APPC, CPI-C, and SNA Sense Codes</i>	Provides general information about APPC, CPI-C, and SNA sense codes that you may encounter when using DB2 Universal Database products.	No form number db2apx70	db2ap
	Available in HTML format only.		
<i>Application Development Guide</i>	Explains how to develop applications that access DB2 databases using embedded SQL or Java (JDBC and SQLJ). Discussion topics include writing stored procedures, writing user-defined functions, creating user-defined types, using triggers, and developing applications in partitioned environments or with federated systems.	SC09-2949 db2a0x70	db2a0
<i>CLI Guide and Reference</i>	Explains how to develop applications that access DB2 databases using the DB2 Call Level Interface, a callable SQL interface that is compatible with the Microsoft ODBC specification.	SC09-2950 db2l0x70	db2l0
<i>Command Reference</i>	Explains how to use the Command Line Processor and describes the DB2 commands that you can use to manage your database.	SC09-2951 db2n0x70	db2n0
<i>Connectivity Supplement</i>	Provides setup and reference information on how to use DB2 for AS/400, DB2 for OS/390, DB2 for MVS, or DB2 for VM as DRDA application requesters with DB2 Universal Database servers. This book also details how to use DRDA application servers with DB2 Connect application requesters.	No form number db2h1x70	db2h1
	Available in HTML and PDF only.		

Table 10. DB2 Information (continued)

Name	Description	Form Number PDF File Name	HTML Directory
<i>Data Movement Utilities Guide and Reference</i>	Explains how to use DB2 utilities, such as import, export, load, AutoLoader, and DPROF, that facilitate the movement of data.	SC09-2955 db2dmx70	db2dm
<i>Data Warehouse Center Administration Guide</i>	Provides information on how to build and maintain a data warehouse using the Data Warehouse Center.	SC26-9993 db2ddx70	db2dd
<i>Data Warehouse Center Application Integration Guide</i>	Provides information to help programmers integrate applications with the Data Warehouse Center and with the Information Catalog Manager.	SC26-9994 db2adx70	db2ad
<i>DB2 Connect User's Guide</i>	Provides concepts, programming, and general usage information for the DB2 Connect products.	SC09-2954 db2c0x70	db2c0
<i>DB2 Query Patroller Administration Guide</i>	Provides an operational overview of the DB2 Query Patroller system, specific operational and administrative information, and task information for the administrative graphical user interface utilities.	SC09-2958 db2dwx70	db2dw
<i>DB2 Query Patroller User's Guide</i>	Describes how to use the tools and functions of the DB2 Query Patroller.	SC09-2960 db2wwx70	db2ww
<i>Glossary</i>	Provides definitions for terms used in DB2 and its components.  Available in HTML format and in the <i>SQL Reference</i> .	No form number db2t0x70	db2t0
<i>Image, Audio, and Video Extenders Administration and Programming</i>	Provides general information about DB2 extenders, and information on the administration and configuration of the image, audio, and video (IAV) extenders and on programming using the IAV extenders. It includes reference information, diagnostic information (with messages), and samples.	SC26-9929 dmbu7x70	dmbu7
<i>Information Catalog Manager Administration Guide</i>	Provides guidance on managing information catalogs.	SC26-9995 db2dix70	db2di



Table 10. DB2 Information (continued)

Name	Description	Form Number PDF File Name	HTML Directory
<i>Information Catalog Manager Programming Guide and Reference</i>	Provides definitions for the architected interfaces for the Information Catalog Manager.	SC26-9997 db2bix70	db2bi
<i>Information Catalog Manager User's Guide</i>	Provides information on using the Information Catalog Manager user interface.	SC26-9996 db2aix70	db2ai
<i>Installation and Configuration Supplement</i>	Guides you through the planning, installation, and setup of platform-specific DB2 clients. This supplement also contains information on binding, setting up client and server communications, DB2 GUI tools, DRDA AS, distributed installation, the configuration of distributed requests, and accessing heterogeneous data sources.	GC09-2957 db2iyx70	db2iy
<i>Message Reference</i>	Lists messages and codes issued by DB2, the Information Catalog Manager, and the Data Warehouse Center, and describes the actions you should take.  You can order both volumes of the Message Reference in the English language in North America with the form number SBOF-8932.	Volume 1 GC09-2978 db2m1x70 Volume 2 GC09-2979 db2m2x70	db2m0
<i>OLAP Integration Server Administration Guide</i>	Explains how to use the Administration Manager component of the OLAP Integration Server.	SC27-0787 db2dpx70	n/a
<i>OLAP Integration Server Metaoutline User's Guide</i>	Explains how to create and populate OLAP metaoutlines using the standard OLAP Metaoutline interface (not by using the Metaoutline Assistant).	SC27-0784 db2upx70	n/a
<i>OLAP Integration Server Model User's Guide</i>	Explains how to create OLAP models using the standard OLAP Model Interface (not by using the Model Assistant).	SC27-0783 db2lpx70	n/a
<i>OLAP Setup and User's Guide</i>	Provides configuration and setup information for the OLAP Starter Kit.	SC27-0702 db2ipx70	db2ip
<i>OLAP Spreadsheet Add-in User's Guide for Excel</i>	Describes how to use the Excel spreadsheet program to analyze OLAP data.	SC27-0786 db2epx70	db2ep

Table 10. DB2 Information (continued)

Name	Description	Form Number PDF File Name	HTML Directory
<i>OLAP Spreadsheet Add-in User's Guide for Lotus 1-2-3</i>	Describes how to use the Lotus 1-2-3 spreadsheet program to analyze OLAP data.	SC27-0785 db2tpx70	db2tp
<i>Replication Guide and Reference</i>	Provides planning, configuration, administration, and usage information for the IBM Replication tools supplied with DB2.	SC26-9920 db2e0x70	db2e0
<i>Spatial Extender User's Guide and Reference</i>	Provides information about installing, configuring, administering, programming, and troubleshooting the Spatial Extender. Also provides significant descriptions of spatial data concepts and provides reference information (messages and SQL) specific to the Spatial Extender.	SC27-0701 db2sbx70	db2sb
<i>SQL Getting Started</i>	Introduces SQL concepts and provides examples for many constructs and tasks.	SC09-2973 db2y0x70	db2y0
<i>SQL Reference, Volume 1 and Volume 2</i>	Describes SQL syntax, semantics, and the rules of the language. This book also includes information about release-to-release incompatibilities, product limits, and catalog views.  You can order both volumes of the <i>SQL Reference</i> in the English language in North America with the form number SBOF-8933.	Volume 1 SC09-2974 db2s1x70  Volume 2 SC09-2975 db2s2x70	db2s0
<i>System Monitor Guide and Reference</i>	Describes how to collect different kinds of information about databases and the database manager. This book explains how to use the information to understand database activity, improve performance, and determine the cause of problems.	SC09-2956 db2f0x70	db2f0
<i>Text Extender Administration and Programming</i>	Provides general information about DB2 extenders and information on the administration and configuring of the text extender and on programming using the text extenders. It includes reference information, diagnostic information (with messages) and samples.	SC26-9930 desu9x70	desu9

Table 10. DB2 Information (continued)

Name	Description	Form Number PDF File Name	HTML Directory
<i>Troubleshooting Guide</i>	Helps you determine the source of errors, recover from problems, and use diagnostic tools in consultation with DB2 Customer Service.	GC09-2850 db2p0x70	db2p0
<i>What's New</i>	Describes the new features, functions, and enhancements in DB2 Universal Database, Version 7.	SC09-2976 db2q0x70	db2q0
<b>DB2 Installation and Configuration Information</b>			
<i>DB2 Connect Enterprise Edition for OS/2 and Windows Quick Beginnings</i>	Provides planning, migration, installation, and configuration information for DB2 Connect Enterprise Edition on the OS/2 and Windows 32-bit operating systems. This book also contains installation and setup information for many supported clients.	GC09-2953 db2c6x70	db2c6
<i>DB2 Connect Enterprise Edition for UNIX Quick Beginnings</i>	Provides planning, migration, installation, configuration, and task information for DB2 Connect Enterprise Edition on UNIX-based platforms. This book also contains installation and setup information for many supported clients.	GC09-2952 db2cyx70	db2cy
<i>DB2 Connect Personal Edition Quick Beginnings</i>	Provides planning, migration, installation, configuration, and task information for DB2 Connect Personal Edition on the OS/2 and Windows 32-bit operating systems. This book also contains installation and setup information for all supported clients.	GC09-2967 db2c1x70	db2c1
<i>DB2 Connect Personal Edition Quick Beginnings for Linux</i>	Provides planning, installation, migration, and configuration information for DB2 Connect Personal Edition on all supported Linux distributions.	GC09-2962 db2c4x70	db2c4
<i>DB2 Data Links Manager Quick Beginnings</i>	Provides planning, installation, configuration, and task information for DB2 Data Links Manager for AIX and Windows 32-bit operating systems.	GC09-2966 db2z6x70	db2z6

Table 10. DB2 Information (continued)

Name	Description	Form Number PDF File Name	HTML Directory
<i>DB2 Enterprise - Extended Edition for UNIX Quick Beginnings</i>	Provides planning, installation, and configuration information for DB2 Enterprise - Extended Edition on UNIX-based platforms. This book also contains installation and setup information for many supported clients.	GC09-2964 db2v3x70	db2v3
<i>DB2 Enterprise - Extended Edition for Windows Quick Beginnings</i>	Provides planning, installation, and configuration information for DB2 Enterprise - Extended Edition for Windows 32-bit operating systems. This book also contains installation and setup information for many supported clients.	GC09-2963 db2v6x70	db2v6
<i>DB2 for OS/2 Quick Beginnings</i>	Provides planning, installation, migration, and configuration information for DB2 Universal Database on the OS/2 operating system. This book also contains installation and setup information for many supported clients.	GC09-2968 db2i2x70	db2i2
<i>DB2 for UNIX Quick Beginnings</i>	Provides planning, installation, migration, and configuration information for DB2 Universal Database on UNIX-based platforms. This book also contains installation and setup information for many supported clients.	GC09-2970 db2ixx70	db2ix
<i>DB2 for Windows Quick Beginnings</i>	Provides planning, installation, migration, and configuration information for DB2 Universal Database on Windows 32-bit operating systems. This book also contains installation and setup information for many supported clients.	GC09-2971 db2i6x70	db2i6
<i>DB2 Personal Edition Quick Beginnings</i>	Provides planning, installation, migration, and configuration information for DB2 Universal Database Personal Edition on the OS/2 and Windows 32-bit operating systems.	GC09-2969 db2i1x70	db2i1
<i>DB2 Personal Edition Quick Beginnings for Linux</i>	Provides planning, installation, migration, and configuration information for DB2 Universal Database Personal Edition on all supported Linux distributions.	GC09-2972 db2i4x70	db2i4

Table 10. DB2 Information (continued)

Name	Description	Form Number PDF File Name	HTML Directory
<i>DB2 Query Patroller Installation Guide</i>	Provides installation information about DB2 Query Patroller.	GC09-2959 db2iwx70	db2iw
<i>DB2 Warehouse Manager Installation Guide</i>	Provides installation information for warehouse agents, warehouse transformers, and the Information Catalog Manager.	GC26-9998 db2idx70	db2id
<b>Cross-Platform Sample Programs in HTML</b>			
Sample programs in HTML	Provides the sample programs in HTML format for the programming languages on all platforms supported by DB2. The sample programs are provided for informational purposes only. Not all samples are available in all programming languages. The HTML samples are only available when the DB2 Application Development Client is installed.  For more information on the programs, refer to the <i>Application Building Guide</i> .	No form number	db2hs
<b>Release Notes</b>			
<i>DB2 Connect Release Notes</i>	Provides late-breaking information that could not be included in the DB2 Connect books.	See note #2.	db2cr
<i>DB2 Installation Notes</i>	Provides late-breaking installation-specific information that could not be included in the DB2 books.	Available on product CD-ROM only.	
<i>DB2 Release Notes</i>	Provides late-breaking information about all DB2 products and features that could not be included in the DB2 books.	See note #2.	db2ir

**Notes:**

1. The character *x* in the sixth position of the file name indicates the language version of a book. For example, the file name *db2d0e70* identifies the English version of the *Administration Guide* and the file name *db2d0f70* identifies the French 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
Brazilian Portuguese	b
Bulgarian	u
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
Russian	r
Simp. Chinese	c
Slovenian	l
Spanish	z
Swedish	s
Trad. Chinese	t
Turkish	m

2. Late breaking information that could not be included in the DB2 books is available in the Release Notes in HTML format and as an ASCII file. The HTML version 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/lpp/db2_07_01` on AIX
    - `/opt/IBMDB2/V7.1` on HP-UX, PTX, Solaris, and Silicon Graphics IRIX
    - `/usr/IBMDB2/V7.1` on Linux.
  - On other platforms, see the `RELEASE.TXT` file. This file is located in the directory where the product is installed. On OS/2 platforms, you can also double-click the **IBM DB2** folder and then double-click the **Release Notes** icon.

## Printing the PDF Books

If you prefer to have printed copies of the books, you can print the PDF files found on the DB2 publications CD-ROM. Using the Adobe Acrobat Reader,

you can print either the entire book or a specific range of pages. For the file name of each book in the library, see Table 10 on page 132.

You can obtain the latest version of the Adobe Acrobat Reader from the Adobe Web site at <http://www.adobe.com>.

The PDF files are included on the DB2 publications CD-ROM with a file extension of PDF. To access the PDF files:

1. Insert the DB2 publications CD-ROM. On UNIX-based platforms, mount the DB2 publications CD-ROM. Refer to your *Quick Beginnings* book for the mounting procedures.
2. Start the Acrobat Reader.
3. Open the desired PDF file from one of the following locations:
  - On OS/2 and Windows platforms:  
*x:\doc\language* directory, where *x* represents the CD-ROM drive and *language* represent the two-character country code that represents your language (for example, EN for English).
  - On UNIX-based platforms:  
*/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.

You can also copy the PDF files from the CD-ROM to a local or network drive and read them from there.

## Ordering the Printed Books

You can order the printed DB2 books either individually or as a set (in North America only) by using a sold bill of forms (SBOF) number. To order books, contact your IBM authorized dealer or marketing representative, or phone 1-800-879-2755 in the United States or 1-800-IBM-4YOU in Canada. You can also order the books from the Publications Web page at <http://www.elink.ibm.com/pbl/pbl>.

Two sets of books are available. SBOF-8935 provides reference and usage information for the DB2 Warehouse Manager. SBOF-8931 provides reference and usage information for all other DB2 Universal Database products and features. The contents of each SBOF are listed in the following table:

Table 11. Ordering the printed books

SBOF Number	Books Included
SBOF-8931	<ul style="list-style-type: none"> <li>• Administration Guide: Planning</li> <li>• Administration Guide: Implementation</li> <li>• Administration Guide: Performance</li> <li>• Administrative API Reference</li> <li>• Application Building Guide</li> <li>• Application Development Guide</li> <li>• CLI Guide and Reference</li> <li>• Command Reference</li> <li>• Data Movement Utilities Guide and Reference</li> <li>• Data Warehouse Center Administration Guide</li> <li>• Data Warehouse Center Application Integration Guide</li> <li>• DB2 Connect User's Guide</li> <li>• Installation and Configuration Supplement</li> <li>• Image, Audio, and Video Extenders Administration and Programming</li> <li>• Message Reference, Volumes 1 and 2</li> <li>• OLAP Integration Server Administration Guide</li> <li>• OLAP Integration Server Metaoutline User's Guide</li> <li>• OLAP Integration Server Model User's Guide</li> <li>• OLAP Integration Server User's Guide</li> <li>• OLAP Setup and User's Guide</li> <li>• OLAP Spreadsheet Add-in User's Guide for Excel</li> <li>• OLAP Spreadsheet Add-in User's Guide for Lotus 1-2-3</li> <li>• Replication Guide and Reference</li> <li>• Spatial Extender Administration and Programming Guide</li> <li>• SQL Getting Started</li> <li>• SQL Reference, Volumes 1 and 2</li> <li>• System Monitor Guide and Reference</li> <li>• Text Extender Administration and Programming</li> <li>• Troubleshooting Guide</li> <li>• What's New</li> </ul>
SBOF-8935	<ul style="list-style-type: none"> <li>• Information Catalog Manager Administration Guide</li> <li>• Information Catalog Manager User's Guide</li> <li>• Information Catalog Manager Programming Guide and Reference</li> <li>• Query Patroller Administration Guide</li> <li>• Query Patroller User's Guide</li> </ul>

## DB2 Online Documentation

### Accessing Online Help

Online help is available with all DB2 components. The following table describes the various types of help.



Type of Help	Contents	How to Access...
<i>Command Help</i>	Explains the syntax of commands in the command line processor.	<p>From the command line processor in interactive mode, enter:</p> <pre>? <i>command</i></pre> <p>where <i>command</i> represents a keyword or the entire command.</p> <p>For example, ? catalog displays help for all the CATALOG commands, while ? catalog database displays help for the CATALOG DATABASE command.</p>
<i>Client Configuration Assistant Help</i>	Explains the tasks you can perform in a window or notebook. The help includes overview and prerequisite information you need to know, and it describes how to use the window or notebook controls.	From a window or notebook, click the <b>Help</b> push button or press the <b>F1</b> key.
<i>Command Center Help</i>		
<i>Control Center Help</i>		
<i>Data Warehouse Center Help</i>		
<i>Event Analyzer Help</i>		
<i>Information Catalog Manager Help</i>		
<i>Satellite Administration Center Help</i>		
<i>Script Center Help</i>	<p>From the command line processor in interactive mode, enter:</p> <pre>? <i>XXXnnnnn</i></pre> <p>where <i>XXXnnnnn</i> represents a valid message identifier.</p> <p>For example, ? SQL30081 displays help about the SQL30081 message.</p> <p>To view message help one screen at a time, enter:</p> <pre>? <i>XXXnnnnn</i>   more</pre> <p>To save message help in a file, enter:</p> <pre>? <i>XXXnnnnn</i> &gt; <i>filename.ext</i></pre> <p>where <i>filename.ext</i> represents the file where you want to save the message help.</p>	

Type of Help	Contents	How to Access...
<i>SQL Help</i>	Explains the syntax of SQL statements.	<p>From the command line processor in interactive mode, enter:</p> <pre>help <i>statement</i></pre> <p>where <i>statement</i> represents an SQL statement.</p> <p>For example, help SELECT displays help about the SELECT statement.</p> <p><b>Note:</b> SQL help is not available on UNIX-based platforms.</p>
<i>SQLSTATE Help</i>	Explains SQL states and class codes.	<p>From the command line processor in interactive mode, enter:</p> <pre>? <i>sqlstate</i> or ? <i>class code</i></pre> <p>where <i>sqlstate</i> represents a valid five-digit SQL state and <i>class code</i> represents the first two digits of the SQL state.</p> <p>For example, ? 08003 displays help for the 08003 SQL state, while ? 08 displays help for the 08 class code.</p>

## Viewing Information Online

The books included with this product are in Hypertext Markup Language (HTML) softcopy format. Softcopy format enables you to search or browse the information and provides hypertext links to related information. It also makes it easier to share the library across your site.

You can view the online books or sample programs with any browser that conforms to HTML Version 3.2 specifications.

To view online books or sample programs:

- If you are running DB2 administration tools, use the Information Center.
- From a browser, click **File** —> **Open Page**. The page you open contains descriptions of and links to DB2 information:
  - On UNIX-based platforms, open the following page:

```
INSTHOME/sql1lib/doc/%L/html/index.htm
```

where %L represents the locale name.

- On other platforms, open the following page:

```
sql1lib\doc\html\index.htm
```

The path is located on the drive where DB2 is installed.

If you have not installed the Information Center, you can open the page by double-clicking the **DB2 Information** icon. Depending on the system you are using, the icon is in the main product folder or the Windows Start menu.

### **Installing the Netscape Browser**

If you do not already have a Web browser installed, you can install Netscape from the Netscape CD-ROM found in the product boxes. For detailed instructions on how to install it, perform the following:

1. Insert the Netscape CD-ROM.
2. On UNIX-based platforms only, mount the CD-ROM. Refer to your *Quick Beginnings* book for the mounting procedures.
3. For installation instructions, refer to the `CDNAVnn.txt` file, where *nn* represents your two character language identifier. The file is located at the root directory of the CD-ROM.

### **Accessing Information with the Information Center**

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

You can open the Information Center by double-clicking the Information Center icon. Depending on the system you are using, the icon is in the Information folder in the main product folder or the Windows **Start** menu.

You can also access the Information Center by using the toolbar and the **Help** menu on the DB2 Windows platform.

The Information Center provides six types of information. Click the appropriate tab to look at the topics provided for that type.

<b>Tasks</b>	Key tasks you can perform using DB2.
<b>Reference</b>	DB2 reference information, such as keywords, commands, and APIs.
<b>Books</b>	DB2 books.
<b>Troubleshooting</b>	Categories of error messages and their recovery actions.
<b>Sample Programs</b>	Sample programs that come with the DB2 Application Development Client. If you did not install the DB2 Application Development Client, this tab is not displayed.
<b>Web</b>	DB2 information on the World Wide Web. To access this information, you must have a connection to the Web from your system.

When you select an item in one of the lists, the Information Center launches a viewer to display the information. The viewer might be the system help viewer, an editor, or a Web browser, depending on the kind of information you select.

The Information Center provides a find feature, so you can look for a specific topic without browsing the lists.

For a full text search, follow the hypertext link in the Information Center to the **Search DB2 Online Information** search form.

The HTML search server is usually started automatically. If a search in the HTML information does not work, you may have to start the search server using one of the following methods:

#### **On Windows**

Click **Start** and select **Programs** → **IBM DB2** → **Information** → **Start HTML Search Server**.

#### **On OS/2**

Double-click the **DB2 for OS/2** folder, and then double-click the **Start HTML Search Server** icon.

Refer to the release notes if you experience any other problems when searching the HTML information.

**Note:** The Search function is not available in the Linux, PTX, and Silicon Graphics IRIX environments.

## **Using DB2 Wizards**

Wizards help you complete specific administration tasks by taking you through each task one step at a time. Wizards are available through the Control Center and the Client Configuration Assistant. The following table lists the wizards and describes their purpose.

**Note:** The Create Database, Create Index, Configure Multisite Update, and Performance Configuration wizards are available for the partitioned database environment.

<b>Wizard</b>	<b>Helps You to...</b>	<b>How to Access...</b>
<i>Add Database</i>	Catalog a database on a client workstation.	From the Client Configuration Assistant, click <b>Add</b> .
<i>Backup Database</i>	Determine, create, and schedule a backup plan.	From the Control Center, right-click the database you want to back up and select <b>Backup</b> → <b>Database Using Wizard</b> .

<b>Wizard</b>	<b>Helps You to...</b>	<b>How to Access...</b>
<i>Configure Multisite Update</i>	Configure a multisite update, a distributed transaction, or a two-phase commit.	From the Control Center, right-click the <b>Databases</b> folder and select <b>Multisite Update</b> .
<i>Create Database</i>	Create a database, and perform some basic configuration tasks.	From the Control Center, right-click the <b>Databases</b> folder and select <b>Create</b> —> <b>Database Using Wizard</b> .
<i>Create Table</i>	Select basic data types, and create a primary key for the table.	From the Control Center, right-click the <b>Tables</b> icon and select <b>Create</b> —> <b>Table Using Wizard</b> .
<i>Create Table Space</i>	Create a new table space.	From the Control Center, right-click the <b>Table Spaces</b> icon and select <b>Create</b> —> <b>Table Space Using Wizard</b> .
<i>Create Index</i>	Advise which indexes to create and drop for all your queries.	From the Control Center, right-click the <b>Index</b> icon and select <b>Create</b> —> <b>Index Using Wizard</b> .
<i>Performance Configuration</i>	Tune the performance of a database by updating configuration parameters to match your business requirements.	From the Control Center, right-click the database you want to tune and select <b>Configure Performance Using Wizard</b> .  For the partitioned database environment, from the Database Partitions view, right-click the first database partition you want to tune and select <b>Configure Performance Using Wizard</b> .
<i>Restore Database</i>	Recover a database after a failure. It helps you understand which backup to use, and which logs to replay.	From the Control Center, right-click the database you want to restore and select <b>Restore</b> —> <b>Database Using Wizard</b> .

## Setting Up a Document Server

By default, the DB2 information is installed on your local system. This means that each person who needs access to the DB2 information must install the same files. To have the DB2 information stored in a single location, perform the following steps:

1. Copy all files and subdirectories from `\sql11ib\doc\html` on your local system to a Web server. Each book has its own subdirectory that contains all the necessary HTML and GIF files that make up the book. Ensure that the directory structure remains the same.

2. Configure the Web server to look for the files in the new location. For information, refer to the NetQuestion Appendix in the *Installation and Configuration Supplement*.
3. If you are using the Java version of the Information Center, you can specify a base URL for all HTML files. You should use the URL for the list of books.
4. When you are able to view the book files, you can bookmark commonly viewed topics. You will probably want to bookmark the following pages:
  - List of books
  - Tables of contents of frequently used books
  - Frequently referenced articles, such as the ALTER TABLE topic
  - The Search form

For information about how you can serve the DB2 Universal Database online documentation files from a central machine, refer to the NetQuestion Appendix in the *Installation and Configuration Supplement*.

### Searching Information Online

To find information in the HTML files, use one of the following methods:

- Click **Search** in the top frame. Use the search form to find a specific topic. This function is not available in the Linux, PTX, or Silicon Graphics IRIX environments.
- Click **Index** in the top frame. Use the index to find a specific topic in the book.
- Display the table of contents or index of the help or the HTML book, and then use the find function of the Web browser to find a specific topic in the book.
- Use the bookmark function of the Web browser to quickly return to a specific topic.
- Use the search function of the Information Center to find specific topics. See “Accessing Information with the Information Center” on page 145 for details.

---

## Appendix F. 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 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 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 which has been exchanged, should contact:

IBM Canada Limited  
Office of the Lab Director  
1150 Eglinton Ave. East  
North York, Ontario  
M3C 1H7  
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 information 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 illustrates 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, which may be denoted by an asterisk(\*), are trademarks of International Business Machines Corporation in the United States, other countries, or both.

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

The following terms are trademarks or registered trademarks of other companies:

Microsoft, Windows, and Windows NT are trademarks or registered trademarks of Microsoft Corporation.

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

Tivoli and NetView are trademarks of Tivoli Systems Inc. in the United States, other countries, or both.

UNIX is a registered trademark in the United States, other countries or both and is licensed exclusively through X/Open Company Limited.

Other company, product, or service names, which may be denoted by a double asterisk(\*\*) may be trademarks or service marks of others.



---

# Index

## A

- access profiles
  - adding databases 63
  - client 68
  - creating 68
  - server 68
  - using 68
- accessing data
  - using DB2 Connect 116
  - using Net.Data or JDBC 119
- accessing multiple servers 114
- add database wizard 146, 147
- adding databases
  - manually 66
  - using access profiles 63
  - using Discovery 64
- Administration Server
  - overview 128
- Alert Center 125
- ALTER TABLESPACE 44
- APPC
  - IBM Communications Server for NT 17
  - Microsoft SNA Server 17
  - software requirements 17, 18, 19
  - supported platforms 17
- application development
  - using Net.Data or JDBC 119
- application programs
  - coordinator node 4
  - migrating to DB2 Enterprise - Extended Edition 4

## B

- backup database wizard 146
- books 131, 141

## C

- capacity
  - disk storage 7
- catalog node name
  - naming rules 101
- catalog nodes 8
- client profiles
  - creating 69
  - definition 69
  - importing 70
  - using 69
- clients
  - installing 53

- clients (*continued*)
  - operating systems supported 113
- columns
  - long-field, partitioning key considerations 9
- Command Center
  - entering DB2 commands 90
  - entering SQL statements 90
  - overview 122
- commands
  - db2 list applications 21
  - db2 list tablespaces 44
  - db2 terminate 21
  - db2cc 78
  - db2ckmig 22
  - db2jstrt 76
  - db2sampl 78
  - db2start 46
  - db2stop 21
  - db2uiddl 46
  - sniffle 82
- communication protocols
  - APPC 17, 18, 19
  - IPX/SPX 17, 18, 19
  - Named Pipes 17, 18, 19
  - NetBIOS 17, 18, 19
  - TCP/IP 17, 18, 19
- communications
  - Control Center 126
- concurrency
  - improving with logical nodes 9
- configuration parameters
  - setting DB2 94
  - SYSADM\_GROUP 94
- configure multisite update wizard 146
- configuring
  - DB2 clients
    - using the Client Configuration Assistant (CCA) 61
- configuring client communications
  - setting configuration parameters 47
- Control Center
  - administering DB2 Connect Enterprise Edition 84
  - administering DB2 for OS/390 84

- Control Center (*continued*)
  - as a Java applet 73
  - as a Java application 73
  - components 122
  - configuring to work with a web server 80
  - customizing db2cc.htm 79
  - functional considerations 80
  - JDBC Applet Server 76
  - machine configurations 74
  - overview 122
  - remotely administering
    - clusters 82
  - running as an applet 78
  - running as an application 78
  - setting up to run as an applet 76
  - supported browsers 75
  - supported Java Runtime Environments (JRE) 75
  - troubleshooting information 83
  - UNIX installation tips 80
- coordinator nodes
  - description 4
- cost-based query optimizers
  - description 4
- create database wizard 147
- create table space wizard 147
- create table wizard 147
- creating profiles
  - clients 69
  - server 69
- creating the sample database
  - cataloging a database 44

## D

- Data Links Manager
  - overview 113
- Data Warehousing
  - sample database 39
- database access
  - two-phase commit 115
- database administration tools
  - Control Center 122
  - overview 122
- database alias
  - naming rules 101
- database manager
  - definition 3

- database objects
    - naming rules 102
  - database partition server
    - definition 3
  - databases
    - authentication type 43
    - definition 3
    - migrating 43
    - naming rules 101
    - partition 3
  - DB2 Application Development Client
    - overview 129
  - DB2 clients
    - accessing databases 114
    - changing privileges 94
    - installing 53
    - licensing 53
    - overview 114, 127
    - platforms supported 113
    - pre-version 7 53
    - software requirements 18
    - supported platforms 53
    - Windows 32-bit operating systems 57
  - DB2 Connect
    - overview 110, 116
  - DB2 Enterprise - Extended Edition
    - configuration planning 15
    - hardware configuration 6
  - DB2 Enterprise Edition
    - memory requirements 15
  - DB2 Everywhere
    - overview 107
  - DB2 library
    - books 131
    - Information Center 145
    - language identifier for
      - books 139
    - late-breaking information 140
    - online help 142
    - ordering printed books 141
    - printing PDF books 140
    - searching online
      - information 148
    - setting up document server 147
    - structure of 131
    - viewing online information 144
    - wizards 146
  - DB2 Security Server
    - starting on Windows NT or
      - Windows 2000 77
  - DB2 Universal Database
    - Control Center 122
    - DB2 Performance Monitor 126
    - DB2 Snapshot Monitor 119
  - DB2 Universal Database (*continued*)
    - overview 122
    - platforms supported 122
    - software requirements 17
    - Visual Explain 126
  - db2classes.exe 78
  - db2classes.tar.Z 78
  - DB2SYSTEM
    - naming rules 105
  - db2uiddl command 46
  - db2unins command
    - uninstalling DB2 97
  - Developer Editions
    - overviews 109
  - developing applications
    - using Net.Data or JDBC 119
  - Discovery
    - adding databases 64
  - disk requirements
    - client 16
    - server 16
  - Distributed Computing Environment
    - software requirements 17
  - DMS tablespaces
    - migrating Version 5
      - databases 44
- E**
- Enterprise - Extended Edition
    - capacity 7
    - overview 109
    - security for an instance 10
    - shared nothing definition 3
    - storage 7
  - Enterprise Edition
    - overview 109
  - environment variables 10
  - error handling
    - migration 23
  - Export function 68, 69
- F**
- fixed disks
    - hardware requirements 16
- G**
- gigabyte (GB) 7
- H**
- hardware requirements
    - fixed disk 16
  - host databases
    - overview 110
  - HTML
    - sample programs 139
- I**
- IBMCATGROUP nodegroup 8
  - IBMDEFAULTGROUP nodegroup 8
  - IBMTEMPGROUP nodegroup 7
  - Import function 68
  - importing profiles
    - client 70
  - index wizard 147
  - Information Center 145
  - installing
    - CID using SystemView LAN 35, 59
    - client 15, 16
    - DB2 Application Development Clients 53
      - Windows 32-bit operating systems 57
    - DB2 clients 53
    - errors 35, 59
    - log 35, 59
    - Netscape browser 145
    - server 15, 16
    - verifying the installation 39
  - instance owner
    - instance owner 11
  - instance-owning machines
    - definition 11
  - instances
    - definition 10
    - multiple on one processor 11
    - naming restrictions 103
  - Intelligent Miner
    - overview 112
  - inter-operator parallelism 5
- J**
- Java Runtime Environment (JRE)
    - defined 73
  - Java Support 119
  - Java Virtual Machine (JVM) 73
  - JDBC Applet Server 76
  - Journal 125
  - JRE
    - supported levels for Control Center 75
- L**
- language identifier
    - books 139
  - late-breaking information 140
  - logical nodes
    - improving concurrency 9
    - purpose 9

- long-field columns
  - partitioning key considerations 9

## M

- managing connections
  - overview 127
  - using the Client Configuration Assistant 127
- managing databases using the Control Center 122
- managing server communications
  - overview 126
- manually adding databases 66
- memory requirements
  - client 15
  - estimating 15
  - recommended 15
  - server 15
- Microsoft SNA Server
  - version required 19
- migrating
  - applications 4
  - databases 22, 43
  - pre-installation tasks 21
  - rebinding packages 47
  - updating database and database manager configuration 47
  - updating statistics 47
  - Version 5 databases with DMS tablespaces 44
- migration
  - post-installation tasks 43

## N

- naming rules
  - database alias 101
  - database objects 102
  - databases 101
  - general 101
  - groups 103
  - instance names 103
  - password 105
  - userids 103
  - username 103
- Net.Data
  - connecting to the Internet 119
  - overview 119
- Net Search Extender
  - overview 112
- NetBIOS
  - code set 99
  - codepage support 99
  - determining codepage 101

- Netscape browser
  - installing 145
- nodegroups
  - IBMCAATGROUP 8
  - IBMDEFAULTGROUP 8
  - multipartition nodegroup 7
- nodes
  - definition 3
  - in the same instance 10

## O

- OLAP Server
  - overview 112
- OLAP Starter Kit
  - sample database 39
- online help 142
- online information
  - searching 148
  - viewing 144

## P

- parameters
  - SYSADM\_GROUP 94
- partitioning key
  - definition 9
- partitioning map
  - overview 9
- partitions
  - definition 3
- passwords
  - naming rules 105
- PDF 140
- performance configuration
  - wizard 147
- Performance Monitor
  - using 126
- Personal Edition
  - overview 108
- planning
  - DB2 configuration 15
  - DB2 Connect configuration 15
- pools
  - agent 9
- printing PDF books 140
- privileges
  - required 94
- product
  - descriptions 107
  - overview 107
- profiles
  - client 68, 69
  - export 68
  - server 68

## Q

- query
  - cost-based optimizer 4

## R

- Relational Connect
  - overview 111
- release notes 140
- restore wizard 147
- restrictions
  - instance name 103

## S

- sample programs
  - cross-platform 139
  - HTML 139
- Satellite Edition
  - overview 108
- scaling the database system
  - definition 9
- Script Center 124
- searching
  - online information 146, 148
- security
  - for an instance 10
- server profiles
  - creating 69
  - definition 68
- setting up document server 147
- shared-nothing configuration
  - definition 3
  - storage 7
- SmartGuides
  - wizards 146
- software requirements
  - communication protocols 17
  - DB2 Application Development Clients 17, 18, 19
  - DB2 clients 17, 18, 19
  - DB2 Connect 17
  - DB2 Universal Database 17
  - Net.Data 18, 19
- Spatial Extender
  - overview 112
- SQL
  - viewing using Visual Explain 126
- Stored Procedure Builder 125
- symmetric multiprocessor
  - architecture 9
- SYSADM
  - controlling 94
- SYSADM\_GROUP parameter 94
- system configuration
  - with DB2 Connect 117

system configuration (*continued*)  
with DB2 Universal  
Database 114

## T

tables

- maximum size 7
- multinode nodegroups 9
- single-node nodegroups 9

TCP/IP

- configuring on OS/2 81
- enabling localhost on OS/2 81
- enabling loopback on OS/2 81
- verifying on OS/2 82

Tivoli Enterprise

- overview 113

Tools Setting 125

transparent parallelism 4

## U

uninstalling DB2 97

username

- naming rules 103

## V

verifying the installation

- server 39

viewing

- online information 144

Visual Explain

- overview 126

## W

Warehouse Manager

- overview 111

Windows 2000

- starting the security server 77

Windows NT

- starting the security server 77

wizards

- add database 146, 147
- backup database 146
- completing tasks 146
- configure multisite update 146
- create database 147
- create table 147
- create table space 147
- index 147
- performance configuration 147
- restore database 147

Workgroup Edition

- overview 109

working with DB2 data 113

workstation name (nname)

- naming rules 104



---

## Contacting IBM

If you have a technical problem, please review and carry out the actions suggested by the *Troubleshooting Guide* before contacting DB2 Customer Support. This guide suggests information that you can gather to help DB2 Customer Support to serve you better.

For information or to order any of the DB2 Universal Database products contact an IBM representative at a local branch office or contact any authorized IBM software remarketer.

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

- 1-800-237-5511 for customer support
- 1-888-426-4343 to learn about available service options

---

## Product Information

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) or 1-800-3IBM-OS2 (1-800-342-6672) to order products or get general information.
- 1-800-879-2755 to order publications.

**<http://www.ibm.com/software/data/>**

The DB2 World Wide Web pages provide current DB2 information about news, product descriptions, education schedules, and more.

**<http://www.ibm.com/software/data/db2/library/>**

The DB2 Product and Service Technical Library provides access to frequently asked questions, fixes, books, and up-to-date DB2 technical information.

**Note:** This information may be in English only.

**<http://www.elink.ibm.com/pbl/pbl/>**

The International Publications ordering Web site provides information on how to order books.

**<http://www.ibm.com/education/certify/>**

The Professional Certification Program from the IBM Web site provides certification test information for a variety of IBM products, including DB2.

**ftp.software.ibm.com**

Log on as anonymous. In the directory /ps/products/db2, you can find demos, fixes, information, and tools relating to DB2 and many other products.

**comp.databases.ibm-db2, bit.listserv.db2-l**

These Internet newsgroups are available for users to discuss their experiences with DB2 products.

**On CompuServe: GO IBMDB2**

Enter this command to access the IBM DB2 Family forums. All DB2 products are supported through these forums.

For information on how to contact IBM outside of the United States, refer to Appendix A of the *IBM Software Support Handbook*. To access this document, go to the following Web page: <http://www.ibm.com/support/>, and then select the IBM Software Support Handbook link near the bottom of the page.

**Note:** In some countries, IBM-authorized dealers should contact their dealer support structure instead of the IBM Support Center.





Part Number: CT7YQNA



Printed in the United States of America  
on recycled paper containing 10%  
recovered post-consumer fiber.

GC09-2963-00



CT7YQNA

