

IBM DB2 Universal Database
Enterprise - Extended Edition for Windows NT**



Quick Beginnings

Version 6

IBM DB2 Universal Database
Enterprise - Extended Edition for Windows NT**



Quick Beginnings

Version 6

Before using this information and the product it supports, be sure to read the general information under "Appendix G. Notices" on page 121.

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, 1999. All rights reserved.**

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

Contents

Welcome to DB2 Universal Database!	v
Conventions	v

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

Chapter 1. Introduction to DB2 Enterprise - Extended Edition	3
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 Communication Manager	11
High Availability	12
Virtual Interface Architecture	12

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

Chapter 2. Planning for Installation	15
Memory Requirements	15
DB2 Universal Database Requirements	15
DB2 Client Requirements	16
Disk Requirements	16
Estimating Fixed Disk Requirements	16
Software Requirements	17
Server Product Requirements	17
Client Product Requirements	18
Possible Client-to-Server Connectivity Scenarios	20

Chapter 3. Installing DB2 Universal Database Enterprise - Extended Edition on Windows NT	21
Before You Begin	21
Performing the Installation	24
Verifying the Installation	27
Software Registration	28

Part 3. Installing and Configuring DB2 Clients 31

Chapter 4. Installing DB2 Clients	33
--	-----------

Chapter 5. Installing DB2 Clients on Windows 32-Bit Operating Systems	35
Before You Begin	35
Installation Steps	35

Chapter 6. Configuring Client-to-Server Communications Using the Client Configuration Assistant	39
Configuration Steps	40
Adding a Database Using a Profile	40
Adding a Database Using Discovery	42
Adding a Database Manually	44
Creating and Using Profiles	46
Server Profiles	46
Client Profiles	47

Part 4. Appendixes 51

Appendix A. Basic Task Knowledge	53
Starting First Steps	53
Starting the Software Registration Tool	53
Starting the Client Configuration Assistant	53
Starting the Control Center	54
Entering Commands Using the Command Center	54
Entering Commands Using the Command Line Processor	56
Command Line Mode	56
Interactive Input Mode	57
Working with the System Administrative Group	57

Appendix B. Migrating from Previous Versions and Releases	59
Migrating from Previous Versions of DB2	59
Pre-Installation Migration Steps	60
Step 1. Back Up Databases	60

Step 2. Verify that Databases Can Be Migrated	61	Managing Instances and Database Objects Using the Control Center	96
Post-Installation Steps	67	Managing Communications on the Server	99
Migrating Databases	68	Monitoring Databases Using DB2 Performance Monitor	99
Optional Post Migration Actions	70	Viewing SQL Access Plans Using Visual Explain	100
Appendix C. National Language Support (NLS)	73	Managing Connections to Databases Using the Client Configuration Assistant	100
Code Page and Language Support	73	Understanding the Administration Server	101
Appendix D. Naming Rules	75	Developing Applications Using the DB2 Software Developer's Kit	101
General Naming Rules	75	Running Your Own Applications	102
Database, Database Alias, and Catalog Node Name Rules	75	Appendix F. How the DB2 Library Is Structured.	105
Object Name Rules	76	Completing Tasks with SmartGuides . . .	105
Username, User ID, Group Name, and Instance Name Rules	77	Accessing Online Help	106
Workstation Name (nname) Rules	77	DB2 Information – Hardcopy and Online	108
DB2SYSTEM Naming Rules	79	Viewing Online Information	115
Password Rules	79	Accessing Information with the Information Center	116
Appendix E. About the DB2 Family of Products	81	Setting Up a Document Server	117
Working with DB2 Data	85	Searching Online Information	118
Accessing DB2 Data from Remote Clients	86	Printing the PostScript Books.	118
Accessing Multiple DB2 Servers. . . .	87	Ordering the Printed Books	119
Accessing Host or AS/400 DB2 Data from the Desktop Using DB2 Connect Enterprise Edition	88	Appendix G. Notices	121
Accessing DB2 Data from the Web Using Java	92	Trademarks	122
Accessing DB2 Data from the Web Using Net.Data	93	Trademarks of Other Companies	122
Accessing DB2 Data from Host and AS/400 Client Machines	95	Index	125
Administering Instances and Databases with the DB2 Administration Tools. . . .	95	Contacting IBM	129

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 set up 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 F. How the DB2 Library Is Structured” on page 105.



- 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, or Windows NT.
- The term *Windows 9x* refers to Windows 95 or Windows 98.
- The term *DB2 client* refers to a DB2 Run-Time Client or a DB2 Administration Client.
- The term *DB2 Universal Database* refers to DB2 Universal Database on OS/2, UNIX, and Windows 32-bit operating systems, unless otherwise stated.
- The term *database partition server* is also referred to as a *node*.

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. It achieves this function by using system resources, including CPU, memory, disk, 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 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 one of the database partition servers (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.

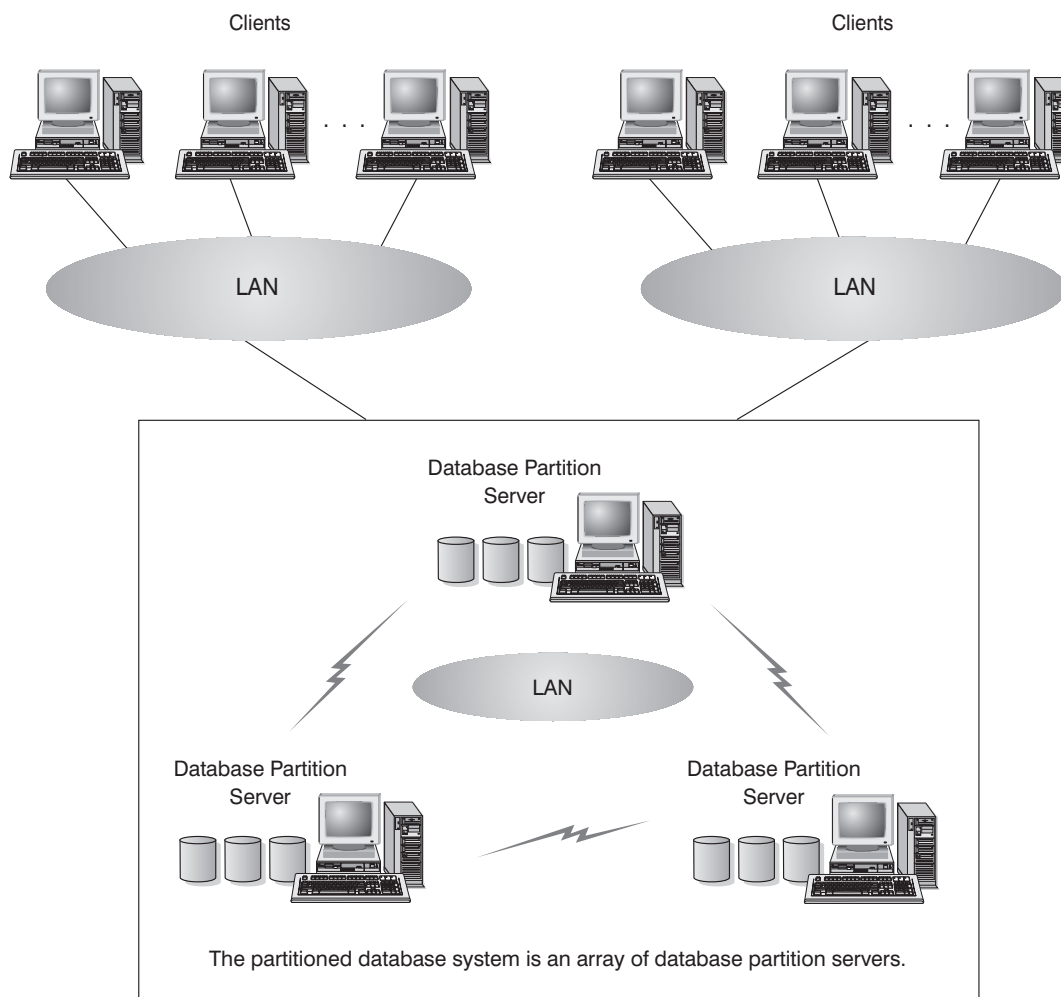


Figure 1. DB2 Enterprise - Extended Edition Hardware Configuration

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 64 GB (gigabytes) per database partition, where giga is 10^9 . For example, in a database that has 128 partitions, the maximum size of one table is 8 TB (terabytes, where tera is 10^{12}).

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 tables in the default nodegroups IBMDEFAULTGROUP and IBMCATGROUP.

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 shows an example of a database in which there are three nodegroups. Nodegroup 1 is a multipartition nodegroup made of four database partitions, and nodegroups 2 and 3 are both single-partition nodegroups.

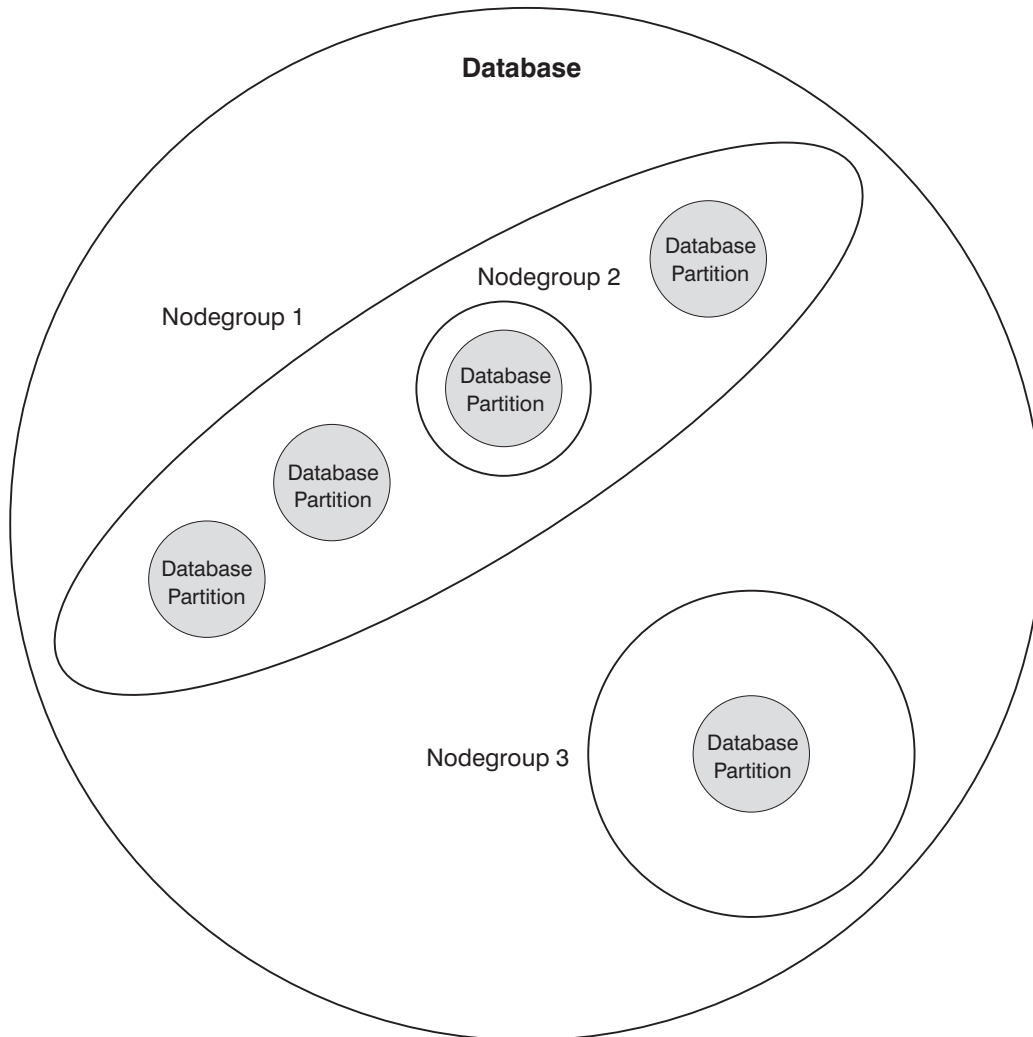


Figure 2. Nodegroups in a Database

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 several database partition servers running on the same machine. This means that the configuration can contain more nodes than machines. In these cases, the machine is said to be running *multiple logical nodes* if they participate in the *same* instance.

With multiple logical node support, you can choose from three types of configurations:

- A standard configuration, where each machine has only one database partition server
- A multiple logical node configuration, where a machine has more than one database partition server
- A configuration where several logical nodes run on each of several machines.

Configurations that use multiple logical nodes are useful when the system runs queries on a machine that has symmetric multiprocessor (SMP) architecture. The ability to configure multiple logical nodes on a machine is also useful if a machine fails. If a machine fails (causing the database partition server or servers on it to fail), you can restart the database partition server (or servers) to run on another machine, using the **db2start nodenum** command or the Start Node option in the Control Center. This ensures that user data remains available. For more information, refer to the *Command Reference* or the Control Center's online help.

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.

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 on page 11, 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 appear to 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.

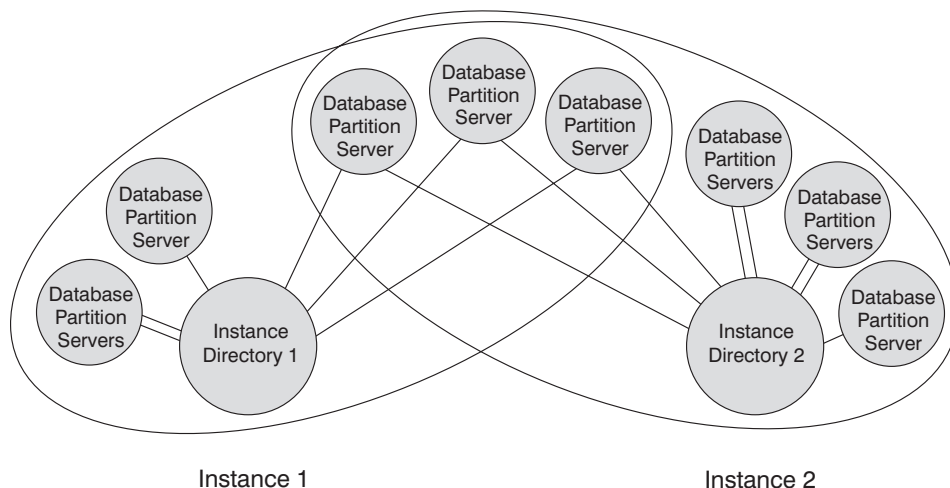


Figure 3. Two Instances

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 Communication 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 `fcnum_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 NT, you implement failover support using Microsoft Cluster Service (MSCS). To use MSCS, you require Windows NT Server 4.0, Enterprise Edition, with the MSCS component installed. For more information, refer to the *Administration Guide*.

Virtual Interface Architecture

On Windows NT, 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



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 NT” on page 21.

For information on the DB2 family of products, see “Appendix E. About the DB2 Family of Products” on page 81.

There are many components that you might want to use in your environment. Use the product and planning information in this section to ensure that your system meets the prerequisites and to decide which components you want to install.

Before you begin your DB2 product installation, you should determine the requirements for the system that you are planning to install and configure.

Memory Requirements

This section shows the *suggested* amount of memory that is required to run a DB2 product. The memory requirements listed here are estimates; the actual amounts required depend on the functions you are using. Use this information to plan for systems with a large number of concurrent clients or databases.

For information about memory requirements during the day-to-day operations of your databases, refer to the *Administration Guide*.

DB2 Universal Database Requirements

Complete the provided worksheet in Table 1 and calculate the approximate values for memory that your system will require.

Table 1. Memory Requirements for DB2 Universal Database

Number of Clients Connecting to a Server	
___ 5 Concurrent Connections	32 MB
___ 10 Concurrent Connections	40 MB
___ 25 Concurrent Connections	48 MB
___ 50 Concurrent Connections	88 MB
___ DB2 Administration Tools	30 MB
Total Memory Requirements	___ MB

DB2 Client Requirements

The amount of memory you require to run a DB2 Run-Time Client is approximately 16 MB. If you are planning to run a DB2 Administration Client, your client workstation should have approximately 32 MB of available memory.

Disk Requirements

This section shows the *minimum* amount of disk space that is required to install your DB2 product and components. It does not include the disk requirements necessary for the operating system, application development tools, and communications products. Consult each product's documentation for these values. Estimates for disk space requirements are listed here; the actual amounts required depend on the functions you are using.

For information about space requirements for data, refer to the *Administration Guide*.

Estimating Fixed Disk Requirements

To estimate the disk requirements for a particular configuration, add the recommended minimum disk sizes for the products and components that you want to install. Include an allowance for your application data.

Server Components

Use Table 2 to estimate the amount of disk space you need to install DB2 and associated components on your operating system.

Table 2. Estimating Disk Requirements

	Recommended Minimum Disk (MB)
DB2 Universal Database Enterprise - Extended Edition for Windows NT	
DB2 Universal Database	34 MB
DB2 GUI Tools	68 MB
DB2 Connect Support (this is not available with DB2 Workgroup Edition)	1 MB
Online documentation in HTML format (English)	31 MB
Far-East Code Page Conversion Support	5 MB
Total Disk Space Required	__ MB

Client Components

Use Table 3 to estimate the amount of disk space you need on each of your client workstations.



If you are planning to use a DB2 client that is not listed in this table, but is supported by DB2, refer to the *Installation and Configuration Supplement* for more information.



If you are installing the Control Center or the online documentation on a Windows NT FAT or a Windows 9x FAT16 disk partition, you may require additional amounts of disk space.

Table 3. Disk Requirements for Client Components

Recommended Minimum Disk (MB)	
Windows 32-bit Operating Systems	
DB2 Run-Time Client (Base)	11 MB
- Client Configuration Assistant	6 MB
or	
DB2 Administration Client (Base)	11 MB
- DB2 Administration Tools	68 MB
- Client Configuration Assistant	6 MB
Total Disk Space Required	__ MB

Software Requirements

This section outlines the software required to run DB2 products.

Server Product Requirements

Table 4 on page 18 lists the operating system and communications software required for DB2 Universal Database.

Table 4. Software Requirements

Product	Hardware/Software Requirements	Communications
Windows 32-bit Operating Systems		
• DB2 Universal Database	• Windows NT Version 4.0 with Service Pack 3, or later	<p>APPC, TCP/IP, and MPTN (APPC over TCP/IP)</p> <ul style="list-style-type: none"> • The Windows NT base operating system provides NetBIOS, IPX/SPX, Named Pipes, and TCP/IP connectivity. <p>For SNA (APPC) connectivity, one of the following communication products is required:</p> <ul style="list-style-type: none"> • IBM Communications Server for Windows NT Version 5.01 or later • IBM Personal Communications for Windows NT Version 4.1 or later • Microsoft SNA Server Version 4 Service Pack 2 or later <p>Notes:</p> <ol style="list-style-type: none"> 1. If SNA two-phase commit is necessary, then IBM Communications Server for Windows NT Version 5.01 or later is required. 2. If you plan to use DCE (Distributed Computing Environment) with Version 6 of DB2 Universal Database, you will need: <ul style="list-style-type: none"> • A DCE product that is at OSF DCE level 1.1 with IBM DCE for Windows NT Version 2.0. • To ensure that if you are connecting to DB2 for OS/390 V5.1 database that it is enabled for DCE support using OS/390 DCE Base Services Version 3. <p>In addition, DB2 also supports Gradient PC-DCE for Windows 32-bit operating systems Version 2.0 Runtime Media Kit.</p> <ol style="list-style-type: none"> 3. If you plan to use LDAP (Lightweight Directory Access Protocol), you require the IBM eNetwork LDAP Directory Client Version 3.1. For more information, refer to the <i>Administration Guide</i>. 4. If you plan to use the ADSTAR Distributed Storage Manager (ADSM) facilities for backup and restore of your databases, you require the ADSM Client Version 3 or later. 5. If you have the IBM Antivirus program installed on your operating system, it must be Version 3.0 or later. 6. If you plan to use the Simple Network Management Protocol (SNMP) subagent, you require DPI 2.0 provided by IBM Netfinity Agent. 7. If you plan to set up Failover support, then you require Windows NT Enterprise Edition V4.0 with the Microsoft Cluster Server support.

Client Product Requirements

Table 5 on page 19 lists the software requirements needed for a DB2 Administration Client, DB2 Run-Time Client, or a DB2 Software Developer's Kit.

Table 5. Software Requirements for Clients

Component	Hardware/Software Requirements	Communications
<ul style="list-style-type: none"> • DB2 Run-Time Client for Windows 9x • DB2 Administration Client for Windows 9x • DB2 Software Developer's Kit for Windows 9x 	<ul style="list-style-type: none"> • Windows 95 4.00.950 or later • Windows 98 	<p>IPX/SPX, Named Pipes, NetBIOS, or TCP/IP</p> <ul style="list-style-type: none"> • The Windows 9x base operating system provides NetBIOS, IPX/SPX, TCP/IP, and Named Pipes connectivity. Note: IPX/SPX connectivity is only supported to Windows NT servers. • If you plan to use LDAP (Lightweight Directory Access Protocol), you require the IBM eNetwork LDAP Directory Client Version 3.1. For more information, refer to the <i>Administration Guide</i>.
<ul style="list-style-type: none"> • DB2 Run-Time Client for Windows NT • DB2 Administration Client for Windows NT • DB2 Software Developer's Kit for Windows NT 	<ul style="list-style-type: none"> • Windows NT Version 4.0 with Service Pack 3 or later • Windows Terminal Server (can only run the DB2 Run-Time Client) 	<p>APPC, IPX/SPX, Named Pipes, NetBIOS, or TCP/IP</p> <ul style="list-style-type: none"> • The Windows NT base operating system provides NetBIOS, IPX/SPX, TCP/IP, and Named Pipes connectivity. • For APPC connectivity, you require one of the following products: <ul style="list-style-type: none"> – IBM eNetwork Communications Server for Windows NT V5.01 or later. – IBM eNetwork Personal Communications for Windows NT V4.2 or later. – Microsoft SNA Server Version 4 Service Pack 2 or later – Wall Data Rumba • If you plan to use DCE (Distributed Computing Environment) with Version 6 of DB2 Universal Database, you will need to ensure that if you are connecting to DB2 for OS/390 V5.1 database that it is enabled for DCE support using OS/390 DCE Base Services Version 3. • If you plan to use LDAP (Lightweight Directory Access Protocol), you require the IBM eNetwork LDAP Directory Client Version 3.1. For more information, refer to the <i>Administration Guide</i>. • If you plan to use the ADSTAR Distributed Storage Manager (ADSM) facilities for backup and restore of your databases, you require the ADSM Client Version 3 or later. • If you have the IBM Antivirus program installed on your operating system, it must be Version 3.0 or later.

Possible Client-to-Server Connectivity Scenarios

The following table shows the communication protocols that can be used when connecting a specific LAN, host or AS/400 DB2 client to a specific DB2 server or DB2 Connect server.



DB2 Workgroup, DB2 Enterprise, and DB2 Enterprise - Extended Editions can service requests from host or AS/400 clients (DRDA ARs).

Table 6. Possible Client-to-Server Connectivity Scenarios

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

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



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 6.0. Go to “Appendix B. Migrating from Previous Versions and Releases” on page 59 for more information.

This section describes how to create a partitioned database system. You will first install the instance-owning database partition server and create (or migrate) a partitioned database instance. On another machine, you will install a database partition server that will participate in this instance.

If you want to install a DB2 Run-Time Client or a DB2 Administration Client, go to “Chapter 4. Installing DB2 Clients” on page 33.

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 NT 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`: is the drive where you installed Windows NT).

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, 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`: is the drive where you installed Windows NT).

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 paulz.torolab.ibm.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.

- ___ 2. A domain user account that belongs to the *Local Administrators* group on the machine where you are going to perform the installation. The user account you specify must be locally defined 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 user rights, refer to the Windows NT online help.



You can perform the installation without the *Act as part of the operating system* advanced user right, however, the setup program will be unable to validate the account that you specify for the Administration Server. We recommend that any user account used to install this product have this advanced user right.

- ___ 3. 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 set up a user account with the username *db2admin* and password *db2admin*. You can accept these values, or provide your own. 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 D. Naming Rules" on page 75.



If you use the default user account *db2admin*, and do not change the default password on the setup window, you should change this password immediately following the installation. This password is the default used for any installation, and therefore is well known. Using this password could pose a security risk to your network.

You must also change the password for the *DB2-DB2DAS00* service to match the new password that you specify for the *db2admin* username. To change the password that is associated with the Administration Server, entering the following command:

```
db2admin setid username password
```

- ___ 4. To verify that DB2 installed correctly, you will need to have a user account that belongs to the System Administrative (SYSADM) group, is 8 characters or less, and complies with all of DB2's naming rules.

By default, any user 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 57. For more information on valid DB2 usernames, see “Appendix D. Naming Rules” on page 75.

Performing the Installation

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

- Step 1. Log on to the system with the domain user account that you created 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 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 auto-start, see the tip that follows.



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

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

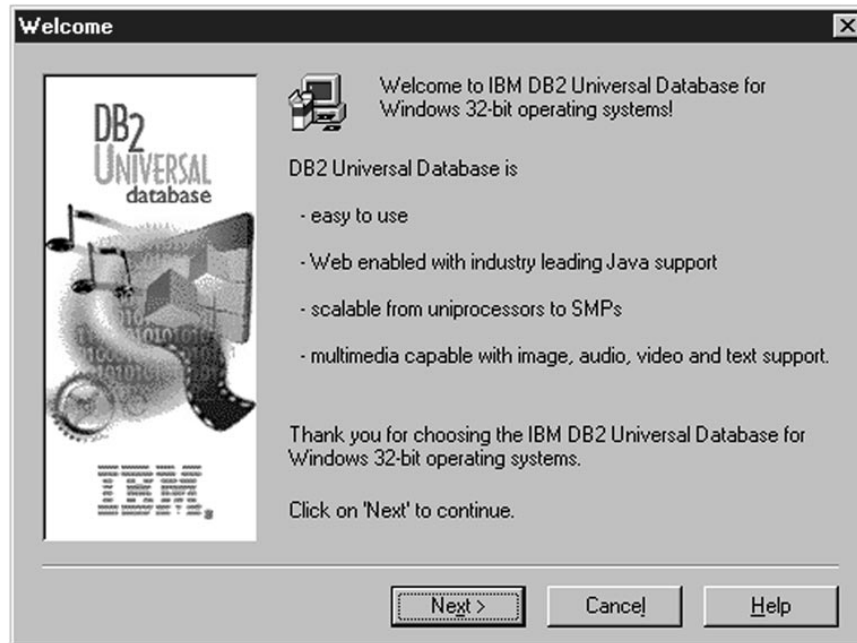
`x:\setup /i language`

where:

- *x:* represents your CD-ROM drive
- *language* represents the country code for your language (for example, EN for English). Table 8 on page 73 lists the code for each available language.

- c. Click on **OK**.

Step 4. The Welcome window opens.



Step 5. Respond to the setup program's prompts. Online help is available to guide you through the remaining steps. Invoke online help by clicking on the **Help** push button or pressing the **F1** key at any time. You can click on the **Cancel** push button 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.

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 and complete these steps, on another machine, to install a database partition server that will participate in the instance you just created.

If you have finished running the setup program on two different machines, creating both the instance-owning database partition server and a database partition server to participate in the instance, you are finished creating a partitioned database system. Go to “Verifying the Installation” on page 27.

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). If you selected the **Automatically start the DB2 instance at boot time** option during the install, the service's start-up type was set to **Automatic**, otherwise, it was set to **Manual**.

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 6 single-partition instances, they were migrated to the Version 6 single-partition format. If the setup program detected any pre-Version 6 partitioned database instances, they were migrated to the Version 6 multipartition 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*.

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 SAMPLE database on your system, and accessing data from the database.

To verify the installation, if you installed First Steps, perform 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 53.
- Step 3. Click on the **Create the SAMPLE Database** graphic button on the main panel of First Steps.



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 on **OK**.

- Step 4. Once the database is created, click on the **View the SAMPLE database** graphic button to select data from the SAMPLE database. This action launches the Command Center. The Command Center allows you to use a supplied script to view data from the database. Click on the **Execution** icon (the gears icon in the top left corner of the screen) to begin the query.
- Step 5. Click on the **Work with the SAMPLE Database** graphic button to start the Control Center. The Control Center allows you to perform administration tasks on different instance and database objects. For more information, see “Administering Instances and Databases with the DB2 Administration Tools” on page 95.

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 **db2sampl** 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:

```
db2sampl F:
```

This command may take a few minutes to process. 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 catalogued with the database alias SAMPLE when it is created.

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

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

```
db2cmd
db2 connect to sample
db2 "select * from staff where dept = 20"
db2 connect reset
```

For information about entering DB2 commands, see “Entering Commands Using the Command Center” on page 54 or “Entering Commands Using the Command Line Processor” on page 56.



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.

Software Registration

The Software Registration tool is used to register IBM Software with IBM and update IBM's Software Customer Profiles. A customer profile contains identification and demographic information about the users of each IBM software product. The software registration tool contains information and instructions on how to register IBM software by modem, telephone, Internet, FAX, or postal service. To register your software with IBM, perform the following steps:

Step 1. Start the IBM Software Registration tool. For more information, see “Starting the Software Registration Tool” on page 53.

Step 2. Enter the appropriate information and complete the IBM Software Registration Smart Guide. Click on the **Help** push button if you need assistance.



Upgrading from a Try-and-Buy License

To upgrade from a Try-and-Buy license, perform the following steps:

- Step 1. Start the Control Center. For more information, see “Starting the Control Center” on page 54.
- Step 2. Click on the **License Center** icon from the Control Center toolbar. The License Center opens.
- Step 3. Complete the steps to upgrade your DB2 license. For more information, refer to the License Center’s online help by pressing the **F1** key.

If you are upgrading your DB2 product to a higher entitlement, you must ensure that you uninstall your current installation of DB2 and then reinstall your new licensed product.

For example, if you have a Try-and-Buy version of DB2 Workgroup Edition installed on your machine, and you purchased a license for DB2 Enterprise Edition, you must ensure that you remove the installed copy of DB2 Workgroup Edition and then reinstall DB2 Enterprise Edition.

If you did not install the DB2 tools, you can upgrade your Try-and-Buy license using the **db2licm** command. For more information, refer to the *Command Reference*.

Note: *Your Proof of Entitlement* and *License Information* booklets identify the products for which you are licensed.



If you are migrating from a previous version of DB2, go to “Post-Installation Steps” on page 67 for information on how to complete the migration process.

You are now ready to install a DB2 Run-Time Client or a DB2 Administration Client. Go to “Chapter 4. Installing DB2 Clients” on page 33 for more information.

Part 3. Installing and Configuring DB2 Clients

Chapter 4. Installing DB2 Clients

This section describes how to install a DB2 Administration Client, a DB2 Run-Time Client, or a DB2 Software Developer's Kit on your workstation. For information on how to deploy this product using a distributed installation, refer to the *Installation and Configuration Supplement*.

DB2 Administration Clients are available for the following platforms: AIX, HP-UX, Linux, OS/2, Silicon Graphics IRIX, Solaris, Windows 9x, and Windows NT.

DB2 Run-Time Clients and DB2 Software Developer's Kits are available for the following platforms: AIX, HP-UX, Linux, OS/2, Silicon Graphics IRIX, Solaris, Windows 9x, and Windows NT.

Clients that connect to a server must have the appropriate DB2 Administration Client, DB2 Run-Time Client, or DB2 Software Developer's Kit installed.

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



Go to the section that provides installation instructions for the DB2 client that you want to install:

- "Chapter 5. Installing DB2 Clients on Windows 32-Bit Operating Systems" on page 35

For information on how to install all other Version 6 DB2 clients or DB2 Software Developer's Kits, refer to the *Installation and Configuration Supplement*.

To download installation packages for all supported DB2 clients, which include pre-Version 6 DB2 clients such as DOS, Macintosh, SCO OpenServer, SCO Unixware, SINIX, and Windows 3.x, connect to the IBM DB2 Client Application Enabler Web site at <http://www.software.ibm.com/data/db2/db2tech/clientpak.html>.

Chapter 5. Installing DB2 Clients on Windows 32-Bit Operating Systems

This section contains the information you need to install a DB2 client or DB2 Software Developer's Kit on Windows 32-bit operating systems.

Before You Begin

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

- ___ 1. Ensure that your system meets all of the memory, hardware, and software requirements to install your DB2 product. For more information, see "Chapter 2. Planning for Installation" on page 15.
- ___ 2. A user account to perform the installation.

Windows 9x

Any valid Windows 9x user.

Windows NT

Any user account that does not belong to the *Guests* group on the machine where the account is defined.

Installation Steps

To install a DB2 Client or Software Developer's Kit, perform the following steps:

- Step 1. Log on to the system with the user account that you created 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 auto-start, see the tip that follows.



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

- a. Click on **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 73 lists the code for each available language.

- c. Click on **OK**.

Step 4. The Welcome window opens.



Step 5. Respond to the setup program's prompts. Online help is available to guide you through the remaining steps. Invoke online help by clicking on the **Help** push button, or by pressing the **F1** key at any time.

You can click on the **Cancel** push button 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 has:

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



To configure your client to access remote servers, go to “Chapter 6. Configuring Client-to-Server Communications Using the Client Configuration Assistant” on page 39.

Chapter 6. Configuring Client-to-Server Communications Using the Client Configuration Assistant



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.

If you have added a new protocol to your network that is not detectable, or wish 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.

The information in this section describes how to use the Client Configuration Assistant (CCA) to configure your OS/2, Windows 9x, or Windows NT clients to access remote servers. 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 53.

With the CCA, you can:

- Configure database connections that applications can use.
- Update or delete existing configured database connections.
- Display the information for existing configured connections.
- Test a connection to a database.
- Enable or disable databases to be configured as CLI or ODBC data sources.
- Export client profiles, or import client and server profiles which contain information for the setup of a client.
- Update client configuration settings.
- Discover remote databases (if enabled).
- Bind user applications and utilities to databases.
- Change your server password.

The CCA provides three methods to set up a database connection to a server:

- Using a profile.
- Searching the network for databases.
- Manually entering database and communication protocol information for a server.



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

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 D. Naming Rules” on page 75.



If you are adding a database to a system that has a DB2 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 57.

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 53.

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 Database** or **Add** push button to configure a connection.



Go to the section that describes the configuration method that you want to use:

- “Adding a Database Using a Profile”.
- “Adding a Database Using Discovery” on page 42.
- “Adding a Database Manually” on page 44.

Adding a Database Using a Profile

A profile contains information about instances and databases on a system, and databases within each instance. For information on profiles, see “Creating and Using Profiles” on page 46.

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

Step 1. Select the **Use an access profile** radio button and click on the **Next** push button.

Step 2. Click on the ... push button and select a profile.

Step 3. Enter a local database alias name in the **Database alias** field and add a comment that describes this database in the **Comment** field. Click on the **Next** push button.



If you do not specify a database alias name, the default will be the same as the remote database alias name.

Step 4. Register this database as an ODBC data source.



If you are not planning to use ODBC, click on the **Done** push button and go to the next step.

- 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.
- c. If you would like to create an ODBC data source file to share database access, select the **Create a file data source** check box and enter the path and filename for this file in the **File data source name** field.
- d. Click on the **Optimize for application** drop down box and select the application for which you want to tune the ODBC settings.
- e. Click on the **Done** push button to add the database that you selected. The Confirmation window opens.

Step 5. Click on the **Test Connection** push button to test the connection. The Connect to DB2 Database window opens.

If the database that you want to add was not added successfully, click on the **Change** push button to change any settings that you may have incorrectly specified or click on the **Help** push button for more information. For even more detailed information, refer to the *Troubleshooting Guide*.

Step 6. Enter your user ID and password and click on **OK**. If the connection is successful, a message confirming the connection appears.

Step 7. You are now able to use this database. To finish using the Add Database SmartGuide, click on the **Close** push button; you can then add more databases by clicking on the **Add** push button or exit the CCA.

Adding a Database Using Discovery



This option cannot return information about pre-Version 5 DB2 systems 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 the 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 on the **Next** push button.
- Step 2. Click on the [+] sign beside the **Known Systems** icon to list all the systems known to your client.
- Step 3. Click on 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 on 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 on the **Next** push button, 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 on the **Add System** push button

- b. Enter the required communication protocol parameters for the remote Administration Server and click on **OK**. A new system is added. For more information, click on the **Help** push button.
- c. Select the database that you want to add and click on the **Next** push button.

Step 4. Enter a local database alias name in the **Database alias** field and add a comment that describes this database in the **Comment** field. Click on the **Next** push button.



If you do not specify a database alias name, the default will be the same as the remote database alias name.

Step 5. Register this database as an ODBC data source.



If you are not planning to use ODBC, click on the **Done** push button and go to the next step.

- 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.
- c. If you would like to create an ODBC data source file to share database access, select the **Create a file data source** check box and enter the path and filename for this file in the **File data source name** field.
- d. Click on the **Optimize for application** drop down box and select the application for which you want to tune the ODBC settings.
- e. Click on the **Done** push button to add the database that you selected. The Confirmation window opens.

Step 6. Click on the **Test Connection** push button to test the connection. The Connect to DB2 Database window opens.

If the database that you want to add was not added successfully, click on the **Change** push button to change any settings that you may have incorrectly specified or click on the **Help** push button for more information. For even more detailed information, refer to the *Troubleshooting Guide*.

Step 7. Enter your user ID and password and click on **OK**. If the connection is successful, a message confirming the connection appears.

Step 8. You are now able to use this database. To finish using the Add Database SmartGuide, click on the **Close** push button; you can then add more databases by clicking on the **Add** push button or exit the CCA.

Adding a Database Manually

If you have the protocol information for the server that you want to connect to, 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 DB2 database** radio button and click on the **Next** push button.



If you are not using Lightweight Directory Access Protocol (LDAP), this tab will not be shown, skip this step and go to Step 3.

Step 2. 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 on the **Next** push button.
- If you would like to maintain the DB2 directories globally at an LDAP server, select the **Add database using LDAP** radio button and click on the **Next** push button.

Step 3. Select the radio button that corresponds to the protocol that you want to use from the **Protocol** list and click on the **Next** push button.



If you have installed a DB2 Connect product or the DB2 Connect Server Support feature available with DB2 Enterprise Edition or DB2 Enterprise - Extended Edition, and have selected the **TCP/IP** or **APPC** radio button, select the radio button from the **Target operating system** list that corresponds to the type of system where the database that you are trying to connect to resides.

Step 4. Enter the required communication protocol parameters and click on the **Next** push button. For more information, click on the **Help** push button.

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. Add a comment that describes this database in the **Comment** field and click on the **Next** push button.

If you are adding a database that resides on a host or AS/400 system, the **Database name** field will appear as follows:

OS/390 or MVS

Location name

AS/400

RDB name

VM or VSE

DBNAME



When a database is created on the remote server, if a database alias is not specified during database creation, the database is created with a database alias=*database_name*; otherwise, the database alias is the name that was specified when the database was created.

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

- 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.
- c. If you would like to create an ODBC data source file to share database access, select the **Create a file data source** check box and enter the path and filename for this file in the **File data source name** field.
- d. Click on the **Optimize for application** drop down box and select the application for which you want to tune the ODBC settings.
- e. Click on the **Done** push button to add the database that you selected. The Confirmation window opens.



If you are adding a database using the *manual* method, you can also select the **Node Options** tab, the **Security Options** tab, or the **Host or AS/400 Options** tab (if you are adding a host or AS/400 database) to specify more detailed catalog information. Refer to the online help in each tab for more information.

If you want to use the default settings for the node, security, and host or AS/400 (if applicable) options, click on the **Done** push button and go to the next step.

Step 7. Click on the **Test Connection** push button to test the connection. The Connect to DB2 Database window opens.

If the database that you want to add was not added successfully, click on the **Change** push button to change any settings that you may

have incorrectly specified or click on the **Help** push button for more information. For even more detailed information, refer to the *Troubleshooting Guide*.

- Step 8. Enter your user ID and password and click on **OK**. If the connection is successful, a message confirming the connection appears.
- Step 9. You are now able to use this database. To finish using the Add Database SmartGuide, click on the **Close** push button; you can then add more databases by clicking on the **Add** push button or exit the CCA.



You can use the Export feature 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 feature 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. You can use either a *Server Profile* or a *Client Profile* to configure database connections on a client.



Go to the section that describes the type of profile that you want to use:

- “Server Profiles”.
- “Client Profiles” on page 47.

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.

To generate a Server Profile, use the Export Server Profile function provided in the Control Center. When a profile is generated, it includes instances that

have the *discover_inst* database manager configuration parameter and databases with the *discover_db* database configuration parameter set to *ENABLE*. The *discover* parameter in the Administration Server's configuration file must be set to either *SEARCH* or *KNOWN* to generate a profile for a server system.

The default settings for these configuration parameters enable you to discover all instances and databases on your network. For information on setting the *discover_inst*, *discover_db*, and *discover* configuration parameters, refer to the *Administration Guide*.



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 Control Center” on page 54.

Step 2. Select the system that you want to create a profile for and click on the right mouse button.



If the system that you want to create a profile for is not shown, select the **Systems** icon, click on the right mouse button, and select the **Add** option. Click on 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 click on the **Export** push button.



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 40.

Client Profiles

Information in a Client Profile can be used to configure clients using the Import function in the Client Configuration Assistant (CCA). These clients can import all or a subset of the configuration information in a profile.



This scenario assumes that the database connections configured on one client will be exported and used to set up one or more clients.

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, 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 53.

Step 2. Click on the **Export** push button. 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 on **OK**, and go to Step 6.
- 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 on **OK**, and go to Step 6.
- 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 on **OK**, and go to the next step.

Step 4. Select the databases to be exported from the **Available DB2 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 on 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 on the **Help** push button.

- Step 6. Click on **OK**. The Export Client Profile window opens.
- Step 7. Enter a path and file name for this Client Profile and click on **OK**. The DB2 Message window opens.
- Step 8. Click on **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 53.
- Step 2. Click on the **Import** push button. The Select Profile window opens.
- Step 3. Select a Client Profile to import and click on **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 all the databases defined 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 on **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* online document.

- Step 6. You are presented with a list of systems, instances, and databases. Select the database that you want to add and click on the **Next** push button.
- Step 7. Enter a local database alias name in the **Database alias** field and add a comment that describes this database in the **Comment** field. Click on the **Next** push button.



If you do not specify a database alias name, the default will be the same as the remote database alias name.

- Step 8. Register this database as an ODBC data source.



If you are not planning to use ODBC, click on the **Done** push button and go to the next step.

- Ensure that the **Register this database for ODBC** check box is selected.
- 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.
- c. If you would like to create an ODBC data source file to share database access, select the **Create a file data source** check box and enter the path and filename for this file in the **File data source name** field.
 - d. Click on the **Optimize for application** drop down box and select the application for which you want to tune the ODBC settings.
 - e. Click on the **Done** push button to add the database that you selected. The Confirmation window opens.
- Step 9. Click on the **Test Connection** push button to test the connection. The Connect to DB2 Database window opens.
- If the database that you want to add was not added successfully, click on the **Change** push button to change any settings that you may have incorrectly specified or click on the **Help** push button for more information. For even more detailed information, refer to the *Troubleshooting Guide*.
- Step 10. Enter your user ID and password and click on **OK**. If the connection is successful, a message confirming the connection appears.
- Step 11. You are now able to use this database. To finish using the Add Database SmartGuide, click on the **Close** push button; you can then add more databases by clicking on the **Add** push button or exit the CCA.

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 Software Registration Tool”.
 - “Starting the Client Configuration Assistant”.
 - “Starting the Control Center” on page 54.
 - “Entering Commands Using the Command Center” on page 54.
 - “Entering Commands Using the Command Line Processor” on page 56.
 - “Working with the System Administrative Group” on page 57.
-

Starting First Steps

Start First Steps as follows:

Windows 9x or Windows NT

Click on **Start** and select **Programs->DB2 for Windows->First Steps**

Starting the Software Registration Tool

Start the Software Registration Tool as follows:

Windows 9x or Windows NT

Click on **Start** and select **Programs->DB2 for Windows->Registration**

Starting the Client Configuration Assistant

Start the Client Configuration Assistant (CCA) as follows:

Windows 9x or Windows NT

Click on **Start** and select **Programs->DB2 for Windows->Client Configuration Assistant**



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

Starting the Control Center

Start the Control Center in one of the following ways:

On Windows 9x or Windows NT

- Enter the **db2cc** command at a command prompt.
- or
- Click on **Start** and select **Programs->DB2 for Windows->Control Center**.



You can also run the Control Center as an applet through a web browser. For more information, refer to the Control Center readme, which can be found in the `x:\sql1ib\cc\prime` directory, where `x:` is drive where you installed your DB2 product.

Entering Commands Using the Command Center

This section describes how to enter commands using the Command Center. The Command Center provides an interactive window to:

- 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. You do this by invoking Visual Explain in the interactive window.
- 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.

Start the Command Center in one of the follows ways:

- Click on the **Command Center** icon in the Control Center.
- Enter the **db2cctr** command.
- Launch the Command Center from the toolbar as follows:

On Windows 9x or Windows NT

- Click on the **Start** button and select **Programs->DB2 for Windows->Command Center**.

The Command Center contains a large input area in which you enter your commands. To run the commands you have entered, click on the **Execution** icon (the gears icon), or press **CTRL+Enter**.



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 the **Enter** key to start the next command on a new line. The default termination character is a semicolon (;). To specify a different termination character, click on the **Tools Settings** icon in the menu toolbar.

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 **Execution** icon (or pressed **CTRL+Enter**), the Command Center switches to the Results window which informs you how the commands are proceeding.

To recall commands that you have entered, select the **Script** Tab, click on the drop down box, and select a command.

To save commands as scripts, select **Script->Save as** from the menu bar. For more information, click on the **Help** push button or press the **F1** key.



If you want to store commonly used SQL statements or DB2 commands as scripts, click on the **Script Center** icon from the main tool bar. 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:

Command Line Mode

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 Command Line Mode) 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*.

Command Line Mode

To invoke a DB2 Command window, do the following:

Windows 9x or Windows NT

Click on **Start** and select **Programs->DB2 for Windows->Command Window**



You can also invoke the DB2 command window by entering the `db2cmd` command at an MS-DOS 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"
```

If you need to enter a long command that does not fit on one line, use the line continuation character, `\`. When you have reached the end of the line, press the **Enter** key to continue entering the command on the next line. 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 9x or Windows NT

Click on **Start**, and select **Programs->DB2 for Windows->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 an MS-DOS 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
```

If you need to enter a long command that does not fit on one line, use the line continuation character, \. When you have reached the end of the line, press the **Enter** key to continue entering the command on the next line. 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 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 file to specify a new group. If you want that group enumerated on the local machine, the you must also set the `DB2_GRP_LOOKUP` registry variable.



By default, in a Windows NT domain environment, only domain users that belong to the Administrators group at the Primary Domain Controller (PDC) have SYSADM authority on an instance. 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 PDC, we suggest that you create a global group and add the users (both domain and local) that you want to grant SYSADM authority. 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*.

Appendix B. Migrating from Previous Versions and Releases

This section describes how to migrate previous versions of DB2 to the Version 6 format.

DB2 Enterprise - Extended Edition Version 6 supports the migration of DB2 Common Server Version 2.x, Database Server Version 4.x, and DB2 Universal Database Version 5.x to a format usable by Version 6.

Migrating from Previous Versions of DB2

This section describes the steps required to migrate either a DB2 Common Server Version 2.x, a Database Server Version 4.x, or a DB2 Universal Database Version 5.x single-partition database system to a Version 6 partitioned database system. It also includes information on how to update a Version 6 single-partition database system to the Version 6 multipartition format.



The migration process for Database Server Version 4 is identical to that used for DB2 Common Server Version 2. Whenever Version 2 is mentioned in this section, the same information also applies to Version 4.

When you migrate from a Version 5.x database system to a DB2 Enterprise - Extended Edition Version 6 database system, depending on the installation option you choose, the setup program will migrate all instances to the Version 6 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.

If you are migrating a DB2 Version 2.x database to DB2 Enterprise - Extended Edition, the database is relocated from the `\instance_name` directory to the `\instance_name\NODE0000` directory. A node is created with a default name of NODE0000.



This section contains the following information:

- “Pre-Installation Migration Steps”.
- “Post-Installation Steps” on page 67.



To avoid difficulties during database migration, read this entire appendix before starting the migration process.

Pre-Installation Migration Steps

These steps will help you to ensure that all databases on your system can be migrated to the new DB2 Version 6 format. Before installing your new version of DB2, perform the following steps:

Step 1. Back Up Databases

You should back up all databases before installing your new version of DB2. To back up the databases, perform the following steps:

Step 1. Complete all database transactions.

Step 2. Ensure all applications disconnect from each database.

To list all applications that are connected to a database, 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
```

To force all applications to disconnect from the database, enter the **db2 force applications all** command.

Step 3. Ensure 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
```

Step 4. Make a backup copy of all databases. For more information on backing up databases, refer to the *Administration Guide* for the DB2 version you are backing up. For the syntax of the backup command, refer to the *Command Reference* for the DB2 version you are backing up.



Make sure that this is the most recent backup copy of the database before you start the next procedure.

Step 5. Stop the database manager by entering the **db2stop** command.

Step 2. Verify that Databases Can Be Migrated

DB2 provides the **db2ckmig** command to check that databases can be migrated. This command must be run prior to installation. The command is located on the product CD-ROM.


To run the **db2ckmig** command:

1. Insert the CD-ROM into the drive.



If you are installing DB2 on a Windows NT workstation, the setup program might be started automatically via the operating system's autorun feature. In this case, do not proceed with the install. Instead, cancel and proceed to the next step.

2. Enter the **db2ckmig** command to verify that the databases on your system can be correctly migrated. The syntax of the command is as follows:



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 the **/p** parameter is specified.

/p password Specifies the password of the user account used to connect to the database. This parameter must be specified if the **/u** parameter is specified.



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.

3. 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, see Table 7 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.

Table 7. Correcting Error Messages

Error	Action
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.

Table 7. Correcting Error Messages (continued)

Error	Action
The Version 2 database contains database objects that have a schema name of SYSCAT, SYSSTAT, or SYSFUN	<p>These schema names are reserved for the Version 6 database manager. To correct this error, perform the following steps:</p> <p>Step 1. Back up the database.</p> <p>Step 2. Export the data from the database object (catalogs or tables).</p> <p>Step 3. Drop the object.</p> <p>Step 4. Recreate the object with another schema name.</p> <p>Step 5. Import/Load the data into the object.</p> <p>Step 6. Run the db2ckmig command against the database again, ensuring that the database passes the db2ckmig check.</p> <p>Step 7. Make a backup copy of the database. For more information, refer to the <i>Administration Guide</i>.</p>

Table 7. Correcting Error Messages (continued)

Error	Action
<p>The Version 2 database contains database objects that have a dependency on the SYSFUN.DIFFERENCE function. Possible violated database objects are:</p> <ul style="list-style-type: none"> • Constraint • Function • Trigger • View 	<p>The SYSFUN.DIFFERENCE function must be dropped and recreated during database migration. However, if there is a database object that is dependent on this function, migration will fail. To correct this error:</p> <p>Constraint Enter the alter table command to drop the constraint.</p> <p>Function Enter the drop function command to drop the function dependent on SYSFUN.DIFFERENCE.</p> <p>Trigger Enter the drop trigger command to drop the trigger.</p> <p>View Enter the drop view command to drop the view.</p> <p>Note: Any package dependent on the SYSFUN.DIFFERENCE function will be marked inoperative after migration. Therefore, the db2ckmig command will not report any package that is dependent on the SYSFUN.DIFFERENCE function.</p>

Table 7. Correcting Error Messages (continued)

Error	Action
The database contains user-defined distinct types (UDTs) that use the type name BIGINT, DATALINK, REAL or REFERENCE.	<p>These data type names are reserved for the Version 6 database manager. To correct this error, perform the following steps:</p> <p>Step 1. Back up the database.</p> <p>Step 2. Export the data from any tables that are dependent on the data types.</p> <p>Step 3. Drop any tables dependent on the data types, and then drop the data types. These drops may drop other objects such as views, indexes, triggers, or functions.</p> <p>Step 4. Create data types with different type names and recreate the tables using the new data type names. Recreate any dropped views, indexes, triggers, or functions.</p> <p>Step 5. Import/Load the data into the object.</p> <p>Step 6. Run the db2ckmig command against the database again, ensuring that the database passes the db2ckmig check.</p> <p>Step 7. Make a backup copy of the database. For more information, refer to the <i>Administration Guide</i>.</p>

Table 7. Correcting Error Messages (continued)

Error	Action
Structured type and function have the same name.	<p>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 have to be dropped and recreated using another name. To correct this error, perform the following steps:</p> <p>Step 1. Back up the database.</p> <p>Step 2. Export the data from any tables that are dependent on the structured types or functions.</p> <p>Step 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.</p> <p>Step 4. Create structured types or functions with different type or function names and recreate the tables using the new data type or function names. Recreate any dropped views, indexes, triggers, or functions.</p> <p>Step 5. Import/Load the data into the object.</p> <p>Step 6. Run the db2ckmig command against the database again, ensuring that the database passes the db2ckmig check.</p> <p>Step 7. Make a backup copy of the database. For more information, refer to the <i>Administration Guide</i>.</p>

Migration Considerations for the DB2 Version 2.x User Exit Program



These instructions apply only to the DB2 Version 2.x **db2uexit** user exit program. If you are not using the Version 2.x **db2uexit** user exit program, skip this section and go to “Installing DB2 Version 6” on page 67.

DB2 Version 6 uses the **db2uexit** user exit program to archive and retrieve log files. For more information on the **db2uexit** interfaces, refer to the *Administration Guide*.

The following should be considered before installing DB2 Enterprise - Extended Edition Version 6 if you are using the Version 2.x **db2uexit** user exit program.

- If the Version 2.x user exit program, **db2uexit.exe**, is found in the `\sqllib\bin` directory before installation, it will remain in this directory after the installation completes. The **db2uext2.exe** program will also be installed in this directory. Its function is to invoke the **db2uexit.cmd** or **db2uexit.exe** user exit programs using the Version 2.x interface. This allows the old user exit program to be used on Version 6.
- If **db2uexit.exe** is in a directory other than the `sqllib\bin` directory, it will remain there after installation, but **db2uext2.exe** will not be installed in the `sqllib\bin` directory. Following installation, if you want to use the old user exit program, you will have to copy it to the `sqllib\bin` directory, then copy **db2uext2.v2** from the `sqllib\misc` directory to the `sqllib\bin` directory, and rename it to **db2uext2.exe**.

If you are migrating from DB2 Version 2.x, you should modify your user exit program to use the DB2 Version 6 interfaces. The new user exit program **db2uexit** should replace **db2uext2** in the `sqllib\bin` directory.

Installing DB2 Version 6

After you have successfully completed the pre-installation steps, you can now start installing DB2 Enterprise - Extended Edition Version 6.

During the installation of DB2 Enterprise - Extended Edition Version 6, instance migration occurs for instances created in previous versions of DB2. For installation instructions, see “Chapter 3. Installing DB2 Universal Database Enterprise - Extended Edition on Windows NT” on page 21.

Post-Installation Steps

After installing DB2 Version 6, you can now migrate databases and complete other migration activities. Do not scale the database system before migrating all the databases; otherwise, the database migration will fail.

You should perform the following steps after DB2 is installed to complete the migration process:

Pre-Version 6	<ol style="list-style-type: none"> 1. Migrate your databases. See “Migrating Databases” for more information. 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 69 for details.
DB2 Universal Database Version 5.x Enterprise - Extended Edition	<p>Migrate your databases.</p> <p>See “Migrating Databases” for more information.</p>
DB2 Universal Database Version 6	<p>Update all instances that you want to use as partitioned database systems.</p> <p>See “Updating Single-Partition Instances for use in Partitioned Database Systems” on page 69 for details.</p>

Migrating Databases

If you are updating a single-partition Version 6 instance to the Version 6 multipartition format, you do not need to migrate the databases in that instance.

When pre-Version 6 databases are migrated for use in a Version 6 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 pre-Version 6 databases owned by an instance, perform the following steps:

- Step 1. Log in with a user account that has SYSADM authority. See “Working with the System Administrative Group” on page 57 for more information.
- Step 2. Ensure that the databases you want to migrate are cataloged. To retrieve a list of all catalogued databases on your system, enter the **db2 list database directory** command.
- Step 3. Migrate the database using the **db2 migrate database** command. For more information refer to the *Command Reference*.

Updating Single-Partition Instances for use in Partitioned Database Systems

After migrating databases, you are then ready to update any 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 \sql1ib 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 6, because at the end of the restore, the database is migrated to Version 6.

- **Migration of Unique Indexes (db2uiddl)**

DB2 Versions 5 and 6 support deferred checking for duplicate index key values until the end of UPDATE statements. This ensures that temporary duplicate index key values which *may* be present in mid-UPDATE, but no longer are present at the end of the UPDATE, will not cause the statement to fail.

With DB2 Version 2 the same UPDATE statement may fail because checking for duplicate key index values is performed row by row as the statement processes the table. For example, if a row with value 1 is changed to value 2, but a row with value 2 already exists, a duplicate value 2 will be detected causing the DB2 Version 2 UPDATE statement to fail.

Version 2.x and 5.x unique indexes are not automatically migrated to Version 6 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 6 semantics. You have to convert unique indexes to Version 6 semantics only if support for deferred uniqueness checking is required.

To convert unique indexes, you need to perform the following steps:

- Step 1. Log in with a user account that has SYSADM authority. See “Working with the System Administrative Group” on page 57 for more information.
- Step 2. Start the database manager by entering the **db2start** command.
- Step 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.
- Step 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.

- Step 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.
- Step 6. Execute the output file, generated by the **db2uiddl** command, as a DB2 Command Line Processor command file, using a command similar to the following:

```
db2 -tvf filename
```

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

- Step 7. Disconnect from the database by entering the **db2 connect reset** command.

DB2 interprets the re-creation of an existing unique index using the **db2uiddl** command to signal that the index is ready to be converted to Version 6 semantics.

- **Update Statistics**

When database migration is completed, the old statistics that are used to optimize query performance are retained in the catalogs. However, Version 6 of DB2 has statistics that are modified or do not exist in Versions 2.x or 5.x. To take advantage of these, 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 6 database manager.

For better performance we recommend that you run the **db2rbind** command to rebuild all packages stored in the database. In DB2 Version 6 this command has a new option, **all**, which, when specified, rebinds all packages (valid and invalid). If the **all** option is not specified with the **db2rbind** command, only those packages marked as invalid are rebound. 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 6 defaults or to other values during database migration. The same is true for database manager configuration parameters which may have changed to Version 6 defaults or to other values. 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 6, 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 2 or Version 5 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 migration utility, **db2exmig**, will preserve any user added columns on the explain tables.

Appendix C. 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.

Note: The code page values in the table that follows are also used as directory names on DB2 CD-ROMs. For example, a reference to `x:\language\win32\install` would be `x:\en\win32\install` for English. For more detailed information on the languages and code pages support, refer to the *Administration Guide*.

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

Table 8. Languages and Code Pages (continued)

Country Code	Language
gr	Greek
hu	Hungarian
il	Hebrew
it	Italian
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 D. 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 76
 - “Username, User ID, Group Name, and Instance Name Rules” on page 77
 - “Password Rules” on page 79
 - “DB2SYSTEM Naming Rules” on page 79
 - “Workstation Name (nname) Rules” on page 77
-

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.

When naming a database or database alias, see “General Naming Rules” on page 75.

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 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 75.

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; and
 - column 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

*Username*s or *User IDs* are the identifiers assigned to individual users. When naming users, groups, or instances, see “General Naming Rules” on page 75.

In addition, the name you specify:

- Can contain 1 to 8 characters
- Cannot be any of the following:
 - USERS
 - ADMINS
 - GUESTS
 - PUBLIC
 - LOCAL
- Cannot begin with:
 - IBM
 - SQL
 - SYS
- 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 75.

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 Workstation <i>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 NT Computer name and assigns it to *DB2SYSTEM*.

When creating a *DB2SYSTEM* name, see “General Naming Rules” on page 75.

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 9x or Windows NT
A maximum of 14 characters.

Appendix E. About the DB2 Family of Products

The DB2 Family of Products consists of DB2 Universal Database (UDB) and DB2 Connect, as well as some associated DB2 products. Each product is available in different versions that address the needs of today's marketplace. This section provides a brief description of the IBM DB2 Family of Products:

DB2 Universal Database

DB2 Universal Database Enterprise - Extended Edition

DB2 UDB Enterprise - Extended Edition provides a relational database management system that is web-enabled with Java support; scalable from single processors to symmetric multiprocessors; and multimedia capable with image, audio, video, and text support. With DB2 Universal Database, local and remote client applications can create, update, control, and manage relational databases using Structured Query Language (SQL), DB2 APIs (Application Programming Interfaces), ODBC (Open Database Connectivity), JDBC (Java Database Connectivity), SQLJ (Embedded SQL for Java), or DB2 CLI (Call Level Interface).

DB2 Universal Database Enterprise - Extended Edition offers the ability to partition a database across multiple independent machines of a common platform. To the end-user and application developer, the partitioned database still appears as a single database on a single machine. This fully scalable database system enables an application to use multiple machines for a database that is too large for a single machine to handle efficiently. SQL operations and utilities can execute in parallel both within and between the individual database partitions, which can speed up the execution time of a single query or command.

DB2 Universal Database Enterprise - Extended Edition includes DB2 Connect functionality that allows access to DB2 databases on AS/400, MVS/ESA, OS/390, VM, or VSE systems. This product also includes the Federated Database Object Support feature which provides transparent read access to a collection of heterogeneous and semiautonomous data sources, and the ability to perform Distributed Requests.

This product is currently available for AIX, Solaris, and Windows NT operating systems.

DB2 Universal Database Enterprise Edition

DB2 UDB Enterprise Edition provides a relational database management system that is web-enabled with Java support; scalable from single processors to symmetric multiprocessors; and multimedia

capable with image, audio, video, and text support. With DB2 Universal Database, local and remote client applications can create, update, control, and manage relational databases using Structured Query Language (SQL), DB2 APIs (Application Programming Interfaces), ODBC (Open Database Connectivity), JDBC (Java Database Connectivity), SQLJ (Embedded SQL for Java), or DB2 CLI (Call Level Interface).

DB2 Universal Database Enterprise Edition also includes the DB2 Connect functionality that allows access to DB2 databases on AS/400, MVS/ESA, OS/390, VM, and VSE systems. This product also includes the Federated Database Object Support feature which provides transparent read access to a collection of heterogeneous and semiautonomous data sources, and the ability to perform Distributed Requests.

This product is currently available on AIX, HP-UX, Linux, OS/2, Solaris and Windows NT operating systems.

DB2 Universal Database Workgroup Edition

DB2 UDB Workgroup Edition provides a relational database management system that is web-enabled with Java support; scalable from single processors to symmetric multiprocessors; and multimedia capable with image, audio, video, and text support. With DB2 Universal Database, local and remote client applications can create, update, control, and manage relational databases using Structured Query Language (SQL), DB2 APIs (Application Programming Interfaces), ODBC (Open Database Connectivity), JDBC (Java Database Connectivity), SQLJ (Embedded SQL for Java), or DB2 CLI (Call Level Interface).

This product is currently available for Linux, OS/2, and Windows NT.

DB2 Universal Database Personal Edition

DB2 UDB Personal Edition provides a relational database management system that is web-enabled with Java support; and multimedia capable with image, audio, video, and text support. DB2 Universal Database enables local applications to create, update, control, and manage relational databases using the same rich set of APIs as DB2 Enterprise Edition.

DB2 Personal Edition can also act as a client to access remote DB2 servers and can accept inbound Administration Client requests to remotely manage its resources.

This product is currently available for OS/2, Linux, Windows 9x, and Windows NT.

DB2 Universal Database Satellite Edition

DB2 Satellite Edition is a small-footprint version of DB2 Universal

Database that is appropriate for single-user systems, both mobile and branch offices, that occasionally connect to a DB2 control server to exchange data with corporate systems. Administration of DB2 Satellite Edition is scalable, and can be performed through batch jobs that are stored in a control database that resides on a DB2 Workgroup, DB2 Enterprise, or DB2 Enterprise - Extended Edition server.

This product is currently available for Windows 9x and Windows NT.

DB2 Connect

DB2 Connect Enterprise Edition

DB2 Connect Enterprise Edition is a connectivity server that concentrates and manages connections from multiple desktop clients and web applications to DB2 database servers running on host or AS/400 systems. IBM's DB2 for AS/400, DB2 for OS/390, and DB2 for VSE & VM databases continue to be the systems of choice for managing most critical data for the world's largest organizations. While these host and AS/400 databases manage the data, there is a great demand to integrate this data with applications running on Windows, UNIX, OS/2 and Apple workstations.

DB2 Connect Enterprise Edition enables local and remote client applications to create, update, control, and manage DB2 databases and host systems using Structured Query Language (SQL), DB2 APIs (Application Programming Interfaces), ODBC (Open Database Connectivity), JDBC (Java Database Connectivity), SQLJ (Embedded SQLJ for Java), or DB2 CLI (Call Level Interface). In addition, DB2 Connect supports Microsoft Windows data interfaces such as ActiveX Data Objects (ADO), Remote Data Objects (RDO), and OLE DB.

DB2 Connect Enterprise Edition is currently available for AIX, HP-UX, Linux, OS/2, Solaris, and Windows NT operating systems. These servers provide support for applications running on Windows 3.1, Windows 9x, Windows NT, UNIX (AIX, SCO UnixWare 7, Solaris, HP-UX, Linux, Silicon Graphics IRIX, SINIX), OS/2, and Apple Macintosh workstations.

DB2 Connect Personal Edition

DB2 Connect Personal Edition provides access from a single workstation to DB2 databases residing on servers such as MVS/ESA, OS/390, OS/400, VM and VSE, as well as to DB2 Universal Database servers on Windows NT, UNIX, and OS/2. DB2 Connect Personal Edition provides the same rich set of APIs as DB2 Connect Enterprise Edition, and also features integrated SNA support on all Windows platforms.

This product is currently available for OS/2, Linux, Windows 9x, and Windows NT operating systems.

Associated DB2 Products

DB2 Universal Developer's Edition

DB2 Universal Developer's Edition provides all the tools and software you need to create and test multimedia database client/server applications that can run on any DB2 UDB product.

DB2 Universal Developer's Edition contains a collection of DB2 Universal Database servers, DB2 clients, DB2 Connect products, DB2 Software Developer's Kits, extenders for audio, video, image, and text, and application development tools for all supported operating systems.

DB2 Personal Developer's Edition

DB2 Personal Developer's Edition provides all the tools and software you need to create and test multimedia database applications that run only on DB2 UDB Personal Edition products.

The DB2 Personal Developer's Edition contains a collection of DB2 Universal Database servers, DB2 clients, DB2 Connect Personal Edition, DB2 Software Developer's Kits, extenders for audio, video, image, and text, and application development tools for all supported operating systems.

DB2 Run-Time Client CD-ROMs

DB2 Run-Time Client CD-ROMs contain all the latest DB2 Run-Time Clients. A DB2 Run-Time Client provides the ability for workstations from a variety of platforms to access DB2 databases. These workstations are known as DB2 Run-Time Clients.

DB2 Run-Time Client CD-ROMs are included with all DB2 server and DB2 Developer's Edition product packages.

DB2 Administration Client CD-ROMs

DB2 Administration CD-ROMs contain all the latest DB2 Administration Clients.

A DB2 Administration Client provides the ability for workstations from a variety of platforms to access and administer DB2 databases. These workstations are known as DB2 Administration Clients. The DB2 Administration Client that matches the operating system of the DB2 product that is installed is included as a component with all DB2 Universal Database and DB2 Connect products.

The DB2 Administration Client has all of the features of the DB2 Run-Time Client and also includes all the DB2 Administration GUI tools, documentation, and support for Thin Clients. The DB2

Administration Client CD-ROMs are included with all DB2 server and DB2 Developer's Edition product packages.

DB2 DataJoiner

DB2 DataJoiner allows you to access data residing on multiple and diverse platforms, both IBM and multi-vendor, relational and non-relational, as a single database image. With DB2 DataJoiner, you can access all the data in your enterprise as if it were local.

DB2 Data Links Manager

DB2 Data Links Manager manages data files that are not normally found in a database (for example, engineering blueprints or medical x-rays). These data files can be on a file system outside of the database. Manipulation of this data is managed and controlled by DATALINK values in a DB2 database. Use DB2 Data Links Manager to control access to files that are external to a DB2 database.

This product is currently available for AIX and Windows NT operating systems.

DB2 Query Patroller

DB2 Query Patroller provides query and resource management for decision support systems enabling the success of highly scalable data warehouses. As an application, it takes ODBC queries from a client, analyzes them and then dynamically distributes the workload across different nodes on the DB2 UDB Enterprise - Extended Edition installation.

This product is currently available for AIX and Solaris operating systems.

DB2 Net.Data

IBM Net.Data is an application that allows Web developers to easily build dynamic Internet applications using "Web Macros". Net.Data Web Macros have the simplicity of HTML and the power of dynamic SQL. Net.Data provides database connectivity to a variety of data sources including information stored in relational databases and flat files. Data sources, such as DB2, Oracle, and Sybase, and DRDA enabled data sources, can be on a wide range of platforms.



DB2 Enterprise - Extended Edition, DB2 Enterprise Edition, and DB2 Workgroup Edition are commonly referred to as servers or DB2 servers throughout this book.

Working with DB2 Data

As well as providing a relational database 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. Local applications, and all Java applications (either local or remote), access a database through a DB2 client. All remote applications that are not Java applets must have a DB2 client installed on the client machine before they can access the remote database.

DB2 Version 6 clients are supported on:

- Windows 9x or Windows NT
- UNIX (AIX, HP-UX, Linux, SGI IRIX, and Solaris)
- OS/2



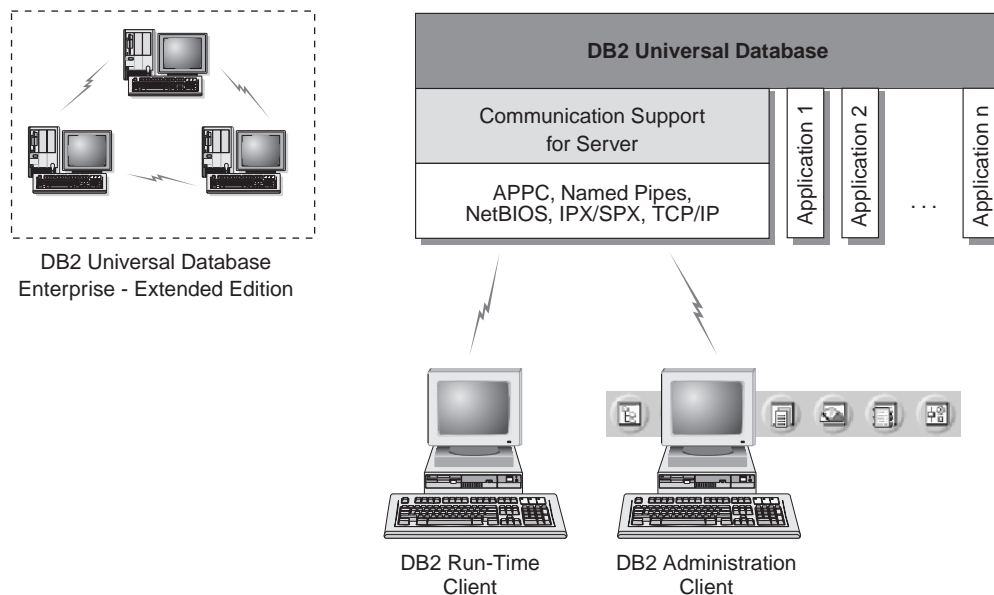
DB2 clients for the following releases and platforms are available for download from the web:

- DB2 Version 1.2 for DOS
- DB2 Version 2.1 for Macintosh
- DB2 Version 2.1 for SCO Open Server
- DB2 Version 5.2 for SCO UnixWare 7
- DB2 Version 2.1 for SINIX
- DB2 Version 5.2 for Windows 3.1

To obtain these clients, connect to the IBM DB2 clients web site at <http://www.software.ibm.com/data/db2/db2tech/clientpak.html>

Figure 4 on page 87 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 4. 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 5 on page 88.

Communicating Between DB2 Universal Database Servers

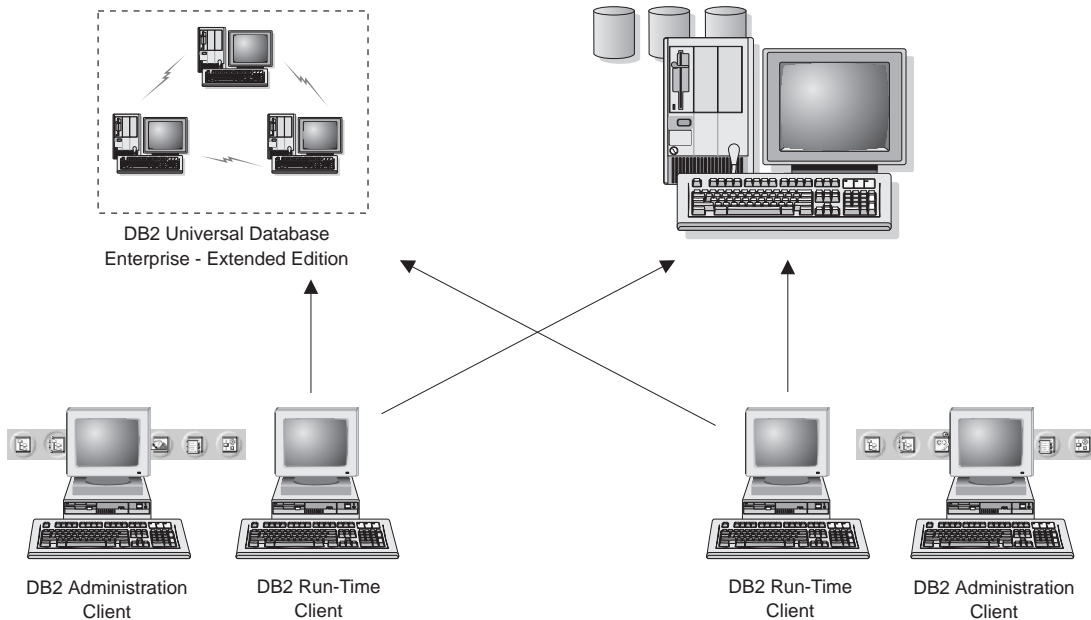


Figure 5. Accessing Data on Multiple Servers

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.



DB2 Universal Database Enterprise Edition and DB2 Universal Database Enterprise - Extended Edition include the **DB2 Connect Server Support** component. All references to DB2 Connect Enterprise Edition also apply to the DB2 Connect Server Support component.

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 monitor such as CICS, Encina, Microsoft Transaction Server (MTS), Tuxedo, Component Broker, and MQSeries are used.

Applications are provided with 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). Use of DRDA allows your applications to establish a fast connection to host and AS/400 databases without expensive host components or proprietary gateways.

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, CICS, Encina, Microsoft Transaction Server) 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 Business Objects, Brio and Cognos, 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.

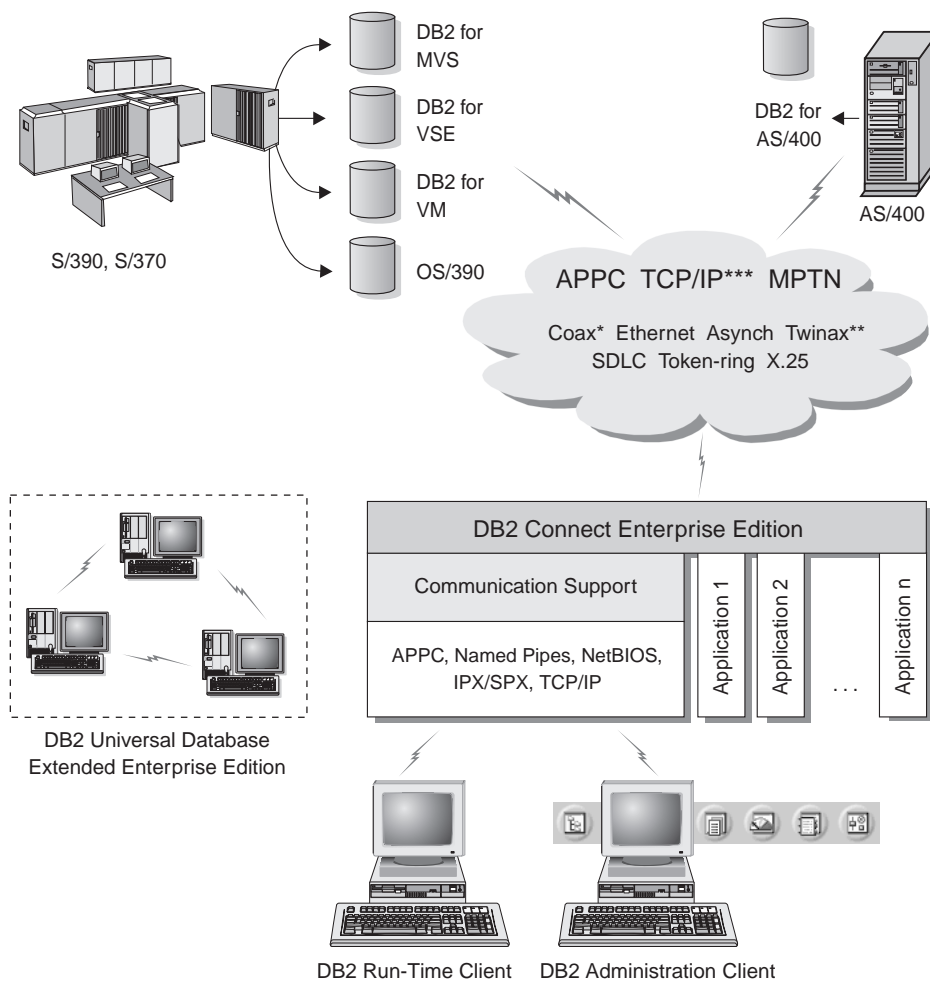
Although DB2 Connect is often installed on an intermediate server machine to connect DB2 clients to a host or AS/400 database, it is also installed on machines where multiple local users want to access the host or AS/400 servers directly. For example, DB2 Connect may be installed on a large machine with many local users. It 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 6 on page 91 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 6. 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 JDBC server and the DB2 client must reside on the same machine as the Web server. The JDBC 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 JDBC server on the machine where the Web server is running. See Figure 7 on page 93 for an example of a Java-enabled browser accessing data from remote DB2 databases.

Accessing DB2 Data Using JDBC

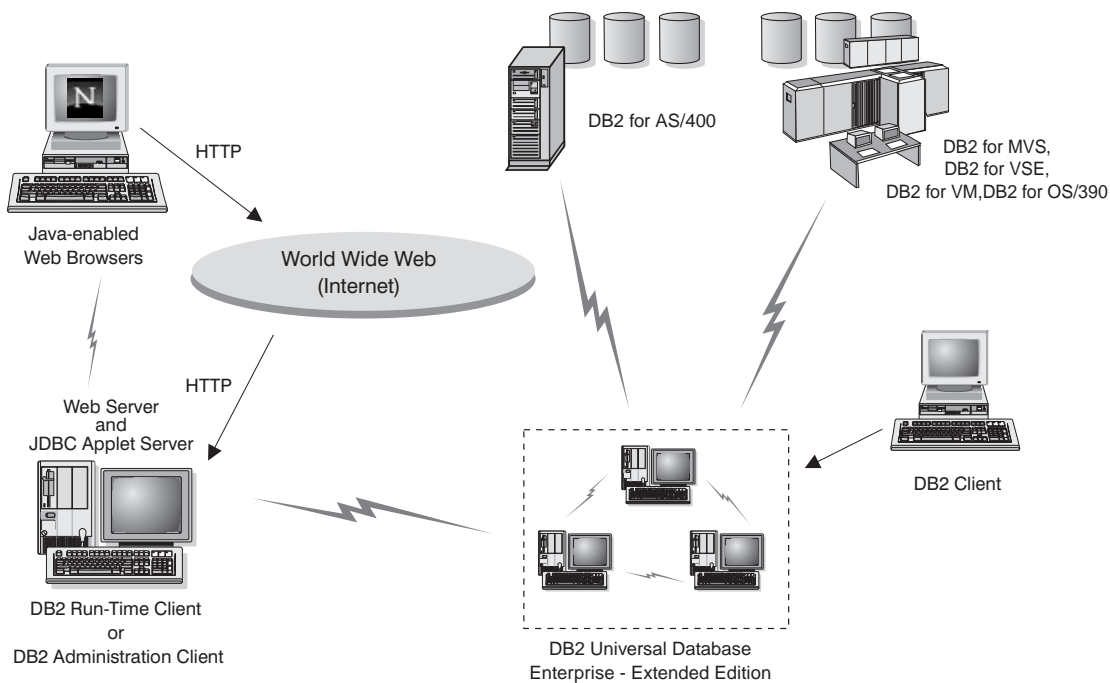


Figure 7. Accessing Data Stored on DB2 Using JDBC

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.software.ibm.com/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.

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.Datapage at <http://www.software.ibm.com/data/net.data>.

Net.Data can be installed with a DB2 server to allow local access to databases. Net.Data can be installed with a 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 8 for an example of a workstation with Net.Data that is being used to access data from a remote DB2 database.

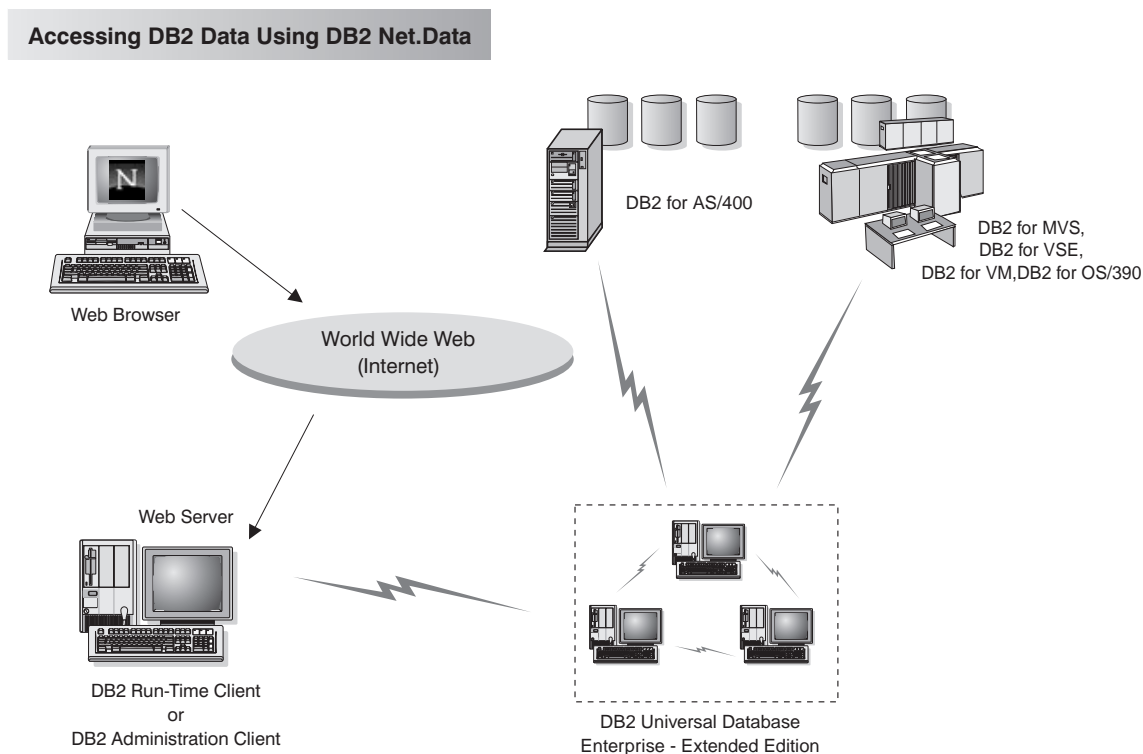


Figure 8. Accessing Internet Data Stored on DB2 Using Net.Data

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 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.

The Control Center for Version 6 has additional support for DB2 UDB for OS/390.

Note: This option is only available on DB2 Enterprise Edition, DB2 Enterprise - Extended Edition, DB2 Connect Personal Edition, and DB2 Connect Enterprise Edition.

If you want to access DB2 for OS/390 functions from the Control Center:

1. Verify the following information with your systems administrator:
 - a. You have a DB2 for OS/390 license (Version 5 or later).
 - b. You are using DB2 UDB for Enterprise Edition.
2. Apply a function modification identifier. Read the DB2 for OS/390 Program Directory. The program directory identifies and describes the contents of FMIDs for each tape or cartridge.

3. Apply any additional service to DB2 as described in the program directory.
4. Ensure that you enabled the stored procedures address space.

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 9 for an example of the main Control Center window.

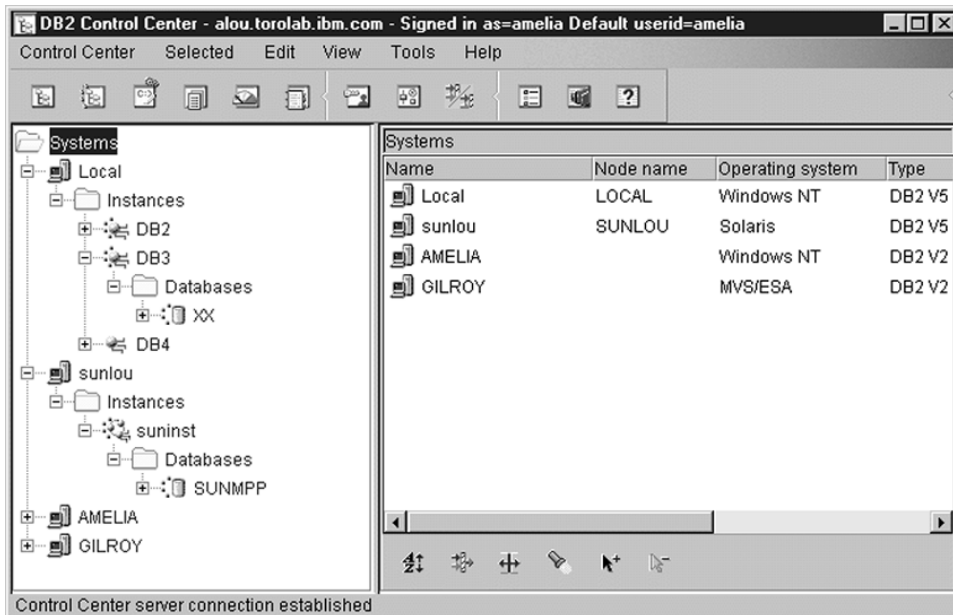


Figure 9. Control Center Main Window

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.

SmartGuides are provided to help you perform complex tasks. For example, a SmartGuide is available to tune the performance of your system. See “Completing Tasks with SmartGuides” on page 105 for descriptions of the various SmartGuides 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.

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.

Tools Setting

Use the Tools Setting to change the settings for the DB2 Administration Tools.

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.

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, 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 99 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 100 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 communication parameters in the database manager 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** option from the pop-up menu.



DB2 Personal Edition does not accept inbound client requests for data. You can only configure inbound communications on a DB2 Personal Edition workstation to allow administrative requests from a DB2 Administration Client.

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 NT Performance Monitor to monitor both database and system performance. For information on how to register DB2 resources and to use the Windows NT 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. This is the preferred method to set up any OS/2, Windows 9x, or Windows NT 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 Client Configuration Assistant, you can:

- Catalog databases so that they can be used by applications. Three methods are available:
 - Search the network for available databases and selecting one. Client access is automatically set up for that database.

- Use a database access profile provided by a database administrator to automatically define your connections. 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.
- 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.

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. These tools are also used by the CCA to catalog databases for a client.

The Administration Server 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 Software Developer's Kit

The DB2 Software Developer's Kit 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 Software Developer's Kit is available for each of the supported operating systems and is currently available in the DB2 Universal Developer's Edition and the DB2 Personal Developer's Edition. Applications that are developed with the DB2 Software Developer's Kit will run on any platform where the equivalent DB2 client component is installed. 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 Software Developer's Kit 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 Software Developer's Kit includes:

- Precompilers for Java, C, C++, COBOL, and FORTRAN.
- Libraries, include files, and code samples to develop applications that use SQLJ and DB2 CLI.
- 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 Software Developer's Kit, 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 Software Developer's Kit 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 F. How the DB2 Library Is Structured

The DB2 Universal Database library consists of SmartGuides, online help, books and sample programs in HTML format. This section describes the information that is provided, and how to access it.

To access product information online, you can use the Information Center. You can view task information, DB2 books, troubleshooting information, sample programs, and DB2 information on the Web. See “Accessing Information with the Information Center” on page 116 for details.

Completing Tasks with SmartGuides

SmartGuides help you complete some administration tasks by taking you through each task one step at a time. SmartGuides are available through the Control Center and the Client Configuration Assistant. The following table lists the SmartGuides.

Note: Create Database, Index, and Configure Multisite Update SmartGuide are available for the partitioned database environment.

SmartGuide	Helps You to...	How to Access...
<i>Add Database</i>	Catalog a database on a client workstation.	From the Client Configuration Assistant, click Add .
<i>Back up Database</i>	Determine, create, and schedule a backup plan.	From the Control Center, click with the right mouse button on the database you want to back up and select Backup->Database using SmartGuide .
<i>Configure Multisite Update SmartGuide</i>	Perform a multi-site update, a distributed transaction, or a two-phase commit.	From the Control Center, click with the right mouse button on the Database icon and select Multisite Update .
<i>Create Database</i>	Create a database, and perform some basic configuration tasks.	From the Control Center, click with the right mouse button on the Databases icon and select Create->Database using SmartGuide .

SmartGuide	Helps You to...	How to Access...
<i>Create Table</i>	Select basic data types, and create a primary key for the table.	From the Control Center, click with the right mouse button on the Tables icon and select Create->Table using SmartGuide .
<i>Create Table Space</i>	Create a new table space.	From the Control Center, click with the right mouse button on the Table spaces icon and select Create->Table space using SmartGuide .
<i>Index</i>	Advise which indexes to create and drop for all your queries.	From the Control Center, click with the right mouse button on the Index icon and select Create->Index using SmartGuide .
<i>Performance Configuration</i>	Tune the performance of a database by updating configuration parameters to match your business requirements.	From the Control Center, click with the right mouse button on the database you want to tune and select Configure using SmartGuide .
<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, click with the right mouse button on the database you want to restore and select Restore->Database using SmartGuide .

Accessing Online Help

Online help is available with all DB2 components. The following table describes the various types of help. You can also access DB2 information through the Information Center. For information see “Accessing Information with the Information Center” on page 116.

Type of Help	Contents	How to Access...
<i>Command Help</i>	Explains the syntax of commands in the command line processor.	From the command line processor in interactive mode, enter: ? <i>command</i> where <i>command</i> is a keyword or the entire command. For example, ? catalog displays help for all the CATALOG commands, while ? catalog database displays help for the CATALOG DATABASE command.

Type of Help	Contents	How to Access...
Control Center Help Client Configuration Assistant Help Event Analyzer Help Command Center Help	Explains the tasks you can perform in a window or notebook. The help includes prerequisite information you need to know, and describes how to use the window or notebook controls.	From a window or notebook, click the Help push button or press the F1 key.
Message Help	Describes the cause of a message, and any action you should take.	<p>From the command line processor in interactive mode, enter:</p> <p><code>? XXXnnnnn</code></p> <p>where <i>XXXnnnnn</i> is a valid message identifier.</p> <p>For example, <code>? SQL30081</code> displays help about the SQL30081 message.</p> <p>To view message help one screen at a time, enter:</p> <p><code>? XXXnnnnn more</code></p> <p>To save message help in a file, enter:</p> <p><code>? XXXnnnnn > filename.ext</code></p> <p>where <i>filename.ext</i> is the file where you want to save the message help.</p>
SQL Help	Explains the syntax of SQL statements.	<p>From the command line processor in interactive mode, enter:</p> <p><code>help statement</code></p> <p>where <i>statement</i> is an SQL statement.</p> <p>For example, help SELECT displays help about the SELECT statement.</p> <p>Note: SQL help is not available on UNIX-based platforms.</p>
SQLSTATE Help	Explains SQL states and class codes.	<p>From the command line processor in interactive mode, enter:</p> <p><code>? sqlstate</code> or <code>? class-code</code></p> <p>where <i>sqlstate</i> is a valid five-digit SQL state and <i>class-code</i> is the first two digits of the SQL state.</p> <p>For example, <code>? 08003</code> displays help for the 08003 SQL state, while <code>? 08</code> displays help for the 08 class code.</p>

DB2 Information – Hardcopy and Online

The table in this section lists the DB2 books. They are divided into two groups:

Cross-platform books

These books contain the common DB2 information for all platforms.

Platform-specific books

These books are for DB2 on a specific platform. For example, there are separate *Quick Beginnings* books for DB2 on OS/2, on Windows NT, and on the UNIX-based platforms.

Cross-platform sample programs in HTML

These samples are the HTML version of the sample programs that are installed with the SDK. They are for informational purposes and do not replace the actual programs.

Most books are available in HTML and PostScript format, or you can choose to order a hardcopy from IBM. The exceptions are noted in the table.

On OS/2 and Windows platforms, HTML documentation files can be installed under the doc\html subdirectory. Depending on the language of your system, some files may be in that language, and the remainder are in English.

On UNIX platforms, you can install multiple language versions of the HTML documentation files under the doc/%L/html subdirectories. Any documentation that is not available in a national language is shown in English.

You can obtain DB2 books and access information in a variety of different ways:

View	See “Viewing Online Information” on page 115.
Search	See “Searching Online Information” on page 118.
Print	See “Printing the PostScript Books” on page 118.
Order	See “Ordering the Printed Books” on page 119.

Name	Description	Form Number File Name for Online Book	HTML Directory
Cross-Platform Books			

Name	Description	Form Number File Name for Online Book	HTML Directory
<i>Administration Guide</i>	<p><i>Administration Guide, Design and Implementation</i> contains information required to design, implement, and maintain a database. It also describes database access using the Control Center(whether local or in a client/server environment), auditing, database recovery, distributed database support, and high availability.</p> <p><i>Administration Guide, Performance</i> contains information that focuses on the database environment, such as application performance evaluation and tuning.</p> <p>You can order both volumes of the <i>Administration Guide</i> in the English language in North America using the form number SBOF-8922.</p>	<p>Volume 1 SC09-2839 db2d1x60</p> <p>Volume 2 SC09-2840 db2d2x60</p>	db2d0
<i>Administrative API Reference</i>	Describes the DB2 application programming interfaces (APIs) and data structures you can use to manage your databases. Explains how to call APIs from your applications.	SC09-2841 db2b0x60	db2b0
<i>Application Building Guide</i>	<p>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.</p> <p>This book combines the <i>Building Applications</i> books for the OS/2, Windows, and UNIX-based environments.</p>	SC09-2842 db2axx60	db2ax
<i>APPC, CPI-C and SNA Sense Codes</i>	<p>Provides general information about APPC, CPI-C, and SNA sense codes that you may encounter when using DB2 Universal Database products.</p> <p>Note: Available in HTML format only.</p>	No form number db2apx60	db2ap

Name	Description	Form Number File Name for Online Book	HTML Directory
<i>Application Development Guide</i>	Explains how to develop applications that access DB2 databases using embedded SQL or JDBC, how to write stored procedures, user-defined types, user-defined functions, and how to use triggers. It also discusses programming techniques and performance considerations. This book was formerly known as the <i>Embedded SQL Programming Guide</i> .	SC09-2845 db2a0x60	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.	SC09-2843 db2l0x60	db2l0
<i>Command Reference</i>	Explains how to use the command line processor, and describes the DB2 commands you can use to manage your database.	SC09-2844 db2n0x60	db2n0
<i>Data Movement Utilities Guide and Reference</i>	Explains how to use the Load, Import, Export, Autoloader, and Data Propagation utilities to work with the data in the database.	SC09-2858 db2dmx60	db2dm
<i>DB2 Connect Personal Edition Quick Beginnings</i>	Provides planning, installing, and configuring information for DB2 Connect Personal Edition.	GC09-2830 db2c1x60	db2c1
<i>DB2 Connect User's Guide</i>	Provides concepts, programming and general usage information about the DB2 Connect products.	SC09-2838 db2c0x60	db2c0
<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, and on how to use DRDA application servers with DB2 Connect application requesters. Note: Available in HTML and PostScript formats only.	No form number db2h1x60	db2h1
<i>Glossary</i>	Provides a comprehensive list of all DB2 terms and definitions. Note: Available in HTML format only.	No form number db2t0x50	db2t0

Name	Description	Form Number File Name for Online Book	HTML Directory
<i>Installation and Configuration Supplement</i>	Guides you through the planning, installation, and set up of platform-specific DB2 clients. This supplement contains information on binding, setting up client and server communications, DB2 GUI tools, DRDA AS, distributed installation, and the configuration of distributed requests and access methods to heterogeneous data sources.	GC09-2857 db2iyx60	db2iy
<i>Message Reference</i>	Lists messages and codes issued by DB2, and describes the actions you should take.	GC09-2846 db2m0x60	db2m0
<i>Replication Guide and Reference</i>	Provides planning, configuration, administration, and usage information for the IBM Replication tools supplied with DB2.	SC26-9642 db2e0x60	db2e0
<i>SQL Getting Started</i>	Introduces SQL concepts, and provides examples for many constructs and tasks.	SC09-2856 db2y0x60	db2y0
<i>SQL Reference, Volume 1 and Volume 2</i>	Describes SQL syntax, semantics, and the rules of the language. 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-8923.	SBOF-8923 Volume 1 db2s1x60 Volume 2 db2s2x60	db2s0
<i>System Monitor Guide and Reference</i>	Describes how to collect different kinds of information about databases and the database manager. Explains how to use the information to understand database activity, improve performance, and determine the cause of problems.	SC09-2849 db2f0x60	db2f0
<i>Troubleshooting Guide</i>	Helps you determine the source of errors, recover from problems, and use diagnostic tools in consultation with DB2 Customer Service.	S10J-8169	db2p0

Name	Description	Form Number File Name for Online Book	HTML Directory
<i>What's New</i>	Describes the new features, functions, and enhancements in DB2 Universal Database, Version 6.0, including information about Java-based tools.	SC09-2851 db2q0x60	db2q0
Platform-Specific Books			
<i>Administering Satellites Guide and Reference</i>	Provides planning, configuration, administration, and usage information for satellites.	GC09-2821 db2dsx60	db2ds
<i>DB2 Personal Edition Quick Beginnings</i>	Provides planning, installation, migration, and configuration information for DB2 Universal Database Personal Edition on the OS/2, Windows 95, and Windows NT operating systems.	GC09-2831 db2i1x60	db2i1
<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. Also contains installing and setup information for many supported clients.	GC09-2834 db2i2x60	db2i2
<i>DB2 for UNIX Quick Beginnings</i>	Provides planning, installation, migration, and configuration information for DB2 Universal Database on UNIX-based platforms. Also contains installing and setup information for many supported clients.	GC09-2836 db2ixx60	db2ix
<i>DB2 for Windows NT Quick Beginnings</i>	Provides planning, installation, migration, and configuration information for DB2 Universal Database on the Windows NT operating system. Also contains installing and setup information for many supported clients.	GC09-2835 db2i6x60	db2i6
<i>DB2 Enterprise - Extended Edition for UNIX Quick Beginnings</i>	Provides planning, installation, and configuration information for DB2 Enterprise - Extended Edition for UNIX. Also contains installing and setup information for many supported clients.	GC09-2832 db2v3x60	db2v3

Name	Description	Form Number File Name for Online Book	HTML Directory
<i>DB2 Enterprise - Extended Edition for Windows NT Quick Beginnings</i>	Provides planning, installation, and configuration information for DB2 Enterprise - Extended Edition for Windows NT. Also contains installing and setup information for many supported clients.	GC09-2833 db2v6x60	db2v6
<i>DB2 Connect Enterprise Edition for OS/2 and Windows NT Quick Beginnings</i>	Provides planning, migration, installation, and configuration information for DB2 Connect Enterprise Edition on the OS/2 and Windows NT operating systems. Also contains installation and setup information for many supported clients. This book was formerly part of the <i>DB2 Connect Enterprise Edition Quick Beginnings</i> .	GC09-2828 db2c6x60	db2c6
<i>DB2 Connect Enterprise Edition for UNIX Quick Beginnings</i>	Provides planning, migration, installation, configuration, and usage information for DB2 Connect Enterprise Edition in UNIX-based platforms. Also contains installation and setup information for many supported clients. This book was formerly part of the <i>DB2 Connect Enterprise Edition Quick Beginnings</i> .	GC09-2829 db2cyx60	db2cy
<i>DB2 Data Links Manager for AIX Quick Beginnings</i>	Provides planning, installation, configuration, and task information for DB2 Data Links Manager for AIX.	GC09-2837 db2z0x60	db2z0
<i>DB2 Data Links Manager for Windows NT Quick Beginnings</i>	Provides planning, installation, configuration, and task information for DB2 Data Links Manager for Windows NT.	GC09-2827 db2z6x60	db2z6
<i>DB2 Query Patroller Administration Guide</i>	Provides administration information on DB2 Query Patrol.	SC09-2859 db2dwx60	db2dw
<i>DB2 Query Patroller Installation Guide</i>	Provides installation information on DB2 Query Patrol.	GC09-2860 db2iwx60	db2iw
<i>DB2 Query Patroller User's Guide</i>	Describes how to use the tools and functions of the DB2 Query Patrol.	SC09-2861 db2wwx60	db2ww

Name	Description	Form Number File Name for Online Book	HTML Directory
Cross-Platform Sample Programs in HTML			
Sample programs in HTML	Provides the sample programs in HTML format for the programming languages on all platforms supported by DB2 for informational purposes (not all samples are available in all languages). Only available when the SDK is installed. See <i>Application Building Guide</i> for more information on the actual programs. Note: Available in HTML format only.	No form number	db2hs/c db2hs/cli db2hs/clp db2hs/cpp db2hs/cobol db2hs/cobol_mf db2hs/fortran db2hs/java db2hs/rexx

Notes:

1. The character in the sixth position of the file name indicates the language of a book. For example, the file name db2d0e60 indicates that the *Administration Guide* is in English. The following letters are used in the file names to indicate the language of a book:

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. For late breaking information that could not be included in the DB2 books:

- On UNIX-based platforms, see the `Release.Notes` file. This file is located in the `DB2DIR/Readme/%L` directory, where `%L` is the locale name and `DB2DIR` is:
 - `/usr/lpp/db2_06_01` on AIX
 - `/opt/IBMdb2/V6.1` on HP-UX, Solaris, SCO UnixWare 7, and Silicon Graphics IRIX
 - `/usr/IBMdb2/V6.1` on Linux.
- On other platforms, see the `RELEASE.TXT` file. This file is located in the directory where the product is installed.
- Under Windows Start menu

Viewing Online Information

The manuals 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 on all platforms other than SCO UnixWare 7:

- If you are running DB2 administration tools, use the Information Center. See “Accessing Information with the Information Center” on page 116 for details.
- Select the Open Page menu item of your Web browser. The page you open contains descriptions of and links to DB2 information:
 - On UNIX-based platforms, open the following page:


```
file:/INSTHOME/sql1lib/doc/%L/html/index.htm
```

where `%L` is 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 on the **DB2 Online Books** icon. Depending on the system you are using, the icon is in the main product folder or the Windows Start menu.

To view online books or sample programs on the SCO UnixWare 7:

- DB2 Universal Database for SCO UnixWare 7 uses the native SCOhelp utility to search the DB2 information. You can access SCOhelp by the following methods:
 - entering the "scohelp" command on the command line,
 - selecting the Help menu in the Control Panel of the CDE desktop or
 - selecting Help in the Root menu of the Panorama desktop

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

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.

Depending on your system, you can access the Information Center from the:

- Main product folder
- Toolbar in the Control Center
- Windows Start menu
- Help menu of the Control Center

The Information Center provides the following kinds of information. Click the appropriate tab to look at the information:

Tasks	Lists tasks you can perform using DB2.
Reference	Lists DB2 reference information, such as keywords, commands, and APIs.
Books	Lists DB2 books.
Troubleshooting	Lists categories of error messages and their recovery actions.
Sample Programs	Lists sample programs that come with the DB2 Software Developer's Kit. If the Software Developer's Kit is not installed, this tab is not displayed.
Web	Lists 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 some search capabilities, so you can look for specific topics, and filter capabilities to limit the scope of your searches.

For a full text search, click the Search button of the Information Center follow the *Search DB2 Books* link in each HTML file.

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 by double-clicking its icon on the Windows or OS/2 desktop.

Refer to the release notes if you experience any other problems when searching the HTML information.

Note: Search function is not available in the Linux and Silicon Graphics environments.

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, use the following instructions:

1. Copy all files and subdirectories from `\sqlib\doc\html` on your local system to a Web server. Each book has its own subdirectory containing 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, see the NetQuestion Appendix in *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. Once you are able to view the book files, you should bookmark commonly viewed topics. Among those, 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 setting up a search, see the NetQuestion Appendix in *Installation and Configuration Supplement* book.

Searching Online Information

To search for information in the HTML books, you can do the following:

- Click on **Search the DB2 Books** at the bottom of any page in the HTML books. Use the search form to find a specific topic. This function is not available in the Linux or Silicon Graphics IRIX environments.
- Click on **Index** at the bottom of any page in an HTML book. Use the index to find a specific topic in the book.
- Display the table of contents or index of 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 116 for details.

Printing the PostScript Books

If you prefer to have printed copies of the manuals, you can decompress and print PostScript versions. For the file name of each book in the library, see the table in “DB2 Information – Hardcopy and Online” on page 108. Specify the full path name for the file you intend to print.

On OS/2 and Windows platforms:

1. Copy the compressed PostScript files to a hard drive on your system. The files have a file extension of .exe and are located in the `x:\doc\language\books\ps` directory, where `x:` is the letter representing the CD-ROM drive and `language` is the two-character country code that represents your language (for example, EN for English).
2. Decompress the file that corresponds to the book that you want. Each compressed book is a self-extracting executable file. To decompress the

book, simply run it as you would run any other executable program. The result from this step is a printable PostScript file with a file extension of .ps.

3. Ensure that your default printer is a PostScript printer capable of printing Level 1 (or equivalent) files.
4. Enter the following command from a command line:

```
print filename.ps
```

On UNIX-based platforms:

1. Mount the CD-ROM. Refer to your *Quick Beginnings* manual for the procedures to mount the CD-ROM.
2. Change to /cdrom/doc/%L/ps directory on the CD-ROM, where /cdrom is the mount point of the CD-ROM and %L is the name of the desired locale. The manuals will be installed in the previously-mentioned directory with file names ending with .ps.Z.
3. Decompress and print the manual you require using the following command:

- For AIX:

```
zcat filename | qprt -P PSprinter_queue
```

- For HP-UX, Solaris, or SCO UnixWare 7:

```
zcat filename | lp -d PSprinter_queue
```

- For Linux:

```
zcat filename | lpr -P PSprinter_queue
```

- For Silicon Graphics IRIX:

```
zcat < filename | lp -d PSprinter_queue
```

where *filename* is the full path name and extension of the compressed PostScript file and *PSprinter_queue* is the name of the PostScript printer queue.

For example, to print the English version of *DB2 for UNIX Quick Beginnings* on AIX, you can use the following command:

```
zcat /cdrom/doc/en/ps/db2ixe60.ps.Z || qprt -P ps1
```

Ordering the Printed Books

You can order the printed DB2 manuals either as a set or individually. There are three sets of books available. The form number for the entire set of DB2 books is SB0F-8926-00. The form number for the books listed under the heading "Cross-Platform Books" is SB0F-8924-00.

Note: These form numbers only apply if you are ordering books that are printed in the English language in North America.

You can also order books individually by the form number listed in “DB2 Information – Hardcopy and Online” on page 108. To order printed versions, 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.

Appendix G. Notices

Any reference to an IBM licensed program in this publication is not intended to state or imply that only IBM's licensed program may be used. Any functionally equivalent product, program or service that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program, or service. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, is the user's responsibility.

IBM may have patents or pending patent applications covering subject matter 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 the

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

Licensees 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.

This publication 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.

Trademarks

The following terms are trademarks or registered trademarks of the IBM Corporation in the United States and/or other countries:

ACF/VTAM	MVS/ESA
ADSTAR	MVS/XA
AISPO	OS/400
AIX	OS/390
AIXwindows	OS/2
AnyNet	PowerPC
APPN	QMF
AS/400	RACF
CICS	RISC System/6000
C Set++	SP
C/370	SQL/DS
DATABASE 2	SQL/400
DataHub	S/370
DataJoiner	System/370
DataPropagator	System/390
DataRefresher	SystemView
DB2	VisualAge
DB2 Connect	VM/ESA
DB2 Universal Database	VSE/ESA
Distributed Relational Database Architecture	VTAM
DRDA	WIN-OS/2
Extended Services	
FFST	
First Failure Support Technology	
IBM	
IMS	
LAN Distance	

Trademarks of Other Companies

The following terms are trademarks or registered trademarks of the companies listed:

C-bus is a trademark of Corollary, Inc. in the United States, other countries, or both.

HP-UX is a trademark of Hewlett-Packard.

Java, HotJava, Solaris, Solstice, and Sun are trademarks of Sun Microsystems, Inc.

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

PC Direct is a trademark of Ziff Communications Company in the United States, other countries, or both and is used by IBM Corporation under license.

ActionMedia, LANDesk, MMX, Pentium, and ProShare are trademarks of Intel Corporation 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 Profile
 - creating 47
 - using 46
- accessing data
 - through DB2 Connect 89
 - through Net.Data or JDBC 92
- accessing multiple servers 87
- Administering Satellites Guide and Reference 112
- Administration Guide 108
- Administration Server
 - overview 101
 - setup 61
- Administrative API Reference 109
- ADSM
 - software requirements 18
- Alert Center 98
- APPC
 - software requirements 18, 19
 - support on Windows NT
 - IBM Communications Server for NT 18
 - Microsoft SNA Server 18
 - supported platforms 18
- APPC, CPI-C and SNA Sense Codes 109
- Application Building Guide 109
- application development
 - using Net.Data or JDBC 92
- Application Development Guide 109
- application programs
 - coordinator node 4
 - migrating to DB2 Enterprise - Extended Edition 5

C

- capacity
 - disk storage 7
- catalog node 8
- CLI Guide and Reference 110
- Client Configuration Assistant
 - administering clients
 - communications 40
 - changing privileges 58
 - client profiles 41
 - exporting 48
 - using 47

- Client Configuration Assistant
 - (continued)
 - overview 40
 - searching the network 40
 - using access profiles 46
- client profile
 - exporting 48
- clients
 - installing 33
 - operating systems supported 86
 - overview 86
- column
 - long-field, partitioning key considerations 9
- Command Center
 - entering DB2 commands 54
 - entering SQL statements 54
 - overview 95
- Command Reference 110
- commands
 - alter table 64
 - db2ckmig 61, 66
 - db2start 70
 - db2uiddl 70
 - drop function 64
 - drop trigger 64
 - drop view 64
- communication protocols
 - APPC 18, 19
 - IPX/SPX 18, 19
 - Named Pipes 18, 19
 - NetBIOS 18, 19
 - TCP/IP 18, 19
- communications
 - configuring the client 40
 - managing
 - Client Configuration Assistant 40
 - Control Center 99
- concurrency
 - improving with logical nodes 10
- configuration parameters
 - setting DB2 58
 - SYSADM_GROUP 58
- configuring client communications
 - setting configuration parameters 71
 - using the Client Configuration Assistant 40

- configuring communications
 - using Client Configuration Assistant 40
- configuring server communications
 - setting DB2COMM 47
- connecting to the Internet using Net.Data 92
- Connectivity Supplement 110
- Control Center
 - components 96
 - overview 95
- coordinator node
 - description 4
- cost-based query optimizer
 - description 4

D

- Data Movement Utilities Guide and Reference 110
- database
 - authentication type 67, 68
 - definition 3
 - migration 68
 - naming rules 75
 - partition 3
- database access
 - two-phase commit 88
- database administration tools
 - Control Center 95
 - overview 95
- database alias
 - naming rules 75
- database manager
 - definition 3
- database objects
 - naming rules 76
- database partition server
 - definition 3
- DB2 client
 - accessing databases 87
 - installing 33
 - Windows 9x 35
 - Windows NT 35
 - overview 86, 87
 - platforms supported 86
 - software requirements 18
- DB2 Connect
 - overview 84, 89

- DB2 Connect Enterprise Edition for OS/2 and Windows NT Quick Beginnings 113
 - DB2 Connect Enterprise Edition for UNIX Quick Beginnings 113
 - DB2 Connect Personal Edition Quick Beginnings 110
 - DB2 Connect User's Guide 110
 - DB2 Data Links Manager for AIX Quick Beginnings 113
 - DB2 Data Links Manager for Windows NT Quick Beginnings 113
 - DB2 Enterprise - Extended Edition
 - configuration planning 16
 - hardware configuration 6
 - DB2 Enterprise - Extended Edition for UNIX Quick Beginnings 112
 - DB2 Enterprise - Extended Edition for Windows NT Quick Beginnings 112
 - DB2 Enterprise Edition
 - memory requirements 16
 - DB2 library
 - books 108
 - Information Center 116
 - language identifier for books 114
 - late-breaking information 115
 - online help 106
 - ordering printed books 119
 - printing PostScript books 118
 - searching online
 - information 118
 - setting up document server 117
 - SmartGuides 105
 - structure of 105
 - viewing online information 115
 - DB2 Personal Edition Quick Beginnings 112
 - DB2 Query Patroller Administration Guide 113
 - DB2 Query Patroller Installation Guide 113
 - DB2 Query Patroller User's Guide 113
 - DB2 Universal Database
 - components included
 - Client Configuration Assistant 100
 - Control Center 95, 96
 - DB2 client 86
 - DB2 Performance Monitor 100
 - DB2 Snapshot Monitor 92
 - DB2 Universal Database *(continued)*
 - components included *(continued)*
 - Visual Explain 100, 100
 - overview 96
 - platforms supported 95
 - software requirements 17
 - db2ckmig command 66
 - db2imigr command 61
 - db2uiddl command 70
 - developing applications
 - using Net.Data or JDBC 92
 - disk requirements
 - client 16
 - server 16
 - Distributed Computing Environment
 - software requirements 18
- E**
- Enterprise - Extended Edition
 - capacity 7
 - storage 7
 - Enterprise - Extended Edition configuration
 - logical nodes 10
 - security for an instance 10
 - shared nothing definition 3
 - environment variables 11
 - error handling
 - migration 62
 - execs
 - db2uexit.exe 66, 67
 - db2uext2.exe 66, 67
- F**
- fixed disks
 - hardware requirements 16
- G**
- GB (gigabyte) 7
 - gigabyte (GB) 7
 - Glossary 110
- H**
- hardware requirements
 - fixed disk 16
- I**
- IBMCATGROUP nodegroup 8
 - IBMDEFAULTGROUP nodegroup 8
 - IBMTMPGROUP 8
 - IBMTMPGROUP nodegroup 7
 - installation
 - CID using SystemView LAN 26, 37
 - clients
 - Windows 9x 35
 - installation *(continued)*
 - clients *(continued)*
 - Windows NT 35
 - disk requirements
 - client 16
 - server 16
 - errors 26, 37
 - log 26, 37
 - memory requirements
 - client 15
 - server 15
 - verifying the installation 27
 - Installation and Configuration Supplement 110
 - instance owner 11
 - instance-owning machine
 - definition 11
 - instances
 - definition 10
 - multiple on one processor 11
 - naming restrictions 77
 - inter-operator parallelism 5
- J**
- Java Support 92
 - Journal 98
- L**
- logical node
 - improving concurrency 10
 - purpose 9
 - supported configurations 10
 - long-field columns
 - partitioning key
 - considerations 9
- M**
- managing connections
 - client
 - overview 100
 - using the Client Configuration Assistant 40, 100
 - server
 - overview 100
 - using the Client Configuration Assistant 40
 - managing databases using the Control Center 96
 - managing server communications
 - overview 99
 - memory requirements
 - client 15
 - estimating 15
 - recommended 15
 - server 15
 - Message Reference 111

- Microsoft SNA Server
 - version required 19
- migration
 - database 68
 - DB2 Version 2.x 59, 60
 - db2chkmig 66
 - of applications 5
 - post-installation 67
 - post-migration
 - rebinding packages 71
 - unique index conversion 70
 - updating database and database manager configuration 71
 - updating statistics 71
 - pre-installation 60
 - previous releases 59, 60
 - user exit program 66

N

- naming rules
 - database 75
 - database alias 75
 - database objects 76
 - general 75
 - groups 77
 - instance names 77
 - password 81
 - userids 77
 - username 35, 77
 - workstation name (nname) 77
- National Language Support (NLS)
 - code set 73
 - codepage support 73
 - determining codepage 75
- Net.Data
 - overview 92
- node
 - definition 3
 - in the same instance 10
- nodegroup
 - IBMCATGROUP 8
 - IBMDEFAULTGROUP 8
 - multipartition nodegroup 7

O

- operating systems
 - client 86

P

- parameters
 - SYSADM_GROUP 58
- partition
 - definition 3
- partitioning key
 - definition 9

- partitioning map
 - overview 9
- password
 - naming rules 81
- Performance Monitor
 - using 99
- planning
 - DB2 configuration 15
 - DB2 Connect configuration 15
- pools
 - agent 10
- privileges
 - required 58
- product
 - descriptions 84
 - overview 84
- programs
 - db2chkmig 61

Q

- query
 - cost-based optimizer 4
- Quick Beginnings for OS/2 112
- Quick Beginnings for UNIX 112
- Quick Beginnings for Windows
 - NT 112

R

- Replication Guide and Reference 111
- restrictions
 - instance name 77

S

- scaling the database system
 - definition 9
- Script Center 97
- security
 - for an instance 10
- setting up client communications
 - using the Client Configuration Assistant 40
- setting up document server 117
- shared-nothing configuration
 - definition 3
 - storage 7
- Software Developer's Kit
 - overview 101
- software requirements
 - communication protocols 17
 - DB2 client 17, 18, 19
 - DB2 Connect 17
 - DB2 Software Developer's Kit 17, 18, 19
 - DB2 Universal Database 17
 - Net.Data 18, 19

SQL

- access plans
 - viewing using Visual Explain 100
- SQL Getting Started 111
- SQL Reference 111
- symmetric multiprocessor
 - architecture 9
- SYSADM
 - privileges
 - controlling 58
 - SYSADM_GROUP parameter 58
 - system configuration
 - with DB2 Connect 90
 - with DB2 Universal Database 87
- System Monitor Guide and Reference 111

T

- table
 - in multinode nodegroup 9
 - in single-node nodegroup 9
 - maximum size 7
- TB (terabyte) 7
- terabyte (TB) 7
- Tools Setting 98
- transparent parallelism 5
- Troubleshooting Guide 111

U

- username
 - naming rules 77
- using the command line processor (CLP)
 - cataloging a database 68

V

- verifying the installation
 - server 27
- Visual Explain
 - overview 99, 100

W

- What's New 111

Contacting IBM

This section lists ways you can get more information from IBM.

If you have a technical problem, please take the time to review and carry out the actions suggested by the *Troubleshooting Guide* before contacting DB2 Customer Support. Depending on the nature of your problem or concern, this guide will suggest information you can gather to help us 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.

Telephone

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

- 1-800-237-5511 to learn about available service options.
- 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.

For information on how to contact IBM outside of the United States, see Appendix A of the IBM Software Support Handbook. You can access this document by accessing the following page:

<http://www.ibm.com/support/>

then performing a search using the keyword “handbook”.

Note that in some countries, IBM-authorized dealers should contact their dealer support structure instead of the IBM Support Center.

World Wide Web

<http://www.software.ibm.com/data/>

<http://www.software.ibm.com/data/db2/library/>

The DB2 World Wide Web pages provide current DB2 information about news, product descriptions, education schedules, and more. The DB2 Product and Service Technical Library provides access to frequently asked questions, fixes, books, and up-to-date DB2 technical information. (Note that this information may be in English only.)

Anonymous FTP Sites

<ftp.software.ibm.com>

Log on as anonymous. In the directory /ps/products/db2, you can find demos, fixes, information, and tools concerning DB2 and many related products.

Internet Newsgroups

comp.databases.ibm-db2, bit.listserv.db2-l

These newsgroups are available for users to discuss their experiences with DB2 products.

CompuServe

GO IBMDB2 to access the IBM DB2 Family forums

All DB2 products are supported through these forums.

To find out about the IBM Professional Certification Program for DB2 Universal Database, go to http://www.software.ibm.com/data/db2/db2tech/db2cert.html
--



Part Number: CT6D3NA



Printed in the United States of America
on recycled paper containing 10%
recovered post-consumer fiber.

GC09-2833-00



CT6D3NA

