

IBM DB2 Universal Database
Enterprise – Extended Edition
for UNIX**



Quick Beginnings

Version 6

IBM DB2 Universal Database
Enterprise – Extended Edition
for UNIX**



Quick Beginnings

Version 6

Before using this information and the product it supports, be sure to read the general information under "Appendix H. Notices" on page 185.

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!	vii
Conventions	vii

Part 1. Introduction to DB2 Universal Database 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 10
Instances 11
Fast Communication Manager 13
High Availability 13

Part 2. DB2 Server Installation and Initial Configuration 15

Chapter 2. Planning for Installation 17
Memory Requirements 17
DB2 Universal Database Requirements 17
DB2 Client Requirements 18
Disk Requirements 18
Estimating Fixed Disk Requirements 18
Software Requirements 22
Server Product Requirements 22
Client Product Requirements 24
Possible Client-to-Server Connectivity Scenarios 27

Chapter 3. Installing and Configuring DB2 Universal Database on AIX 29
Before You Begin 29
Performing the Installation 35
Post Installation Steps 36
Step 1. Create a DB2 Instance 36
Step 2. Create the Administration Server 37
Step 3. Update the Node Configuration File 38

Step 4. Enable the Fast Communications Manager 40
Step 5. Modify the Environment to run DB2 Enterprise - Extended Edition 42
Step 6. Enable the Execution of Remote Commands 45
Step 7. Create Links for DB2 Files (Optional) 45
Step 8. Start the Administration Server 46
Step 9. Create the SAMPLE Database (Optional) 48
Step 10. Install the License Key 49

Chapter 4. Installing and Configuring DB2 Universal Database on Solaris. 51
Before you Begin 51
Performing the Installation 54
Post Installation Steps 57
Step 1. Update the Node Configuration File 57
Step 2. Enable the Fast Communications Manager 59
Step 3. Enable the Execution of Remote Commands 61
Step 4. Create Links for DB2 Files (Optional) 61
Step 5. Start the Administration Server 62
Step 6. Create the SAMPLE Database (Optional) 63
Step 7. Install the License Key 64

Part 3. Installing DB2 Clients 65

Chapter 5. Installing DB2 Clients 67

Chapter 6. Installing DB2 Clients on Windows 32-Bit Operating Systems. 69
Before You Begin 69
Installation Steps 69

Chapter 7. Installing DB2 Clients on OS/2 Operating Systems. 73
Before You Begin 73
Installation Steps 74

Chapter 8. Installing DB2 Clients on UNIX Operating Systems.	75
Before You Begin.	75
Installation Steps.	75
Step 1. Update Kernel Configuration Parameters.	75
Step 2. Mount the CD-ROM.	77
Step 3. Perform the Installation.	80

Part 4. Configuring DB2 Communications 85

Chapter 9. Configuring Client-to-Server Communications Using the Client Configuration Assistant	87
Configuration Steps.	88
Adding a Database Using a Profile.	88
Adding a Database Using Discovery	90
Adding a Database Manually	92
Creating and Using Profiles	94
Server Profiles.	94
Client Profiles.	95

Chapter 10. Configuring Client-to-Server Communications Using the Command Line Processor	99
Configuring TCP/IP on the Client	99
Step 1. Identify and Record Parameter Values	100
Step 2. Configure the Client	101
Step 3. Test the Client-to-Server Connection.	106

Part 5. Appendixes 109

Appendix A. Basic Task Knowledge.	111
Starting the Software Registration Tool	111
Starting the Client Configuration Assistant	111
Starting the Control Center	112
Entering Commands Using the Command Center	113
Entering Commands Using the Command Line Processor	114
Command Line Mode	115
Interactive Input Mode.	116
Working with the System Administrative Group	116

Appendix B. Migrating from Previous Versions and Releases	119
Migrating from Previous Versions of DB2	119
Migrating Instances	120
Step 1. Prepare the DB2 Instance for Migration	120
Step 2. Verify that Databases Can Be Migrated	121
Step 3. Migrate the DB2 Instance	128
Migrating Databases	129
Trouble-Shooting Database Migration Errors	129
Optional Post Migration Actions	130

Appendix C. National Language Support (NLS)	135
Language and Codeset Support for UNIX Operating Systems	135
Code Page and Language Support for OS/2 and Windows Operating Environments	135

Appendix D. Naming Rules	137
General Naming Rules	137
Database, Database Alias, and Catalog Node Name Rules	137
Object Name Rules	138
Username, User ID, Group Name, and Instance Name Rules	139
Workstation Name (nname) Rules	139
DB2SYSTEM Naming Rules	141
Password Rules	141

Appendix E. About the DB2 Family of Products	143
Working with DB2 Data	147
Accessing DB2 Data from Remote Clients	148
Accessing Multiple DB2 Servers.	149
Accessing Host or AS/400 DB2 Data from the Desktop Using DB2 Connect Enterprise Edition	150
Accessing DB2 Data from the Web Using Java	154
Accessing DB2 Data from the Web Using Net.Data	155
Accessing DB2 Data from Host and AS/400 Client Machines	157
Administering Instances and Databases with the DB2 Administration Tools.	157
Managing Instances and Database Objects Using the Control Center	158

Managing Communications on the Server	161
Monitoring Databases Using DB2 Performance Monitor	161
Viewing SQL Access Plans Using Visual Explain	162
Managing Connections to Databases Using the Client Configuration Assistant	162
Understanding the Administration Server	163
Developing Applications Using the DB2 Software Developer's Kit	163
Running Your Own Applications	164
Appendix F. How the DB2 Library Is Structured.	165
Completing Tasks with SmartGuides	165
Accessing Online Help	166
DB2 Information – Hardcopy and Online	168
Viewing Online Information	175
Accessing Information with the Information Center	176
Setting Up a Document Server	177
Searching Online Information	178

Printing the PostScript Books.	178
Ordering the Printed Books	179

Appendix G. Removing DB2 Products	181
Step 1. Stop the Administration Server	181
Step 2. Stop All DB2 Instances	181
Step 3. Remove the Administration Server	182
Step 4. Remove DB2 Instances	183
Step 5. Remove DB2 Products	184
Remove DB2 Products on AIX Systems	184
Remove DB2 Products on Linux Systems	184
Remove DB2 Products on Solaris Systems	184

Appendix H. Notices	185
Trademarks	186
Trademarks of Other Companies	186

Index	189
------------------------	------------

Contacting IBM	193
---------------------------------	------------

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. 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 165.



- 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 Universal Database

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 10.

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. By default, every database partition server in your partitioned database system 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.

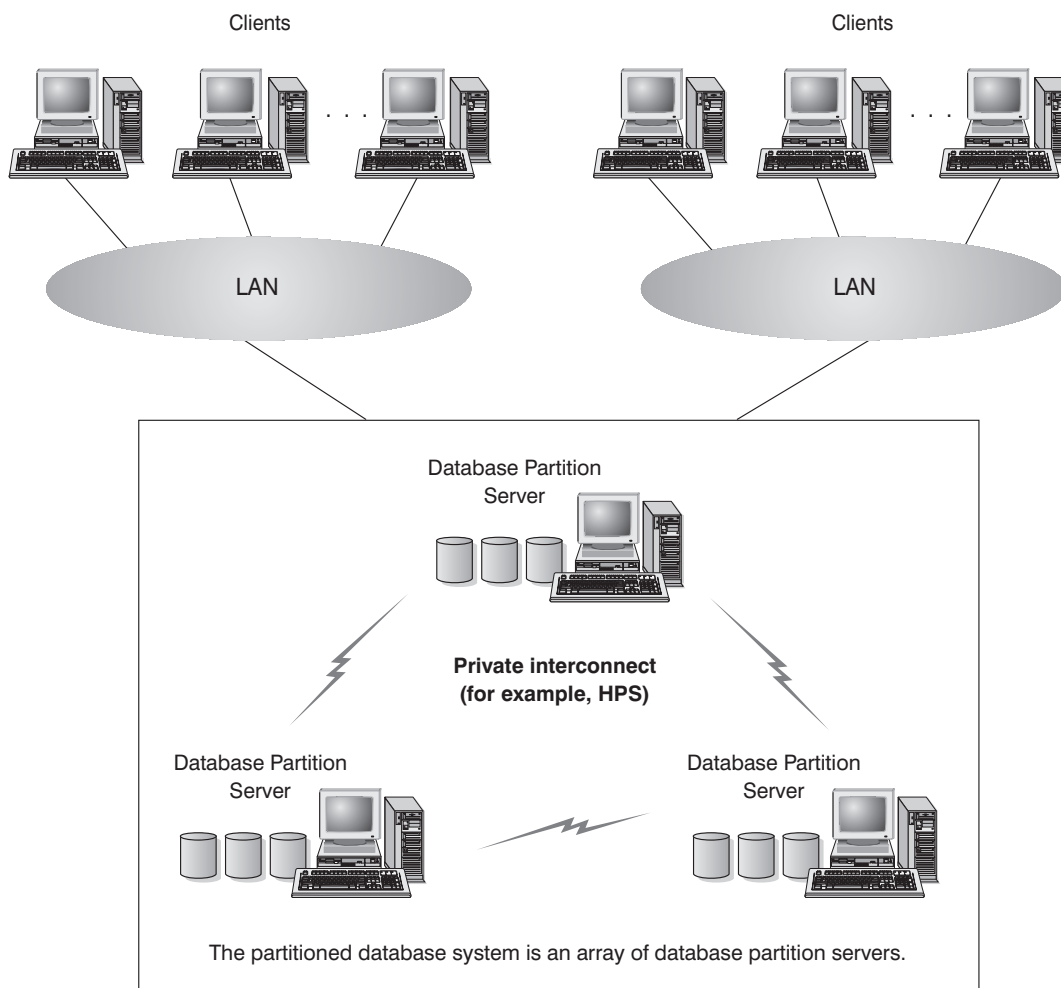


Figure 1. DB2 Enterprise - Extended Edition Hardware Configuration

DB2 EEE can run on a cluster of individual CPUs interconnected by shared memory (symmetric multiprocessors (SMP)), a dedicated high-speed communications switch (for example, HPS), or a LAN. The number of database partition servers in a configuration varies by platform. You should limit the number of database partition servers that communicate over a LAN to 16.

In practice, the number of database partition servers in a configuration is limited by the platform and the management tools available on each platform. For more information about configuration, refer to the *Administration Guide*.

For example, in an IBM RISC/6000 Scalable POWERparallel Systems (RISC/6000 SP) environment that is running AIX, the number of database partition servers is only limited by the possible size of an AIX RISC/6000 SP system.

In a Solaris environment, the number of database partition servers is limited by the size of the machines and the number of those machines that are clustered together. Forty database partition servers could be run on a clustered system of four Ultra Enterprise 6000s with ten CPUs each.

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

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) that is defined in the node configuration file (`db2nodes.cfg`).

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

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

Figure 2 on page 9 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.

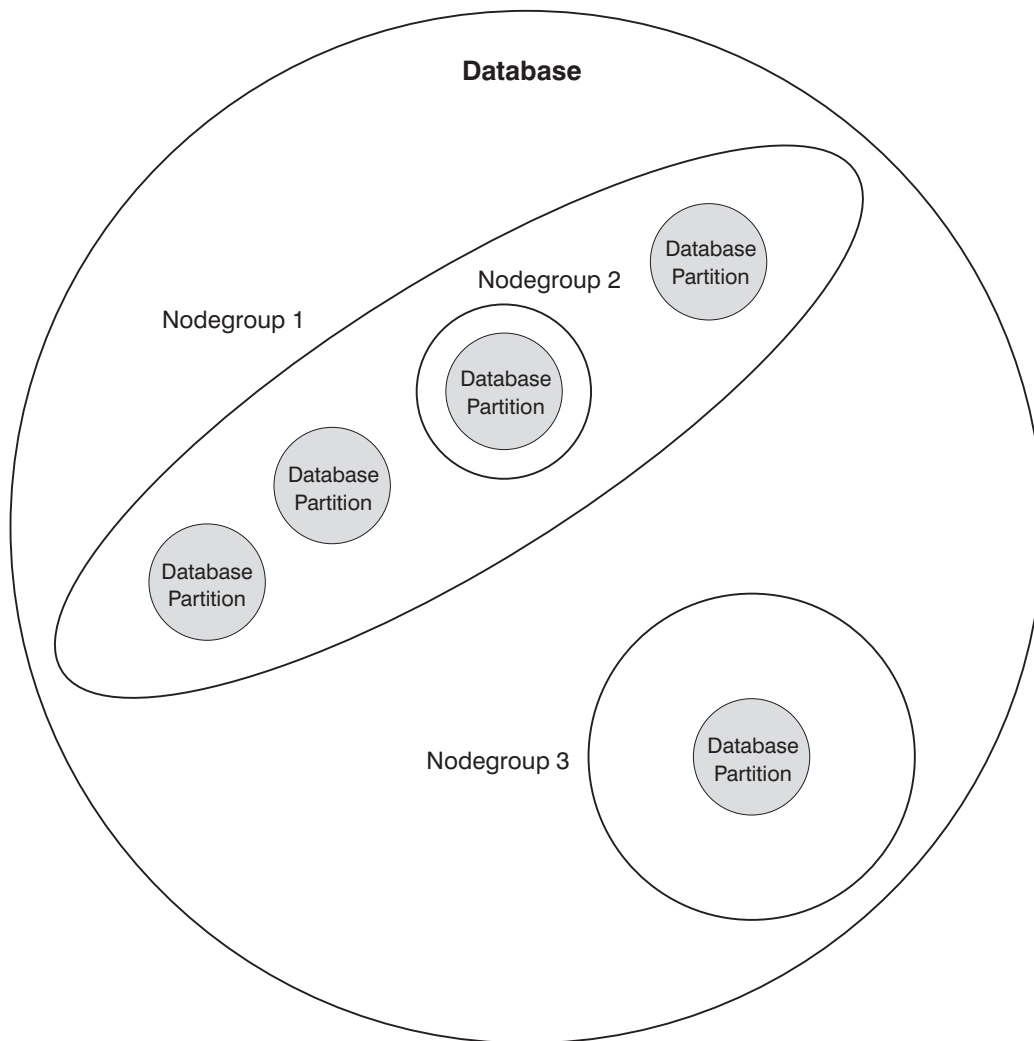


Figure 2. Nodegroups in a Database

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 new nodes have been defined in the `db2nodes.cfg` file, you can add them to a nodegroup in a database. 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.

For more information about setting up logical nodes, see “Step 3. Update the Node Configuration File” on page 38 (for AIX) or Step 1. Update the Node Configuration File (for Solaris), or refer to the *Administration Guide*.

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 (DB2 EEE), an instance is made up of all the database partition servers (nodes) that were defined to take part in a given partitioned database system. The database partition servers are defined in the `db2nodes.cfg` file as *nodes*.

Each instance has different security from other instances on the same machine. This is shown in Figure 3 on page 12, 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.

The `db2nodes.cfg` file for Instance 1 will not list the database partition servers that belong to Instance 2, and the converse.

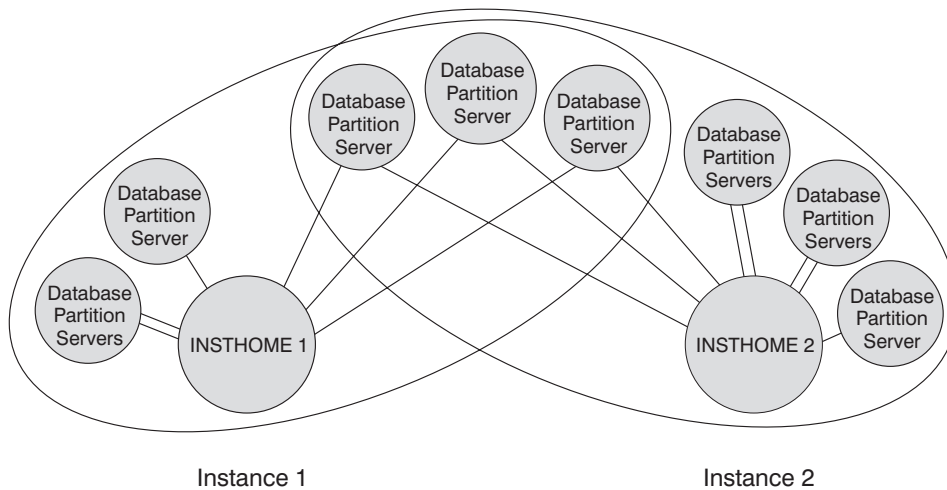


Figure 3. Two Instances

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

- To have distinct test and production environments
- To use different versions of DB2; for example, one instance can use DB2 EEE Version 5.x and another can use DB2 EEE Version 6
- To restrict access to specific databases
- To exploit different database configurations.

Each instance is owned by a user known as the *instance owner*. For information about creating instances, refer to the *Administration Guide*.

The instance owner has *SYSADM* authority for all databases that belong to the instance. Because the instance owner has almost complete control over the instance, this user ID can:

- Issue the **db2start** and **db2stop** commands on the instance
- Modify database configuration parameters
- Modify database manager configuration parameters
- Grant privileges to and revoke them from other users

The instance owner cannot remove an instance. This requires *root* authority.

A one-to-one correspondence exists between an instance and an instance owner; that is, a user cannot own more than one instance. (However, an

instance owner may possess authorizations for other instances, up to, and including, *SYSADM*). In addition to this, each instance must have a separate home directory.

Fast Communication Manager

The Fast Communication Manager (FCM) provides communication support for DB2 Enterprise - Extended Edition. Each database partition server has one FCM daemon to provide communications between database partition servers to handle agent requests, and to deliver message buffers. It consists of:

- A communications process, known as the FCM daemon (*db2fcmdm*)
- Requester functions that run within database manager processes
- Initialization and termination functions

The FCM daemon is started when you start the instance. When the daemon starts, it reads the node configuration file (*INSTHOME/sqllib/db2nodes.cfg*, where *INSTHOME* is the home directory of the instance owner) and defines a well-known address to use in communications.

If communications fail between database partition servers or if they re-establish communications, the FCM daemon 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 specify the number of FCM message buffers with the *fcm_num_buffers* database manager configuration parameter. For a description of this and other FCM parameters, refer to the *Administration Guide*.

High Availability

You can set up your partitioned database system so that if a machine fails, the database server on the failed machine can run on another machine.

On AIX, you implement failover support using IBM's High Availability Cluster Multi-Processing (HACMP). Failover capability allows for the automatic transfer of workload from one processor to another should there be a hardware failure. HACMP provides increased availability through a cluster of processors which share resources such as disks or network access.

On Solaris, you implement failover support using Sun Cluster 2.1. Sun Cluster 2.1 performs both failure detection and the restarting of resources in a clustered environment, as well as failover support for physical disks and IP addresses.

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

Part 2. DB2 Server Installation and Initial Configuration

Chapter 2. Planning for Installation



If you know that your system meets all the hardware and software requirements, and you want to begin installing your DB2 product right away, go to “Chapter 3. Installing and Configuring DB2 Universal Database on AIX” on page 29 or “Chapter 4. Installing and Configuring DB2 Universal Database on Solaris” on page 51.

For information on the DB2 family of products, see “Appendix E. About the DB2 Family of Products” on page 143.

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 on page 18 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	256 MB
___ 10 Concurrent Connections	384 MB
___ 25 Concurrent Connections	512 MB
___ 50 Concurrent Connections	1024 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 on page 19 to estimate the *approximate* 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 AIX	
DB2 Universal Database	55 MB
DB2 GUI Tools	20 MB
DB2 Connect Support	1 MB
Online documentation in HTML format (English)	60 MB
Far-East Code Page Conversion Support	5 MB
Total Disk Space Required	___ MB
DB2 Universal Database Enterprise - Extended Edition for Solaris	
DB2 Universal Database	50 MB
DB2 GUI Tools	20 MB
DB2 Connect Support	1 MB
Online documentation in HTML format (English)	60 MB
Far-East Code Page Conversion Support	5 MB
Total Disk Space Required	___ MB
	The <i>online documentation in HTML format</i> component will install the DB2 documentation in a compressed format. You may need extra disk space temporarily for decompression. Once you have finished the installation, remove the fileset to reclaim the disk space.

Client Components

Use Table 3 on page 20 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 <i>Installation and Configuration Supplement</i> 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)	
AIX	
DB2 Run-Time Client (Base)	25 MB
- Client Configuration Assistant	5 MB
or	
DB2 Administration Client (Base)	25 MB
- DB2 Administration Tools	40 MB
- Client Configuration Assistant	6 MB
Total Disk Space Required	__ MB
HP-UX	
DB2 Run-Time Client (Base)	25 MB
- Client Configuration Assistant	5 MB
or	
DB2 Administration Client (Base)	25 MB
- DB2 Administration Tools	40 MB
- Client Configuration Assistant	6 MB
Total Disk Space Required	__ MB
Linux	
DB2 Run-Time Client (Base)	25 MB
- Client Configuration Assistant	5 MB
or	
DB2 Administration Client (Base)	25 MB
- DB2 Administration Tools	40 MB
- Client Configuration Assistant	6 MB
Total Disk Space Required	__ MB

Table 3. Disk Requirements for Client Components (continued)

Recommended Minimum Disk (MB)	
OS/2	
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
Silicon Graphics IRIX	
DB2 Run-Time Client (Base)	25 MB
- Client Configuration Assistant	6 MB
or	
DB2 Administration Client (Base)	25 MB
- DB2 Administration Tools	40 MB
- Client Configuration Assistant	6 MB
Total Disk Space Required	__ MB
Solaris	
DB2 Run-Time Client (Base)	25 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

Table 3. Disk Requirements for Client Components (continued)

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

Table 4. Software Requirements

Product	Hardware/Software Requirements	Communications
AIX		
• DB2 Universal Database	RISC System/6000 and the following: <ul style="list-style-type: none"> • AIX Version 4.2 or later 	<p>APPC, IPX/SPX, TCP/IP and MPTN (APPC over TCP/IP)</p> <ul style="list-style-type: none"> • For TCP/IP connectivity, no additional software is required. • IPX/SPX connectivity is provided by: <ul style="list-style-type: none"> – AIX base operating system 4.2 or later, which supports DB2 direct addressing. – AIX base operating system 4.3 or later (contains Novell Netware Services for AIX Version 4.1), which supports DB2 direct and file server addressing. • For SNA (APPC) connectivity, one of the following communication products is required: <ul style="list-style-type: none"> – IBM eNetwork Communications Server for AIX V5.0.2.4 – Bull DPX/20 SNA/20 <p>Notes:</p> <ol style="list-style-type: none"> 1. If you want to install the DRDA Application Server function of DB2, you must install the Syncpoint Manager (SPM) and its prerequisites in order to use two-phase commit. 2. If you plan to use DCE (Distributed Computing Environment) with Version 6 of the DB2 Universal Database products, you require a DCE product that is at OSF DCE level 1.1. For DB2 Connect support, you require DB2/MVS Version 5.1 plus its prerequisite, OS/390 DCE Base Services Version 3 for DCE support. With DB2 Connect, you must install DCE Directory Services on the client and the DRDA server. You do not need DCE installed on a DB2 Connect Enterprise Edition server. 3. 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. 4. If you plan to use the Simple Network Management Protocol (SNMP) subagent, you require DPI 2.0 provided by IBM Netfinity Agent. 5. 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.

Table 4. Software Requirements (continued)

Product	Hardware/Software Requirements	Communications
Solaris		
• DB2 Universal Database	<p>Solaris SPARC-based computer and the following:</p> <ul style="list-style-type: none"> • Solaris Version 2.5.1 or later <p>The following patches are required for Solaris version 2.5.1:</p> <ul style="list-style-type: none"> • 101242 Rev. 11 or higher • 103566 Rev. 08 or higher • 103600 Rev. 13 or higher • 103640 Rev. 20 or higher <p>The following patches are required for Solaris version 2.6:</p> <ul style="list-style-type: none"> • 105568 Rev. 12 or higher • 105210 Rev. 13 or higher • 105181 Rev. 06 or higher 	<p>APPC, IPX/SPX, or TCP/IP</p> <ul style="list-style-type: none"> • TCP/IP is provided with the Solaris base operating system. • IPX/SPX connectivity is provided with SolarNet PC Protocol Services 1.1 with IPX/SPX for the Solaris 2.x Operating Environment. • For APPC connectivity, you require SunLink SNA 9.0 or later, and the following communication products: <ul style="list-style-type: none"> – SunLink P2P LU6.2 9.0 or later – SunLink PU2.1 9.0 or later – SunLink P2P CPI-C 9.0 or later <p>Notes:</p> <ol style="list-style-type: none"> 1. DB2 Connect for Solaris requires Solaris Version 2.6 or later. 2. If you plan to use DCE (Distributed Computing Environment) with Version 6 of the DB2 Universal Database products, you require Transarc DCE Version 2.0 for Solaris 2.6 or higher. With DB2 Connect, you must install DCE Directory Services on the client and the DRDA server. You do not need DCE installed on a DB2 Connect Enterprise Edition server. 3. 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.

Client Product Requirements

Table 5 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 AIX • DB2 Administration Client for AIX • DB2 Software Developer's Kit for AIX 	<p>RISC System/6000 and the following:</p> <ul style="list-style-type: none"> • AIX Version 4.2 or later 	<p>APPC or TCP/IP</p> <ul style="list-style-type: none"> • For APPC connectivity, you require IBM eNetwork Communications Server Version 5.0.2.4 or later for AIX • The AIX base operating system provides TCP/IP connectivity, if selected during install. <p>Note: If you plan to use DCE (Distributed Computing Environment), you require a DCE product that is provided by the AIX Version 5 operating system with its latest DCE PTF.</p>

Table 5. Software Requirements for Clients (continued)

Component	Hardware/Software Requirements	Communications
<ul style="list-style-type: none"> DB2 Run-Time Client for HP-UX 10.20 DB2 Administration Client for HP-UX 10.20 DB2 Software Developer's Kit for HP-UX 	<p>HP 9000 Series 700 or 800 system and the following:</p> <ul style="list-style-type: none"> HP-UX Version 10.20 or later <p>The following patches are required:</p> <ul style="list-style-type: none"> For HP-UX Version 10.20 <ul style="list-style-type: none"> PHSS_10556 PHSS_10436 PHSS_10053 PHSS_10113 For systems with the ANSI C or C++ compilers: <ul style="list-style-type: none"> PHSS_10261 PHSS_7505 PHSS_9096 for C++ 	<p>APPC or TCP/IP</p> <ul style="list-style-type: none"> For APPC connectivity, HP-UX Version 10.20 requires the following: <ul style="list-style-type: none"> SNAPLUS2 and the following components: <ul style="list-style-type: none"> SNAPLUS2 Link Version A.10.10 SNAPLUS2 API Version A.10.10 <p>Note: If you plan to use DCE (Distributed Computing Environment) with Version 6 of the DB2 Universal Database products, you require a DCE product that is at OSF DCE level 1.1, which is provided by the HP-UX Version 11 and later operating system.</p> <p>With DB2 Connect, you must install DCE Directory Services on the client and the DRDA server. You do not need DCE installed on a DB2 Connect Enterprise Edition server.</p>
<ul style="list-style-type: none"> DB2 Run-Time Client for HP-UX for 11.00 DB2 Administration Client for HP-UX for 11.00 DB2 Software Developer's Kit for HP-UX 	<p>HP 9000 Series 700 or 800 system and the following:</p> <ul style="list-style-type: none"> HP-UX Version 11.00 or later 	<p>APPC or TCP/IP</p> <ul style="list-style-type: none"> For APPC connectivity, you require either of the following: <ul style="list-style-type: none"> SNAPLUS2 Link R6.11.00.00 SNAPLUS2 API R6.11.00.00 <p>Note: If you plan to use DCE (Distributed Computing Environment), you require a DCE product that is provided by the HP-UX Version 11 base operating system.</p>
<ul style="list-style-type: none"> DB2 Run-Time Client for Linux DB2 Administration Client for Linux DB2 Software Developer's Kit for Linux 	<ul style="list-style-type: none"> Linux kernel 2.0.35 or higher; <i>glibc</i> Version 2.0.7 or higher; <i>pdks</i> package (required to run the DB2 command line processor); and <i>libstdc++</i> Version 2.8.0 or higher. <p>To install DB2, you will need <i>rpm</i>.</p>	<p>TCP/IP</p> <ul style="list-style-type: none"> The Linux base operating system provides TCP/IP connectivity, if selected during install. APPC connectivity is not provided in this release.

Table 5. Software Requirements for Clients (continued)

Component	Hardware/Software Requirements	Communications
<ul style="list-style-type: none"> DB2 Run-Time Client for Silicon Graphics IRIX DB2 Administration Client for Silicon Graphics IRIX DB2 Software Developer's Kit for Silicon Graphics IRIX 	<ul style="list-style-type: none"> Silicon Graphics IRIX, Version 6.x, and the following filesets: <ul style="list-style-type: none"> – eoe.sw.oampkg – eoe.sw.svr4net The following patches are required for Versions 6.2 and 6.3: <ul style="list-style-type: none"> – 2791.0 – 3778.0 	<p>TCP/IP</p> <ul style="list-style-type: none"> The Silicon Graphics IRIX base operating system provides TCP/IP connectivity.
<ul style="list-style-type: none"> DB2 Run-Time Client for Solaris DB2 Administration Client for Solaris DB2 Software Developer's Kit for Solaris 	<p>Solaris SPARC-based computer and the following:</p> <ul style="list-style-type: none"> Solaris Version 2.5.1 or later <p>The following patches are required for Solaris version 2.5.1:</p> <ul style="list-style-type: none"> • 101242 Rev. 11 or higher • 103566 Rev. 08 or higher • 103600 Rev. 13 or higher • 103640 Rev. 20 or higher <p>The following patches are required for Solaris version 2.6:</p> <ul style="list-style-type: none"> • 105568 Rev. 12 or higher • 105210 Rev. 13 or higher • 105181 Rev. 06 or higher 	<p>APPC or TCP/IP</p> <ul style="list-style-type: none"> For APPC connectivity, you require SunLink SNA 9.0 or later and the following: <ul style="list-style-type: none"> – SunLink P2P LU6.2 9.0 or later – SunLink PU2.1 9.0 or later – SunLink P2P CPI-C 9.0 or later The Solaris base operating system provides TCP/IP connectivity. <p>Note: If you plan to use DCE (Distributed Computing Environment), you require a DCE product with Transarc DCE Version 1.1 for Solaris 2.5 and 2.5.1, patch level 18 or higher.</p>
<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>.

Table 5. Software Requirements for Clients (continued)

Component	Hardware/Software Requirements	Communications
<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 and Configuring DB2 Universal Database on AIX



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. See “Appendix B. Migrating from Previous Versions and Releases” on page 119 for information.

This section describes how to install and configure an AIX-based partitioned database system. If you want to install a DB2 Run-Time Client or a DB2 Administration Client, go to “Chapter 5. Installing DB2 Clients” on page 67. For information on how to deploy this product using a distributed installation or other operating-system installation methods, refer to the *Installation and Configuration Supplement*.

Before You Begin



The instructions in this section assume that you are installing DB2 Enterprise - Extended Edition (DB2 EEE) for AIX in an RS/6000 SP environment.

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 17.
- ___ 2. The hostname of the Control workstation and all the workstations that will participate in your partitioned database system. To resolve a workstation’s hostname, log on to the workstation and enter the following command:

hostname

Make note of this output for each workstation.



If you are planning to use every RS/6000 SP workstation in your partitioned database system, skip this item and proceed to the next requirement.

- ___ 3. Have a file that lists the hostnames for all of the RS/6000 SP workstations that will participate in your partitioned database system. Set up the working collective to distribute commands to this list of workstations by performing the following steps:

- Step a. Log on to the Control workstation as a user with root authority
- Step b. Create a file called `eeelist.txt` that will list the *hostnames* for all of the workstations that you want to install this product on.



To resolve a workstation's hostname, enter the **hostname** command.

For example, assume that you wanted to install this product on two SP nodes called `workstation1` and `workstation2`. The contents of this file would be:

```
workstation1
workstation2
```

- Step c. Update the working collective. To update this list, enter the following command:

```
export WCOLL=/eeelist.txt
```

where *eeelist.txt* is the name of the file that you created that lists the RS/6000 SP workstations where you want to install DB2 EEE.

- Step d. Verify that the names in the working collective are indeed the workstations that you want to install this product on by entering the following command:

```
dsh -q
```

You will receive output similar to the following:

```
Working collective file /eeelist.txt:
workstation1
workstation2
Fanout: 64
```

- 4. Verify that Network File System (NFS) is running on each workstation that will participate in your partitioned database system by entering the following command:

```
lssrc -g nfs
```

The Status field for each process should indicate active.

- 5. Ensure that the `rpc.statd` and `rpc.lockd` processes are active on each workstation that will participate in your partitioned database system by entering the following command:

```
ps -ef | grep rpc
```

- 6. Have a file system that is available to all of the RS/6000 SP workstations that will participate in your partitioned database system. This file system will be the home file system for the instance owner and

the Administration Server. To create, NFS export, and NFS mount this file system, perform the following steps:

- Step a. Log on to a workstation that will participate in your partitioned database system as a user with root authority. Typically, this would be a dedicated RS/6000 SP workstation where the home directory of your DB2 users will reside.
- Step b. Create a home file system for your partitioned database system called /home2, by performing the following steps:



In our examples, we will assume that you created a file system called home2 and its mount point is /home2.

- 1) Enter the **smit jfs** command.
- 2) Click on the **Add a Journaled File System** icon.
- 3) Click on the **Add a Standard Journaled File System** icon.
- 4) Select a volume group from the **Volume Group Name** list where you want this file system to physically reside.
- 5) Set the **SIZE of file system (in 512-byte blocks) (Num.)** field to 180 000 (this is about 90 MB).
- 6) Enter a mount point for this file system in the **MOUNT POINT** field. For example, /home2.
- 7) Set the **Mount AUTOMATICALLY at system restart** field to yes.
The remaining fields can be left to the default settings. For more information, refer to your *AIX Administration Guide*.
- 8) Click on **OK**.

- Step c. Mount the file system that you just created by entering the following command:

```
mount /home2
```



You should not mount this file system using the **AMD** or **automounter** utilities. These utilities can cause NFS mounting or locking problems in a partitioned database system.

- Step d. NFS export the /home2 file system so that it is available to all of the RS/6000 SP workstations that will participate in your partitioned database system by performing the following steps:
- 1) Enter the **smit nfs** command.
 - 2) Click on the **Network File System (NFS)** icon.
 - 3) Click on the **Add a Directory to Exports List** icon.
 - 4) Enter the pathname and directory to export (for example, /home2) in the **PATHNAME of directory to export** field.

- 5) Enter the name of each workstation that will participate in your partitioned database system in the **HOSTS allowed root access** field. You recorded these *hostnames* in “Before You Begin” on page 29.



If you are using a high speed interconnect, we recommend that you specify the high speed interconnect names for each workstation in this field as well.

The remaining fields can be left to the default settings. For more information, refer to your *AIX Administration Guide*.

- 6) Click on **OK**.

Step e. Log out.

Step f. Log on to *each* workstation that will participate in your partitioned database system and NFS mount the file system that you exported by performing the following steps:

- 1) Enter the **smit nfs** command.
- 2) Click on the **Network File System (NFS)** icon.
- 3) Click on the **Add a File System for Mounting** icon.
- 4) Enter the pathname of the mount point in the **PATHNAME of the mount point (Path)** field.

The path name of the mount point is where you should create the home directory for the instance owner and the Administration Server. For example, /home2.

- 5) Enter the pathname of the remote directory in the **PATHNAME of the remote directory** field.

For our example, you should enter the same value that you entered in the **PATHNAME of the mount point (Path)** field.

- 6) Enter the *hostname* of the machine where you exported the file system in the **HOST where the remote directory resides** field.

This is the hostname of the machine where the file system that you are mounting was created.



To improve performance, you may want to NFS mount the file system that you created over a high speed interconnect. If you want to mount this file system using a high speed interconnect, you must enter its name in the **HOST where remote directory resides** field.

You should be aware that if the high speed interconnect ever becomes unavailable for some reason, every workstation that participates in your partitioned database system will lose access to these user's home directory.

- 7) Set the **MOUNT now, add entry to /etc/filesystems or both?** field to both.
- 8) Set the **/etc/filesystems entry will mount the directory on system RESTART** field to yes.
- 9) Set the **MODE for this NFS file system** field to read-write.
- 10) Set the **Mount file system soft or hard** field to soft.
 A soft mount means that the machine *will not* try for an infinite period of time to remotely mount the directory. A hard mount means that your machine will infinitely try to mount the directory, this could cause problems in the event of a system crash. We recommend that you set this field to soft.
 The remaining fields can be left to the default settings. For more information, refer to your *AIX Administration Guide*.
- 11) Ensure that this file system is mounted with the **Allow execution of SUID and sgid programs in this file system?** field set to Yes. This is the default setting.
- 12) Click on **OK**.

Step g. Log out.

— 7. Create three separate groups and user accounts for the:

- DB2 instance owner
- user that will execute fenced UDFs (user defined functions) or stored procedures
- Administration Server.

The usernames you create must conform to both your operating system's naming rules, and those of DB2. For more information on naming rules, see "Appendix D. Naming Rules" on page 137.

To create all three of these users, perform the following steps:

Step a. Log on to the Control workstation as a user with root authority.



If you are using NIS and file collections is enabled, you should set up your system to exclude the propagation of any user administration files. For more information, refer to your NIS documentation.

Step b. Create a group for the instance owner (for example, db2iadm1), the user that will execute UDFs or stored procedures (for example, db2fadm1), and the Administration Server (for example, db2asgrp) by entering the following commands:

```
mkgroup id=999 db2iadm1
mkgroup id=998 db2fadm1
mkgroup id=997 db2asgrp
```

Step c. Create a user that belongs to each group that you created in the previous step by entering the following commands:

```
mkuser id=1004 pgrp=db2iadm1 groups=db2iadm1
      home=/home2/db2inst1 db2inst1
mkuser id=1003 pgrp=db2fadm1 groups=db2fadm1
      home=/home2/db2fenc1 db2fenc1
mkuser id=1002 pgrp=db2asgrp groups=db2asgrp
      home=/home2/db2as db2as
```

Step d. Set an initial password for each user that you created by entering the following commands:

```
passwd db2inst1
passwd db2fenc1
passwd db2as
```

Step e. Log out.

Step f. Log on to the Control workstation as each user that you created (db2inst1, db2fenc1, and db2as). You will be prompted to change each user's password since this is the first time that these users have logged onto the system.

Step g. Log out.

Step h. Log on to the Control Workstation as a user with root authority.

Step i. Distribute each group and user account that you created to the other SP workstations that will participate in your partitioned database system by entering the following command:

```
dsh /var/sysman/supper update user.admin
```



If you are installing DB2 EEE on every RS/6000 SP workstation in your environment, you can specify the **-a** flag as follows:

```
dsh -a /var/sysman/supper update user.admin
```

This command will automatically be distributed to each RS/6000 SP workstation that you listed in the `eeelist.txt` file.



If you do not issue the **dsh** command, the *cron* job will automatically create these groups and user accounts when it is automatically run. The default setting is 10 minutes past the hour.

Step j. Log out.

Performing the Installation

To install DB2 Enterprise - Extended Edition (DB2 EEE), perform the following steps:

Step 1. Log on to the Control workstation as a user with root authority.

Step 2. Insert the appropriate CD-ROM into the drive.

Step 3. Create a directory for the CD-ROM. To create a directory, enter the following command:

```
mkdir /cdrom
```

Step 4. Allocate a CD-ROM file system by entering the following command:

```
crfs -v cdrfs -p ro -d cd0
```

Step 5. Mount the CD-ROM file system by entering the following command:

```
mount -v cdrfs -r /dev/cd0 /cdrom
```

Step 6. Start the installation program on each workstation that will participate in your partitioned database system by entering the following command:

```
dsh installp -qagXd /cdrom/db2 db2_06_01.xlic
```



If you are installing DB2 EEE on every RS/6000 SP workstation in your environment, you can specify the **-a** flag as follows:

```
dsh -a installp -qagXd /cdrom/db2 db2_06_01.xlic
```

Step 7. The db2_06_01.xlic fileset will only install those components that are required to run DB2 EEE, this fileset will not install the Control Center and only install the *English* version of the DB2 messages, documentation. To install any non-English DB2 messages, documentation, or the Control Center, perform the following steps:

Step a. Enter the **smit install** command.

Step b. Click on the **Install and Update Software** icon.

Step c. Click on the **Install/Update From All Available Software** icon.

Step d. Enter the CD-ROM directory that you created in the **INPUT device/directory for software** field. For example, /cdrom.

Step e. Click on **OK**.

Step f. Click on the **List** push button beside the * **SOFTWARE to install** field.

Step g. Select the appropriate **DB2 Product Messages**, **DB2 Product Documentation (HTML)**, **DB2 Control Center Help (HTML)**, and the **DB2 Control Center** entries from the list of software to install.



You can use this method to install any other DB2 components that were not installed by default.

Step h. Click on **OK**.

When you have installed the documentation, you will need to decompress and untar the documentation using the **db2insthtml** command. For more information, refer to the *Administration Guide*.

Step 8. Enter the **lslpp -l | grep db2_06_01** command when the installation completes to see the DB2 components that were installed on your system.

Step 9. Log out.

Post Installation Steps

This section describes the remaining steps that you have to perform to get your partitioned database system up and running.

Step 1. Create a DB2 Instance

This section describes how to create an instance using the DB2 Installer program. We recommend that you use the DB2 Installer program to create an instance because it will configure the instance automatically for communications.

For information on how to create an instance using the **db2icrt** command, and configure it for communications, refer to the *Administration Guide* and the *Installation and Configuration Supplement* online document.

To create an instance using the DB2 Installer program, perform the following steps:

- Step 1. Log on as a user with root authority to any RS/6000 SP workstation where you installed DB2 Enterprise - Extended Edition (DB2 EEE).
- Step 2. Enter the **/usr/lpp/db2_06_01/install/db2setup** command to start the DB2 Installer program. The DB2 Installer window opens.
- Step 3. Select the **Create** option and press the **Enter** key.
Press the **Tab** key to change a highlighted option and the **Enter** key to select or deselect the option you want. For more information on using the DB2 Installer program, or for help with any task you are performing with the DB2 Installer, select the **Help** option and press the **Enter** key.
- Step 4. Select the **Create a DB2 Instance** option and press the **Enter** key.

- Step 5. Fill in all of the fields for the username that you created for the instance owner in “Before You Begin” on page 29.
- Step 6. Select the **OK** option and press the **Enter** key.
- Step 7. Fill in all of the fields for the username that you created to execute fenced UDFs and stored procedures in “Before You Begin” on page 29.
- Step 8. Select the **OK** option and press the **Enter** key. You are returned to the Create DB2 Services window.
- Step 9. Select the **OK** option and press the **Enter** key. You will receive a warning message that you have not yet created the Administration Server, you can ignore this message. The Summary Report window opens.
- Step 10. Select the **Continue** option and press the **Enter** key. The DB2 Installer program will prompt you to complete this task. When you have completed this task, you are returned to the DB2 Installer window.
- Step 11. Select the **Close** option and press the **Enter** key.

Step 2. Create the Administration Server

This section describes how to create the Administration Server using the DB2 Installer program. If you are planning to use the Control Center to administer your partitioned database system, you require a running Administration Server.

We recommend that you use the DB2 Installer program to create the Administration Server because it will configure the Administration Server automatically for communications. For information on how to create an instance using the **dasicrt** command, and configure it for communications, refer to the *Administration Guide* and the *Installation and Configuration Supplement* online document.

To create the Administration Server using the DB2 Installer program, perform the following steps:

- Step 1. Log on as a user with root authority to any RS/6000 SP workstation where you installed DB2 Enterprise - Extended Edition (DB2 EEE).
- Step 2. Enter the **/usr/lpp/db2_06_01/install/db2setup** command to start the DB2 Installer program. The DB2 Installer windows opens.
- Step 3. Select the **Create** option and press the **Enter** key.
- Step 4. Select the **Create the Administration Server** option and press the **Enter** key.
- Step 5. Fill in the fields for the username that you created for the Administration Server in “Before You Begin” on page 29.

- Step 6. Select the **OK** option and press the **Enter** key. You are returned to the Create DB2 Services window.
- Step 7. Select the **OK** option and press the **Enter** key. The Summary Report window opens.
- Step 8. Select the **Continue** option and press the **Enter** key. The DB2 Installer program will prompt you to complete this task. When you have completed this task, you are returned to the DB2 Installer window.
- Step 9. Select the **Close** option and press the **Enter** key.

Step 3. Update the Node Configuration File

The node configuration file (`db2nodes.cfg`), located in the `INSTHOME/sqllib/` directory (where `INSTHOME` is the home directory of the instance owner) contains configuration information for all the database partition servers that are participating in the instance. There is a `db2nodes.cfg` file for every multipartition instance. This file contains one entry for each database partition server for a particular DB2 instance. By default, when you create an instance, the `db2nodes.cfg` file is automatically created and an entry for the workstation where the instance was created is added to the file.

The format of the `db2nodes.cfg` file is as follows:

```
nodenum    hostname    logical port    netname
```

where:

- nodenum** A unique number, between 0 and 999, that identifies a database partition server in a partitioned database system.
- When you want to scale your partitioned database system, you add an entry for each database partition server to the `db2nodes.cfg` file. The *nodenum* that you select for additional database partition servers must be in ascending order, however, gaps can exist in this sequence. You may choose to put a gap between the *nodenum* values if you plan to add an MLN and wish to keep them logically grouped together in this file.
- This entry is required.
- hostname** The TCP/IP *hostname* of that database partition server (which you recorded earlier) for use by the FCM.
- This entry is required.



You can specify the name of a high speed interconnect for the *hostname* value. This allows the **db2start**, **db2stop**, and **db2_all** commands to use the high speed interconnect for communications.

- logical port** Specifies the logical port number for the database partition

server. This field is used to specify a particular database partition server on a workstation that is running MLNs. If there is no entry for this field, the default is 0. However, if you add an entry for the *netname* field, you must enter a number for the *logical port* field.

If you are not using MLNs and you specify an entry for the *netname* field, the entry *must* be set to 0.

If you are using MLNs, the *logical port* you specify *must* start at 0 and continue in ascending order with no gaps (for example, 0,1,2).

This field is only optional when you are not using MLNs or a high speed interconnect.

netname Specifies the hostname or the IP address of the high speed interconnect for FCM communications.

If an entry is specified for this field, all communications between database partition servers (except for communications as a result of the **db2start**, **db2stop**, and **db2_all** commands) is handled through the high speed interconnect.

This parameter is only required if you are using a high speed interconnect for database partition communications.

For example, when you created the DB2 instance db2inst1, on the workstation called workstation1, the db2nodes.cfg file was updated as follows:

```
0          workstation1  0
```

If you want your partitioned database system to contain two workstations, called workstation1 and workstation2, you would update the db2nodes.cfg file as follows:

```
0          workstation1  0
1          workstation2  0
```

If you want your partitioned database system to contain two workstations, called workstation1 and workstation2, and workstation1 was running 3 database partition servers, you would update the db2nodes.cfg file as follows:

```
0          workstation1  0
1          workstation1  1
2          workstation1  2
3          workstation2  0
```

If you want your partitioned database system to contain two workstations, called workstation1 and workstation2 (with workstation2 running 2 database

partition servers), and use a high speed interconnect called switch1 and switch2, you would update the db2nodes.cfg file as follows:

0	workstation1	0	switch1
1	workstation2	0	switch2
2	workstation2	1	switch2

To update the db2nodes.cfg file, perform the following steps:

- Step 1. Log on to the instance as a user with system administrative (SYSADM) authority (by default, the username that you created for the instance owner). For more information, see “Working with the System Administrative Group” on page 116.
- Step 2. Ensure the DB2 instance is stopped by entering the *INSTHOME/sqllib/adm/db2stop* command, where *INSTHOME* is the home directory of the instance owner.



The db2nodes.cfg is locked when the instance is running and can only be edited when the instance is stopped.

- Step 3. Edit the db2nodes.cfg file and add an entry for each database partition server that will participate in your partitioned database system.
- Step 4. Log out.

For more information on the db2nodes.cfg file, refer to the *Administration Guide*.

Step 4. Enable the Fast Communications Manager



If you are using Network Information Services (NIS), NIS+, or any other tools to manage system configuration files, you must perform this step on the master server. After you have completed this step, you should ensure that the changes have taken effect on all the RS/6000 SP workstations in your partitioned database system.

This section describes how to enable communications between the database partition servers that participate in your partitioned database system. Communications between database partition servers is handled by the Fast Communications Manager (FCM). To enable FCM, you must ensure that the port or port range that you reserve in the /etc/services file is available on every workstation in your partitioned database system.

When you created a DB2 instance using the DB2 Installer program, an entry similar to the following was placed in the /etc/services file:

```
DB2_instance_name      60000/tcp
```

where *instance_name* is the name of the instance that you created.

This entry *must* be of the form:

```
DB2_instance_name    port_number
```

where:

- *instance_name* is the name of the multipartition instance.
- *port_number* is the port number that you reserve for database partition server communications.



The port 60000 is the default FCM port used by an instance for any instance that was created using the DB2 Installer program.

For example, if you created an instance called `db2inst1`, the entry placed in the `/etc/services` file would be:

```
DB2_db2inst1    60000/tcp
```



If the `/etc/services` file is part of file collections, this entry will be overwritten by the `supper` program.

If you are planning to run multiple logical nodes (MLNs), you must reserve a range of consecutive ports (one for each logical node). The consecutive port range that you reserve must be equal to the number of MLNs on the database partition server that is running the most MLNs in your partitioned database system. To reserve a port range for FCM to use, you will need to add another entry to the `/etc/services` file that specifies the end of the port range. This entry *must* be of the format:

```
DB2_instance_name_END    port_number
```

- *instance_name* is the name of the multipartition instance.
- *port_number* is the port number the last port in the port range that you reserved for database partition server communications.

For example, if you have a database partition server that is running 6 database partitions you will need to specify 6 consecutive ports as your port range. In our example, the entry in the `/etc/services` file would be:

```
DB2_db2inst1    60000/tcp
DB2_db2inst1_END    60005
```



You can add a comment that describes these entries using the `#` comment identifier. For example:

```
DB2_db2inst1    60000/tcp # This is an EEE instance
DB2_db2inst1_END    60005/tcp # that is running 6 MLNs.
```

To enable your partitioned database system for FCM communications, perform the following steps on each machine that will participate in your partitioned database system:

- Step 1. Log on the workstation where you created a DB2 instance as a user with root authority.
- Step 2. Edit the `/etc/services` file and add the port or port range entries that you have reserved for FCM communications.
- Step 3. Log out.
- Step 4. Log on to each additional workstation in your partitioned database system and add the same entries to the `/etc/services` file that you added to the `/etc/services` file on the workstation where you created the DB2 instance.
- Step 5. Log out.

For more information on FCM communications, refer to the *Administration Guide*.



If the `/etc/services` file is part of file collections, you must add these entries to the `/etc/services` file on the Control workstation and update each workstation using the following command:

```
dsh -a /var/sysman/supper update user.admin
```

Step 5. Modify the Environment to run DB2 Enterprise - Extended Edition

This section describes the different environment settings that you need to update on each database partition server that will participate in your partitioned database system before you begin using DB2 EEE. To tune your system's environment to run DB2 EEE, perform the following steps:

- Step 1. Log on to a database partition server as a user with root authority.
- Step 2. Set the AIX `maxuproc` (maximum number of processes per user) device attribute to `1000` on all the workstations that are participating in your partitioned database system by entering the following command:

```
dsh -a chdev -l sys0 -a maxuproc='1000'
```

- Step 3. Set the TCP/IP network parameters on all the workstations that are participating in your partitioned database system to the following values:

```
thewall      = 16384
sb_max       = 1310720
rfc1323      = 1
tcp_sendspace = 655360
tcp_recvspace = 655360
udp_sendspace = 65536
udp_recvspace = 655360
```



To list the current settings of all network-related parameters, enter the **dsh -a no -a** command.

To set a parameter, enter the follow command:

```
dsh -a no -a parameter_name=value
```

where:

- *parameter_name* is the parameter you want to set.
- *value* is the value that you want to set for this parameter.

For example, to set the `tcp_sendspace` parameter to 655360 on every workstation, enter the following command:

```
dsh -a no -a tcp_sendspace=655360
```



The above values are the minimum values for these parameters. If you are already setting any of these network-related parameters to a higher value, do not set the parameter to the lower listed value.

For a list of all the network parameters that are set on a workstation, enter the **no -a** command.

In addition, if you are using a high speed interconnect, you must set the *spoolsize* and *rpoolsz* for *css0* to the following values:

```
spoolsize      8388608
rpoolsz        8388608
```



To list the current settings of these parameters, enter the **dsh -a lsattr -l css0 -E** command.

To set these parameters, enter the following commands:

```
dsh -a /usr/lpp/ssp/css/chgcss -l css0 -a spoolsize=8388608
dsh -a /usr/lpp/ssp/css/chgcss -l css0 -a rpoolsz=8388608
```

If you are not using the `/tftpboot/tuning.cst` file to tune your system, you can use the `/usr/lpp/db2_06_01/misc/rc.local.sample` sample script file to update the network-related parameters. To update the network-related parameters using the sample script file, perform the following steps:

- a. Copy this script file to the `/etc` directory and make it executable by root by entering the following commands:

```
cp /usr/lpp/db2_v6_01/misc/rc.local.sample /etc/rc.local
chown root:sys /etc/rc.local
chmod 744 /etc/rc.local
```
- b. Review the `/etc/rc.local` file and update it if necessary.
- c. Add an entry to the `/etc/inittab` file so that the `/etc/rc.local` script is executed whenever the machine is

rebooted. You can use the **mkitab** command to add an entry to the /etc/inittab file. To add this entry, enter the following command:

```
mkitab "rclocal:2:wait:/etc/rc.local > /dev/console 2>&1"
```

- d. Ensure that /etc/rc.nfs entry is included in the /etc/inittab file by entering the following command:

```
lsitab rcnfs
```

- e. Update the network parameters without rebooting your system by entering the following command:

```
/etc/rc.local
```

- ___ Step 4. Ensure that you have enough paging space for DB2 EEE to run. If you do not have sufficient paging space to run DB2 EEE, the operating system will kill the process that is using the most virtual memory (this is likely to be one of the DB2 EEE processes). To check for available paging space, enter the following command:

```
lsp -a
```

This command will return output similar to the following:

Page Space	Physical Volume	Volume Group	Size	%Used	Active	Auto	Type
paging00	hdisk1	rootvg	60MB	19	yes	yes	lv
hd6	hdisk0	rootvg	60MB	21	yes	yes	lv
hd6	hdisk2	rootvg	64MB	21	yes	yes	lv

We recommend that the paging space available be equal to twice the amount of physical memory installed on your workstation.

- ___ Step 5. If you are creating a small to intermediate size partitioned database system, the number of network file system daemons (NFSDs) on the machine that has the instance owner's home directory should be close to:

of biod on a machine * # of machines in the instance

We recommended that you run 10 biod processes on every machine. According to the above formula, on a four machine system with 10 biod processes, you would use 40 NFSDs.

If you are installing a larger system, you can have up to 120 NFSDs on the machine.

For additional information about NFS, refer to your NFS documentation.

Step 6. Enable the Execution of Remote Commands

In a multipartition instance, each database partition server must have the authority to perform remote commands on all the other database partition servers in the partitioned database system. To do this, you can either create an *INSTHOME/.rhosts* file (where *INSTHOME* is the home directory of the instance owner) or an */etc/hosts.equiv* file.

If you decide to create an */etc/hosts.equiv* file, you will have to create this file on every workstation that participates in your partitioned database system.

If you want to use the *INSTHOME/.rhosts* file, it should contain entries similar to the following:

```
workstation1.torolab.ibm.com db2inst1
workstation2.torolab.ibm.com db2inst1
switch01.torolab.ibm.com      db2inst1
switch02.torolab.ibm.com      db2inst1
```

For more information about the *INSTHOME/.rhosts* or the */etc/hosts.equiv* file, refer to your operating system's documentation.

Step 7. Create Links for DB2 Files (Optional)

You can use the **db2ln** command to create links for the DB2 files to the */usr/lib* directory, and for the include files to the */usr/include* directory, for a particular version and release level of DB2 Enterprise - Extended Edition (DB2 EEE). Links can be established for only one version of DB2 on a given workstation.

You may want to create these links if you are developing or running applications and want to avoid having to specify the full path to the product libraries and include files.

To create links for the DB2 files, perform the following steps:

Step 1. Log on to the Control workstation as a user with root authority.

Step 2. Enter the following command to create the links for the DB2 files:

```
dsh /usr/lpp/db2_06_01/cfg/db2ln
```

This command will automatically be distributed to each RS/6000 SP workstation that you listed in the *eeelist.txt* file.



If you installed DB2 EEE on every RS/6000 SP workstation in your environment, you can specify the *-a* flag as follows:

```
dsh -a /usr/lpp/db2_06_01/cfg/db2ln
```

Step 3. Log out.

If there are existing links to the `/usr/lib` and `/usr/include` directories from previous versions of DB2 EEE, they will automatically be removed by executing the **db2ln** command to create links for this version of DB2 EEE.

Step 8. Start the Administration Server

This section describes how to start and configure the Administration Server. Before you can use the Client Configuration Assistant (CCA) or the Control Center to administer your partitioned database system, you must ensure the Administration Server is started and the `db2cc1st` daemon is running on every database partition server in your partitioned database system.

When you created the Administration Server using the DB2 Installer program, an entry similar to the following was placed in the `/etc/services` file on the workstation where you created the Administration Server:

```
db2ccmsrv      523/tcp
```



If the `/etc/services` file is part of file collections, this entry will be overwritten by the supper program.

To enable your partitioned database system so that it can be administered by the CCA or the Control Center, perform the following steps:

- Step 1. Log on the workstation where you created the Administration Server as a user with root authority.
- Step 2. Edit the `/etc/services` file add the port that was reserved for the Administration Server, if it does not exist.
- Step 3. Log out.
- Step 4. Log on to each additional workstation in your partitioned database system and add the same entry to the `/etc/services` file that you added to the `/etc/services` file on the workstation where you created the Administration Server.



If the `/etc/services` file is part of file collections, you must add this entry to the `/etc/services` file on the Control workstation and update each workstation using the following command:

```
dsh -a /var/sysman/supper update user.admin
```

- Step 5. Log out.
- Step 6. Log on to any database partition server with the username that you created for the Administration Server (for example, `db2as`).
- Step 7. Enter the following command to start the Administration Server

```
DASINSTHOME/sql1lib/bin/db2admin start
```

where *DASINSTHOME* is the home directory of the Administration Server.

- Step 8. Log out.
- Step 9. Log on to each database partition server that participates in your partitioned database system with the username that you created for the Administration Server (for example, db2as).
- Step 10. Enter the following command to start the db2cc1st daemon
- ```
DASINSTHOME/sql1lib/bin/db2cc1st
```

where *DASINSTHOME* is the home directory of the Administration Server.



Instead of logging on to each database partition server, you can run the **db2cc1st** command on all database partition servers by entering the following commands on any database partition server:

```
export RAHOSTFILE=INSTHOME/sql1lib/db2nodes.cfg
db2_all "DASINSTHOME/sql1lib/bin/db2cc1st"
```

where:

- *INSTHOME* is the home directory of the instance owner.
- *DASINSTHOME* is the home directory of the Administration Server.

- Step 11. Ensure that the db2cc1st daemon is running by entering the following command:

```
ps -ef | grep db2cc1st
```

Error information for the db2cc1st daemon is returned to the syslog file. Errors are written to this file, instead of the db2diag.log file, because the daemon runs independently from the instance.



Instead of logging on to each database partition server, you can run the **ps -ef | grep db2cc1st** command on all database partition servers by entering the following command on any database partition server:

```
db2_all ps -ef | grep db2cc1st
```

- Step 12. Log out.



You can set up the db2cc1st daemon to start when a workstation is restarted by adding the follow command to the */etc/inittab* file.

```
mkitab "db2cc1st:2:once:DASINSTHOME/sql1lib/bin/db2cc1st"
```

where *DASINSTHOME* is the home directory of the Administration Server.

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

## Step 9. Create the SAMPLE Database (Optional)

This section describes how to verify that DB2 Enterprise - Extended Edition (DB2 EEE) is installed and configured correctly by creating the SAMPLE database on your system, and accessing data from the database.



To complete the steps in this section, we recommend that you create a separate file system (for example, /database) on *each* database partition server that participates in your partitioned database system.

For more information on how to create a file system, refer to your operating system's documentation.

To verify that you have installed and configured DB2 EEE, perform the following steps:

Step 1. Log on to the system as a user with system administrative (SYSADM) authority (by default, the username that you created for the instance owner). For more information, see “Working with the System Administrative Group” on page 116.

Step 2. Enter the following command to create the SAMPLE database:

```
INSTHOME/sql1lib/bin/db2samp1 path
```

where:

- *INSTHOME* is the home directory of the instance owner.
- *path* is a *local* file system that exists on every workstation that participates in your partitioned database system.

For example, to create the SAMPLE database on the /database file system, enter the following command:

```
INSTHOME/sql1lib/bin/db2samp1 /database
```

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:

```
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 113 or “Entering Commands Using the Command Line Processor” on page 114.



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.

## Step 10. Install the License Key

This section describes how to install the license key for the DB2 product that you installed.



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

1. Log on to the Control workstation as a user with root authority.
2. The DB2 product license key can be found in the `nodelock` file, located in the `/var/ibm` directory.
3. Update your DB2 product license by entering the following command:

```
/usr/lpp/db2_06_01/cfg/db2licm license_filename
```

where `license_filename` is the full pathname and filename for the license file that corresponds to the product you have purchased.

The name of the license file for this product is `db2udbeee.lic`.

For example, if the CD-ROM is mounted in the `/cdrom` directory and the name of the license file is `db2udbeee.lic`, you would enter the following command:

```
/usr/lpp/db2_06_01/cfg/db2licm /cdrom/db2/license/db2udbeee.lic
```



You can also use the Control Center to manage and monitor your DB2 licensing agreement and to monitor database usage. To manage licensing using the Control Center, perform the following steps:

- Step 1. Start the Control Center. For more information, see “Starting the Control Center” on page 112.
- 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.



---

## Chapter 4. Installing and Configuring DB2 Universal Database on Solaris

This section describes how to install and configure a Solaris-based partitioned database system. If you want to install a DB2 Run-Time Client or a DB2 Administration Client, go to “Chapter 5. Installing DB2 Clients” on page 67. For information on how to deploy this product using a distributed installation or other operating-system installation methods, refer to the *Installation and Configuration Supplement*.

We assume in these instructions that you install and configure DB2 Enterprise - Extended Edition (DB2 EEE) using the DB2 Installer program. We also assume that you select to install the Control Center, create an instance, and create the Administration Server using the DB2 Installer program.

If you want to install your DB2 product using your UNIX operating system's native installation tools, refer to the *Installation and Configuration Supplement*.



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. See “Appendix B. Migrating from Previous Versions and Releases” on page 119 for information.

---

### Before you Begin



The instructions in this section assume that you are installing DB2 Enterprise - Extended Edition (DB2 EEE) for Solaris on one workstation and running Multiple Logical Nodes (MLNs) to create your partitioned database system.

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 17.
- \_\_\_ 2. Have a file system that is available to all of the systems that will participate in your partitioned database system. This file system will be the home file system for the instance owner and the Administration Server. To create a file system, perform the following steps:

- Step a. Select a disk partition or logical volume and use a utility like `newfs` to create this file system. For more information, enter the **man newfs** command.
- Step b. Mount this file system locally and add an entry to the `/etc/ufstab` file so that the this file system is mounted each time the system is rebooted.



If you are installing DB2 EEE on a cluster of Solaris systems, you must add an entry to the `/etc/dfs/dfstab` file to export this file system via NFS. After you have exported this file system, you must mount this file system on each system in the DB2 EEE cluster.

- \_\_\_ 3. Create three separate groups and user accounts for the:
- DB2 instance owner
  - user that will execute fenced UDFs (user defined functions) or stored procedures
  - Administration Server.

The usernames you create must conform to both your operating system's naming rules, and those of DB2. For more information on naming rules, see "Appendix D. Naming Rules" on page 137.



You can have the DB2 Installer create these usernames, or you can create them manually.

To create all three of these users, perform the following steps:

- Step a. Log on to the system as a user with root authority.
- Step b. Create a group for the instance owner (for example, `db2iadm1`), the user that will execute UDFs or stored procedures (for example, `db2fadm1`), and the Administration Server (for example, `db2asgrp`) by entering the following commands:
- ```
groupadd -g 999 db2iadm1
groupadd -g 998 db2fadm1
groupadd -g 997 db2asgrp
```
- Step c. Create a user that belongs to each group that you created in the previous step by entering the following commands:
- ```
useradd -g db2iadm1 -u 1004 -d /home2/db2inst1 -m db2inst1
useradd -g db2fadm1 -u 1003 -d /home2/db2fenc1 -m db2fenc1
useradd -g db2asgrp -u 1002 -d /home2/db2as -m db2as
```
- Step d. Set an initial password for each user that you created by entering the following commands:
- ```
passwd db2inst1
passwd db2fenc1
passwd db2as
```
- Step e. Log out.

- ___ 4. Update the Kernel Configuration Parameters. In order to run DB2 EEE you should update your kernel configuration parameter by performing the following steps:

Step a. Log on to the system as a user with root authority.

Step b. Update the kernel configuration parameters as follows:

Table 7. Solaris Kernel Configuration Parameters (Recommended Values)

Kernel Parameter	Physical Memory			
	64MB - 128MB	128MB - 256MB	256MB - 512MB	512MB+
msgsys:msginfo_msgmax	65 535(1)	65535(1)	65 535(1)	65 535(1)
msgsys:msginfo_msgmnb	65 535(1)	65 535(1)	65 535(1)	65 535(1)
msgsys:msginfo_msgmap	130	258	258	258
msgsys:msginfo_msgmni	128	256	256	256
msgsys:msginfo_msgssz	16	16	16	16
msgsys:msginfo_msgtql	256	512	1024	1024
msgsys:msginfo_msgseg	8 192	16 384	32 768	32 768
shmsys:shminfo_shmmax	67 108 864	134 217 728(2)	268 435 456(2)	536 870 912(2)
shmsys:shminfo_shmseg	16	16	16	16
shmsys:shminfo_shmmni	300	300	300	300
semsys:seminfo_semmni	128	256	512	1024
semsys:seminfo_semmap	130	258	514	1026
semsys:seminfo_semmns	256	512	1024	2048
semsys:seminfo_semnu	256	512	1024	2048

Notes:

- 1) The *msgsys:msginfo_msgmnb* and *msgsys:msginfo_msgmax* parameters must be set to 65535 or larger.
- 2) The *shmsys:shminfo_shmmax* parameters should be set to the suggested value in the above table, or 90% of the physical memory (in bytes), whichever is higher. For example, if you have 196 MB of physical memory in your system, set the *shmsys:shminfo_shmmax* parameter to 184968806 ($196 \times 0.9 \times 1024 \times 1024$).

To set a kernel parameter, add a line at the end of the */etc/system* file as follows:

```
set parameter_name = value
```

For example, to set the value of the *msgsys:msginfo_msgmax* parameter, add the following line to the end of the */etc/system* file:

```
set msgsys:msginfo_msgmax = 65535
```

Sample files for updating the kernel configuration parameters are provided in the `/opt/IBMdb2/V6.1/cfg` directory. The names for these files are as follows:

kernel.param.64MB

for systems with 64MB–124MB of physical memory

kernel.param.128MB

for systems with 128MB–256MB of physical memory

kernel.param.256MB

for systems with 256MB–512MB of physical memory

kernel.param.512MB

for systems with 512MB–1GB of physical memory

kernel.param.hints

for systems with more than 1GB of physical memory

Depending upon the amount of physical memory in your system, append the appropriate kernel configuration parameter file to the `/etc/system` file. If necessary, change the value of the `shmsys:shminfo_shmmax` parameter as described in Note 2 above.

After updating the `/etc/system` file, reboot the system.

Performing the Installation

This section describes how to install DB2 EEE on a Solaris system or a Solaris cluster.



When you use the DB2 Installer program, you should be aware of the following:

- The DB2 Installer's **db2setup** command only works with Bash, Bourne, and Korn shells. Other shells are not supported.
- You can generate a trace log, *db2setup.trc*, to record errors experienced during the installation. Run the **db2setup** command as follows:

```
db2setup -d
```

This creates a trace file, `/tmp/db2setup.trc`.

To install DB2 Enterprise - Extended Edition (DB2 EEE), perform the following steps:

To mount the CD-ROM on Solaris, perform the following steps:

Step 1. Log in as a user with root authority.

Step 2. If the Volume Manager is not running on your system, enter the following commands to mount the CD-ROM:

```
mkdir -p /cdrom/unnamed_cdrom
mount -F hsfs -o ro /dev/dsk/c0t6d0s2 /cdrom/unnamed_cdrom
```

where */cdrom/unnamed_cdrom* represents the CD-ROM mount directory.

Note: If you are mounting the CD-ROM drive from a remote system using NFS, the CD-ROM file system on the remote machine must be exported with root access. You must also mount that file system with root access on the local machine.

If the Volume Manager (vold) is running on your system, the CD-ROM is automatically mounted as:

/cdrom/unnamed_cdrom

Step 3. Log out.

Step 1. Log on to the system as a user with root authority.



If you are installing DB2 EEE on a Solaris cluster, you must log on to each system and install this product, or use the *cconsole* or *ctelnet* utilities.

Step 2. If the Volume Manager (vold) is not running, enter the following commands to mount the CD-ROM:

```
mkdir -p //cdrom/unamed_cdrom
mount -F hsfs -o ro /dev/dsk/c0t6d0s2 /cdrom/unamed_cdrom
```

where */cdrom/unamed_cdrom* is the CD-ROM mount point.

If the Volume Manager (vold) is running on your system, the CD-ROM is automatically mounted as */cdrom/unamed_cdrom*.



If you are mounting a CD-ROM drive from a remote system using NFS, the CD-ROM file system on the remote machine must be exported with root access. You must also mount that file system with root access on the local machine.

Step 3. Insert the appropriate CD-ROM into the drive.

Step 4. Change to the directory where the CD-ROM is mounted by entering the following command:

```
cd /cdrom/unnamed_cdrom
```

where */cdrom/unnamed_cdrom* is the mount point of the CD-ROM on Solaris.

Step 5. Enter the **./db2setup** command to start the DB2 Installer program. The **Install DB2 V6** window opens.



It will take some time for the DB2 Installer program to start up, as it is scanning your system for information.

+----- Install DB2 V6 -----+

Select the products you are licensed to install. Your Proof of Entitlement and License Information booklet identify the products for which you are licensed.

To see the preselected components or customize the selection, select Customize for the product.

<input type="checkbox"/> DB2 Administration Client	: Customize. . .
<input checked="" type="checkbox"/> DB2 UDB Enterprise - Extended Edition	[Customize. . .]
<input type="checkbox"/> DB2 Software Developer's Kit	: Customize. . .

To choose a language for the following components, select Customize for the product.

DB2 Product Messages	[Customize. . .]
DB2 Product Library	[Customize. . .]

[OK] [Cancel] [Help]

+-----+

- Step 6. From the product list on the *Install DB2 V6* screen, select the **DB2 Enterprise - Extended Edition** option and press the **Enter** key. Press the **Tab** key to change the highlighted option and the **Enter** key to select or deselect an option. To display the components for a DB2 product that you want to install, select the **Customize** option. To go back to a previous windows at any time, select the **Cancel** option.



To refresh the current screen, press the **F5** key or **Ctrl+L**.

- Step 7. When you have finished selecting the DB2 product and its components, select **OK** to continue with the installation. For more information or assistance during the installation of any DB2 product or component, select the **Help** option.
- Step 8. Log out.



When the installation completes, the software is installed in the /opt/IBMdb2/V6.1 directory.

Post Installation Steps

This section describes that remaining steps that you have to perform to get your partitioned database system up and running.

Step 1. Update the Node Configuration File

The node configuration file (db2nodes.cfg), located in the *INSTHOME*/sql1lib/ directory (where *INSTHOME* is the home directory of the instance owner) contains configuration information for all the database partition servers that are participating in the instance. There is a db2nodes.cfg file for every multipartition instance. This file contains one entry for each database partition server for a particular DB2 instance. By default, when you create an instance, the db2nodes.cfg file is automatically created and an entry for the workstation where the instance was created is added to the file.

The format of the db2nodes.cfg file is as follows:

```
nodenum    hostname    logical port    netname
```

where:

nodenum A unique number, between 0 and 999, that identifies a database partition server in a partitioned database system.

When you want to scale your partitioned database system, you add an entry for each database partition server to the db2nodes.cfg file. The *nodenum* that you select for additional database partition servers must be in ascending order, however, gaps can exist in this sequence. You may choose to put a gap between the *nodenum* values if you plan to add an MLN and wish to keep them logically grouped together in this file.

This entry is required.

hostname The TCP/IP *hostname* of that database partition server (which you recorded earlier) for use by the FCM.

This entry is required.



You can specify the name of a high speed interconnect for the *hostname* value. This allows the **db2start**, **db2stop**, and **db2_all** commands to use the high speed interconnect for communications.

logical port Specifies the logical port number for the database partition

server. This field is used to specify a particular database partition server on a workstation that is running MLNs. If there is no entry for this field, the default is 0. However, if you add an entry for the *netname* field, you must enter a number for the *logical port* field.

If you are not using MLNs and you specify an entry for the *netname* field, the entry *must* be set to 0.

If you are using MLNs, the *logical port* you specify *must* start at 0 and continue in ascending order with no gaps (for example, 0,1,2).

This field is only optional when you are not using MLNs or a high speed interconnect.

netname Specifies the hostname or the IP address of the high speed interconnect for FCM communications.

If an entry is specified for this field, all communications between database partition servers (except for communications as a result of the **db2start**, **db2stop**, and **db2_all** commands) is handled through the high speed interconnect.

This parameter is only required if you are using a high speed interconnect for database partition communications.

For example, when you created the DB2 instance db2inst1, on the workstation called workstation1, the db2nodes.cfg file was updated as follows:

```
0          workstation1  0
```

If you want your partitioned database system to contain two workstations, called workstation1 and workstation2, you would update the db2nodes.cfg file as follows:

```
0          workstation1  0
1          workstation2  0
```

If you want your partitioned database system to contain two workstations, called workstation1 and workstation2, and workstation1 was running 3 database partition servers, you would update the db2nodes.cfg file as follows:

```
0          workstation1  0
1          workstation1  1
2          workstation1  2
3          workstation2  0
```

If you want your partitioned database system to contain two workstations, called workstation1 and workstation2 (with workstation2 running 2 database

partition servers), and use a high speed interconnect called switch1 and switch2, you would update the db2nodes.cfg file as follows:

0	workstation1	0	switch1
1	workstation2	0	switch2
2	workstation2	1	switch2

To update the db2nodes.cfg file, perform the following steps:

- Step 1. Log on to the instance as a user with system administrative (SYSADM) authority (by default, the username that you created for the instance owner). For more information, see “Working with the System Administrative Group” on page 116.
- Step 2. Ensure the DB2 instance is stopped by entering the *INSTHOME/sqllib/adm/db2stop* command, where *INSTHOME* is the home directory of the instance owner.



The db2nodes.cfg is locked when the instance is running and can only be edited when the instance is stopped.

- Step 3. Edit the db2nodes.cfg file and add an entry for each database partition server that will participate in your partitioned database system.
- Step 4. Log out.

For more information on the db2nodes.cfg file, refer to the *Administration Guide*.

Step 2. Enable the Fast Communications Manager

This section describes how to enable communications between the database partition servers that participate in your partitioned database system. Communications between database partition servers is handled by the Fast Communications Manager (FCM). To enable FCM, you must ensure that the port or port range that you reserve in the /etc/services file is available on every workstation in your partitioned database system.

When you created a DB2 instance using the DB2 Installer program, an entry similar to the following was placed in the /etc/services file:

```
DB2_instance_name      60000/tcp
```

where *instance_name* is the name of the instance that you created.

This entry *must* be of the form:

```
DB2_instance_name      port_number
```

where:

- *instance_name* is the name of the multipartition instance.
- *port_number* is the port number that you reserve for database partition server communications.



The port 60000 is the default FCM port used by an instance for any instance that was created using the DB2 Installer program.

For example, if you created an instance called db2inst1, the entry placed in the /etc/services file would be:

```
DB2_db2inst1      60000/tcp
```

If you are planning to run multiple logical nodes (MLNs), you must reserve a range of consecutive ports (one for each logical node). The consecutive port range that you reserve must be equal to the number of MLNs on the database partition server that is running the most MLNs in your partitioned database system. To reserve a port range for FCM to use, you will need to add another entry to the /etc/services file that specifies the end of the port range. This entry *must* be of the format:

```
DB2_instance_name_END    port_number
```

- *instance_name* is the name of the multipartition instance.
- *port_number* is the port number the last port in the port range that you reserved for database partition server communications.

For example, if you have a database partition server that is running 6 database partitions you will need to specify 6 consecutive ports as your port range. In our example, the entry in the /etc/services file would be:

```
DB2_db2inst1      60000/tcp
DB2_db2inst1_END  60005
```



You can add a comment that describes these entries using the # comment identifier. For example:

```
DB2_db2inst1      60000/tcp # This is an EEE instance
DB2_db2inst1_END  60005/tcp # that is running 6 MLNs.
```

To enable your partitioned database system for FCM communications, perform the following steps on each machine that will participate in your partitioned database system:

- Step 1. Log on the workstation where you created a DB2 instance as a user with root authority.
- Step 2. Edit the /etc/services file and add the port or port range entries that you have reserved for FCM communications.
- Step 3. Log out.

Step 4. Log on to each additional workstation in your partitioned database system and add the same entries to the `/etc/services` file that you added to the `/etc/services` file on the workstation where you created the DB2 instance.

Step 5. Log out.

For more information on FCM communications, refer to the *Administration Guide*.

Step 3. Enable the Execution of Remote Commands

In a multipartition instance, each database partition server must have the authority to perform remote commands on all the other database partition servers in the partitioned database system. To do this, you can either create an `INSTHOME/.rhosts` file (where `INSTHOME` is the home directory of the instance owner) or an `/etc/hosts.equiv` file.

If you decide to create an `/etc/hosts.equiv` file, you will have to create this file on every workstation that participates in your partitioned database system.

If you want to use the `INSTHOME/.rhosts` file, it should contain entries similar to the following:

```
workstation1.torolab.ibm.com db2inst1
workstation2.torolab.ibm.com db2inst1
switch01.torolab.ibm.com      db2inst1
switch02.torolab.ibm.com      db2inst1
```

For more information about the `INSTHOME/.rhosts` or the `/etc/hosts.equiv` file, refer to your operating system's documentation.

Step 4. Create Links for DB2 Files (Optional)

You can use the **db2ln** command to create links for the DB2 files to the `/usr/lib` directory, and for the include files to the `/usr/include` directory, for a particular version and release level of DB2 Enterprise - Extended Edition (DB2 EEE). Links can be established for only one version of DB2 on a given workstation.

You may want to create these links if you are developing or running applications and want to avoid having to specify the full path to the product libraries and include files.

To create links for the DB2 files, perform the following steps:

Step 1. Log on to the Control workstation as a user with root authority.

Step 2. Enter the following command to create the links for the DB2 files:

```
dsh -a /opt/IBMDB2/V6.1/cfg/db2ln
```

Step 3. Log out.

If there are existing links to the `/usr/lib` and `/usr/include` directories from previous versions of DB2 EEE, they will automatically be removed by executing the **db2ln** command to create links for this version of DB2 EEE.

Step 5. Start the Administration Server

This section describes how to start and configure the Administration Server. Before you can use the Client Configuration Assistant (CCA) or the Control Center to administer your partitioned database system, you must ensure the Administration Server is started and the `db2cc1st` daemon is running on every database partition server in your partitioned database system.

When you created the Administration Server using the DB2 Installer program, an entry similar to the following was placed in the `/etc/services` file on the workstation where you created the Administration Server:

```
db2ccmsrv      523/tcp
```

To enable your partitioned database system so that it can be administered by the CCA or the Control Center, perform the following steps:

- Step 1. Log on the workstation as a user with root authority.
- Step 2. Open the `/etc/services` file and verify that the port was reserved for the Administration Server.
- Step 3. Log out.
- Step 4. Log on to the workstation with the username that you created for the Administration Server (for example, `db2as`).
- Step 5. Enter the following command to start the Administration Server

```
DASINSTHOME/sql1lib/bin/db2admin start
```

where *DASINSTHOME* is the home directory of the Administration Server.

- Step 6. Enter the following command to start the `db2cc1st` daemon

```
DASINSTHOME/sql1lib/bin/db2cc1st
```

where *DASINSTHOME* is the home directory of the Administration Server.

- Step 7. Ensure that the `db2cc1st` daemon is running by entering the following command:

```
ps -ef | grep db2cc1st
```

Error information for the `db2cc1st` daemon is returned to the `syslog` file. Errors are written to this file, instead of the `db2diag.log` file, because the daemon runs independent from the instance.

Step 8. Log out.

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



You can set up the db2cc1st daemon to start when a workstation is restarted by adding the follow command to the `/etc/inittab` file.

```
cl:234:once:DASINSTHOME/sql1lib/bin/db2cc1st"
```

where *DASINSTHOME* is the home directory of the Administration Server.

Step 6. Create the SAMPLE Database (Optional)

This section describes how to verify that DB2 Enterprise - Extended Edition (DB2 EEE) is installed and configured correctly by creating the SAMPLE database on your system, and accessing data from the database.



To complete the steps in this section, we recommend that you create a separate file system (for example, `/database`) on *each* database partition server that participates in your partitioned database system.

For more information on how to create a file system, refer to your operating system's documentation.

To verify that you have installed and configured DB2 EEE, perform the following steps:

Step 1. Log on to the system as a user with system administrative (SYSADM) authority (by default, the username that you created for the instance owner). For more information, see "Working with the System Administrative Group" on page 116.

Step 2. Enter the following command to create the SAMPLE database:

```
INSTHOME/sql1lib/bin/db2saml path
```

where:

- *INSTHOME* is the home directory of the instance owner.
- *path* is a *local* file system that exists on every workstation that participates in your partitioned database system.

For example, to create the SAMPLE database on the `/database` file system, enter the following command:

```
INSTHOME/sql1lib/bin/db2saml /database
```

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:

```
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 113 or “Entering Commands Using the Command Line Processor” on page 114.



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.

Step 7. Install the License Key

DB2 software registration is handled automatically if you installed your DB2 product from a CD-ROM using the DB2 Installer program. If you installed DB2 using your UNIX operating system’s native installation tools, you must use the **db2licm** command to register DB2. Refer to the *Command Reference* for more information.

Part 3. Installing DB2 Clients

Chapter 5. 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 6. Installing DB2 Clients on Windows 32-Bit Operating Systems" on page 69
- "Chapter 7. Installing DB2 Clients on OS/2 Operating Systems" on page 73
- "Chapter 8. Installing DB2 Clients on UNIX Operating Systems" on page 75

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 6. 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 17.
- ___ 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:

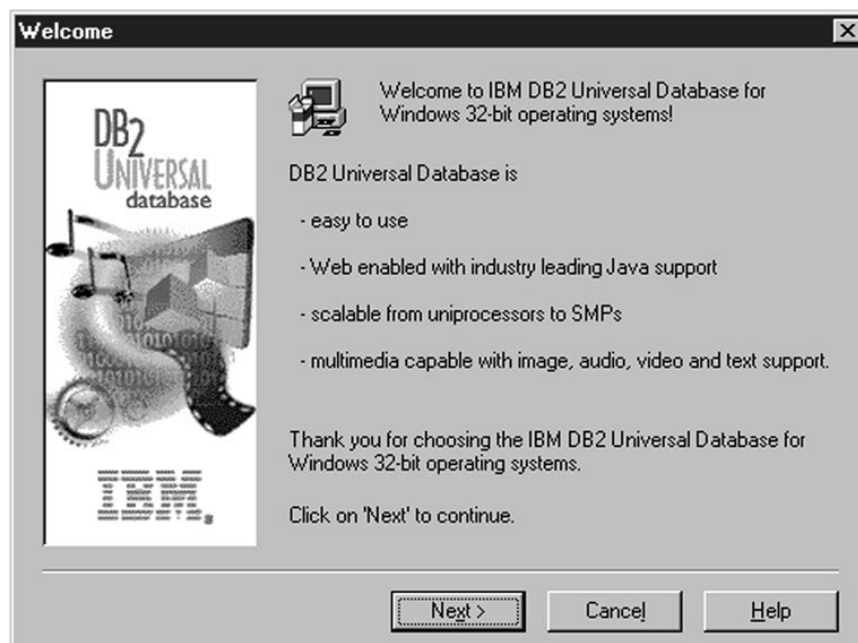
`x:\setup /i language`

where:

- *x*: represents your CD-ROM drive
- *language* represents the country code for your language (for example, EN for English). Table 15 on page 136 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 9. Configuring Client-to-Server Communications Using the Client Configuration Assistant” on page 87.

Chapter 7. Installing DB2 Clients on OS/2 Operating Systems

This section contains the information that you need to install a DB2 client or Software Developer's Kit on OS/2 operating systems. If you have a pre-Version 6 DB2 client for OS/2, the WIN-OS/2 support installed will be kept at its current level.



If you want to run Windows 3.x applications on your OS/2 system, you must also install the DB2 Client Application Enabler for Windows 3.x on your system. For more information, connect to the IBM DB2 Client Application Enabler Web site at <http://www.software.ibm.com/data/db2/db2tech/clientpak.html>.

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 17.
- ___ 2. A user ID to perform the installation.
If UPM is installed, the user ID you specify must have *Administrator* or *Local Administrator* authority. Create a user ID with these characteristics if necessary.
If UPM is not installed, DB2 will install it and set up the userid USERID with password PASSWORD.
- ___ 3. 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 116. For more information on valid DB2 usernames, see "Appendix D. Naming Rules" on page 137.

Installation Steps


To install a DB2 client or Software Developer's Kit for OS/2, perform the following steps:

- Step 1. Insert the appropriate CD-ROM into the drive.
- Step 2. Open an OS/2 command window and set the directory to your CD-ROM drive by entering the following command:


```
x:\install
```

where x: represents your CD-ROM drive.

- Step 3. The Welcome window opens.



- Step 4. Respond to the install 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.



For information on errors encountered during installation, see the 11.log and 12.log files. These files store general information and error messages resulting from installation and uninstall activities. By default, these files are 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*.



To configure your client to access remote servers, go to "Chapter 9. Configuring Client-to-Server Communications Using the Client Configuration Assistant" on page 87.

Chapter 8. Installing DB2 Clients on UNIX Operating Systems

This section contains the information that you need to install a DB2 client or Software Developer's Kit on UNIX operating systems.



If you want to remove a DB2 product, go to "Appendix G. Removing DB2 Products" on page 181.

Before You Begin

Before you begin installing a DB2 client or Software Developer's Kit using the DB2 Installer program, you need to gather the following 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 17.
- ___ 2. A username for the default DB2 instance. We recommend that you create a new group and use it as the primary group for the DB2 instance owner.



When you use the DB2 Installer program, you should be aware of the following:

- The DB2 Installer's **db2setup** command only works with Bash, Bourne, and Korn shells. Other shells are not supported.
- You can generate a trace log, *db2setup.trc*, to record errors experienced during the installation. Run the **db2setup** command as follows:

```
db2setup -d
```

This creates a trace file, */tmp/db2setup.trc*.

Installation Steps

To install a DB2 client or Software Developer's Kit on AIX, HP-UX, Linux, Silicon Graphics IRIX, and Solaris systems, perform the following steps:

Step 1. Update Kernel Configuration Parameters

To run a DB2 client or Software Developer's Kit on a HP-UX, or Solaris system, you may have to update some kernel configuration parameters.



Go to the section that describes the kernel configuration parameters for the DB2 client that you want to install:

- “Recommended Values for HP-UX Version 10 and Version 11”
- “Recommended Values for Solaris” on page 77

You do *not* need to update any kernel configuration parameters to run a DB2 client or Software Developer’s Kit on an AIX, Linux, or SGI IRIX workstation. If you are installing a DB2 client or Software Developer’s Kit on AIX, Linux, or SGI IRIX, go to “Step 2. Mount the CD-ROM” on page 77.

Recommended Values for HP-UX Version 10 and Version 11

Table 8 lists the recommended values for HP-UX kernel configuration parameters. These values are valid for HP-UX 10 and HP-UX 11.

Table 8. HP-UX Kernel Configuration Parameters (Recommended Values)

Kernel Parameter	Recommended Value
msgseg	8192
msgmnb	65535 (1)
msgmax	65535 (1)
msgssz	16

Notes:

1. Parameters msgmnb and msgmax must be set to 65535 or higher.
2. To maintain the interdependency among kernel parameters, change parameters in the same sequence in which they appear in Table 8.

To change a value, perform the following steps:

- Step 1. Enter the **sam** command to start the System Administration Manager (SAM) program.
- Step 2. Double-click on the **Kernel Configuration** icon.
- Step 3. Double-click on the **Configurable Parameters** icon.
- Step 4. Double-click on the parameter that you want to change and enter the new value in the **Formula/Value** field.
- Step 5. Click on **OK**.
- Step 6. Repeat these steps for all of the kernel configuration parameters that you want to change.
- Step 7. When you are finished setting all of the kernel configuration parameters, select **Action->Process New Kernel** from the action menu bar.

The HP-UX operating system automatically reboots after you change the values for the kernel configuration parameters.



Go to “Step 2. Mount the CD-ROM” to continue with the installation.

Recommended Values for Solaris

Table 9 lists the recommended values for Solaris kernel configuration parameters.

Table 9. Solaris Kernel Configuration Parameters (Recommended Values)

Kernel Parameter	Recommended Value
msgsys:msginfo_msgmax	65535 (1)
msgsys:msginfo_msgmnb	65535 (1)
msgsys:msginfo_msgseg	8192
msgsys:msginfo_msgssz	16

Notes:

1. Parameters `msgsys:msginfo_msgmnb` and `msgsys:msginfo_msgmax` must be set to 65535 or higher.

To set a kernel parameter, add a line at the end of the `/etc/system` file as follows:

```
set parameter_name = value
```

where *parameter_name* represents the parameter you want to change.

For example, to set the value of parameter `msgsys:msginfo_msgmax`, add the following line to the end of the `/etc/system` file:

```
set msgsys:msginfo_msgmax = 65535
```

After changing the kernel parameters, reboot the system so that the changes can take effect.



Go to “Step 2. Mount the CD-ROM” to continue with the installation.

Step 2. Mount the CD-ROM

To install your DB2 product using the DB2 Installer program, you must first mount the CD-ROM.



Go to the section that describes the mounting instructions for a DB2 client that you want to install:

- “Mounting on AIX Systems”
- “Mounting on HP-UX Systems” on page 79
- “Mounting on Linux Systems” on page 79
- “Mounting on Silicon Graphics IRIX” on page 79
- “Mounting on Solaris Systems” on page 80

Mounting on AIX Systems

To mount the CD-ROM on AIX, perform the following steps:

- Step 1. Log in as a user with root authority.
- Step 2. Insert the CD-ROM in the drive.
- Step 3. Create a directory to mount the CD-ROM by entering the following command:

```
mkdir -p /cdrom
```

where `cdrom` represents the CD-ROM mount directory.

- Step 4. Allocate a CD-ROM file system by entering the following command:

```
smitty storage
```

- Step 5. Select **File Systems**.
- Step 6. Select **Add / Change / Show / Delete File Systems**.
- Step 7. Select **CDROM File Systems**.
- Step 8. Select **Add CDROM File System**.
- Step 9. Select **Device Name**.



Device names for CD-ROM file systems must be unique. If there is a duplicate device name, you may need to delete a previously-defined CD-ROM file system or use another name for your directory.

- Step 10. In the pop-up window, enter the following as **mount point**:

```
/cdrom
```

- Step 11. Mount the CD-ROM file system by entering the following command:

```
smit mountfs
```

- Step 12. Enter a value in the **FILE SYSTEM** name field. For example, the name could be `/dev/cd0`.
- Step 13. Enter a value in the **Directory over which to mount** field. For example, this value could be `/cdrom`.
- Step 14. Enter a value in the **Type of Filesystem** field. For example, this value could be `cdrfs`.

Step 15. Set the **Mount as READ-ONLY system** to Yes.

Step 16. Click on **OK**.

Step 17. Log out.



Go to “Step 3. Perform the Installation” on page 80 to continue with the installation.

Mounting on HP-UX Systems

To mount the CD-ROM on HP-UX, perform the following steps:

Step 1. Log in as a user with root authority.

Step 2. Insert the CD-ROM in the drive.

Step 3. Mount it by entering the following commands:

```
mkdir /cdrom
/usr/sbin/mount /dev/dsk/c0t2d0 /cdrom
```

where /cdrom represents the CD-ROM mount point.

Step 4. Log out.



The CD-ROM may also be mounted using the System Administration (SAM) tool. Consult your HP-UX documentation for more information about SAM.



Go to “Step 3. Perform the Installation” on page 80 to continue with the installation.

Mounting on Linux Systems

To mount the CD-ROM on Linux, perform the following steps:

Step 1. Log in as a user with root authority.

Step 2. Insert the CD-ROM in the drive and mount it with a command similar to the following:

```
mount -t iso9660 -o ro /dev/cdrom /cdrom
```

where /cdrom is the mount point of the CD-ROM.

Step 3. Log out.



Go to “Step 3. Perform the Installation” on page 80 to continue with the installation.

Mounting on Silicon Graphics IRIX

Perform the following steps to mount the CD-ROM on Silicon Graphics IRIX operating systems:

1. Log in as a user with root authority.
2. Insert the appropriate CD-ROM into the drive and mount it using the **mount** command as follows:

```
mount -t iso9660 device mount_point
```

For example, to mount the CD-ROM as /cdrom, enter the following commands:

```
mkdir /cdrom
mount -t iso9660 /dev/scsi/sc0d710 /cdrom
```

3. Log out.



Go to “Step 3. Perform the Installation” to continue with the installation.

Mounting on Solaris Systems

To mount the CD-ROM on Solaris, perform the following steps:

- Step 1. Log in as a user with root authority.
- Step 2. If the Volume Manager is not running on your system, enter the following commands to mount the CD-ROM:

```
mkdir -p /cdrom/unnamed_cdrom
mount -F hsfs -o ro /dev/dsk/c0t6d0s2 /cdrom/unnamed_cdrom
```

where */cdrom/unnamed_cdrom* represents the CD-ROM mount directory.

Note: If you are mounting the CD-ROM drive from a remote system using NFS, the CD-ROM file system on the remote machine must be exported with root access. You must also mount that file system with root access on the local machine.

If the Volume Manager (vold) is running on your system, the CD-ROM is automatically mounted as:

```
/cdrom/unnamed_cdrom
```

- Step 3. Log out.



Go to “Step 3. Perform the Installation” to continue with the installation.

Step 3. Perform the Installation

After you mount the CD-ROM file system, use the DB2 Installer program to install a DB2 product.



If you are installing a DB2 client from a remote server, it is better to use the **telnet** command to open a telnet session instead of using the **rlogin** command to connect to your remote server.

- Step 1. Log in as user with root authority.
- Step 2. Insert the appropriate CD-ROM into the drive.
- Step 3. Change to the directory where the CD-ROM is mounted by entering the following command:

```
cd /cdrom
```

where /cdrom is the mount point of the CD-ROM drive.

- Step 4. Change to the directory where the install image for the DB2 product that you want to install is located.

DB2 install images are available in the following directories (assuming that the mount point for the CD-ROM is /cdrom):

AIX	/cdrom/db2/aix
HP-UX Version 10	/cdrom/db2/hpux10
HP-UX Version 11	/cdrom/db2/hpux11
Linux	/cdrom/db2/linux
SGI	/cdrom/db2/sgi
Solaris	/cdrom/unnamed_cdrom/db2/solaris

Step 5. Enter the **./db2setup** command to start the DB2 Installer program.
The Install DB2 V6 window opens.

```

+----- Install DB2 V6 -----+
|
| Select the products you are licensed to install. Your Proof of
| Entitlement and License Information booklet identify the products for
| which you are licensed.
|
| To see the preselected components or customize the selection, select
| Customize for the product.
| [*] DB2 Administration Client [ Customize. . . ]
|
|
| To choose a language for the following components, select Customize for
| the product.
| DB2 Product Messages [ Customize. . . ]
| DB2 Product Library [ Customize. . . ]
|
| [ OK ] [ Cancel ] [ Help ]
+-----+

```

Step 6. From the product list on the **Install DB2 V6** window, select the product that you want to install and select **OK**.
Press the Tab key to change the highlighted option and the Enter key to select or deselect the option you want.
To display the required and optional components for a product you want to install, select **Customize**. To go back to a previous window at any time, select **Cancel**. You can invoke the online help by clicking on **Help**.



To refresh the current screen, press the **F5** key or **Ctrl+L**.

This completes the installation. The DB2 software is installed in the *DB2DIR* directory,

where *DB2DIR* = /usr/lpp/db2_06_01 on AIX
 = /opt/IBMdb2/V6.1 on Solaris



You can use the DB2 Installer program to create an additional instance, create an Administration Server, or add additional products or components after your initial installation. To create or add a new instance, an Administration Server, or additional DB2 products and components, enter the following command:

On AIX

```
/usr/lpp/db2_06_01/install/db2setup
```

On Solaris

```
/opt/IBMdb2/V6.1/install/db2setup
```

On Linux

```
/usr/IBMdb2/V6.1/install/db2setup
```



To configure your client to access a remote DB2 server, see “Chapter 10. Configuring Client-to-Server Communications Using the Command Line Processor” on page 99.

Part 4. Configuring DB2 Communications

Chapter 9. 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 111.

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 137.



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 116.

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 111.

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 90.
- “Adding a Database Manually” on page 92.

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 94.

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 95.

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 112.

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 88.

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 111.
- 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 111.
- 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.

Chapter 10. Configuring Client-to-Server Communications Using the Command Line Processor



To configure a client to communicate with a server, the remote server must be configured to accept in-bound requests for the communication protocol that you want to use. By default, the installation program automatically detects and configures any protocols running on your server.

If you have added a new protocol to your network, or wish to change any of the default settings on the server, refer to the *Installation and Configuration Supplement*.

This section describes how to configure a client to communicate with a server using the command line processor.



If you plan to use an OS/2 or Windows 32-bit client to communicate with a server, the Client Configuration Assistant (CCA) makes it easy to automate configuration and administration tasks. If you have installed the CCA, it is recommended that you use this tool to configure your DB2 clients for communications.

See “Chapter 9. Configuring Client-to-Server Communications Using the Client Configuration Assistant” on page 87 for more information.

For instructions on entering DB2 commands, see “Entering Commands Using the Command Center” on page 113 or “Entering Commands Using the Command Line Processor” on page 114.



Go to the section that describes how to configure communications to access a remote server using the communication protocol of your choice:

- For TCP/IP, see “Configuring TCP/IP on the Client”.
- For APPC, refer to the *Installation and Configuration Supplement*.

Configuring TCP/IP on the Client

This section assumes that TCP/IP is functional on the client and server workstations. See “Software Requirements” on page 22 for the communication protocol requirements for your platform. See “Possible Client-to-Server Connectivity Scenarios” on page 27 for the supported communication protocols for your particular client and server.

To set up TCP/IP communications on a DB2 client, perform the following steps:

Step 1. Identify and record parameter values.

Step 2. Configure the client:

- a. Resolve the server's host address.
- b. Update the services file.
- c. Catalog a TCP/IP node.
- d. Catalog the database.

Step 3. Test the connection between the client and server.



Due to the characteristics of the TCP/IP protocol, TCP/IP may not be immediately notified of the failure of a partner on another host. As a result, a client application accessing a remote DB2 server using TCP/IP, or the corresponding agent at the server, may sometimes appear to be hung. DB2 uses the TCP/IP SO_KEEPALIVE socket option to detect when there has been a failure and the TCP/IP connection has been broken.

If you are experiencing problems with your TCP/IP connection, refer to the *Troubleshooting Guide* for information on how to adjust this parameter and other common TCP/IP problems.

Step 1. Identify and Record Parameter Values

As you proceed through the configuration steps, complete the *Your Value* column in the following table. You can fill in some of the values before you start configuring this protocol.

Table 10. TCP/IP Values Required at the Client

Parameter	Description	Sample Value	Your Value
Host Name • Hostname (<i>hostname</i>) or • IP address (<i>ip_address</i>)	Use the <i>hostname</i> or <i>ip_address</i> of the remote server workstation. To resolve this parameter: • Enter the hostname command at the server to obtain the <i>hostname</i> . • Contact your network administrator to obtain the <i>ip_address</i> or enter the ping <i>hostname</i> command.	serverhost or 9.21.15.235	

Parameter	Description	Sample Value	Your Value
Service Name <ul style="list-style-type: none"> • Connection Service name (<i>svcname</i>) or • Port number/Protocol (<i>port_number/tcp</i>) 	<p>Values required in the <code>services</code> file.</p> <p>The Connection Service name is an arbitrary name that represents the Connection port number (<i>port_number</i>) on the client.</p> <p>The port number for the client must be the same as the port number that the <i>svcname</i> parameter maps to in the <code>services</code> file at the server. (The <i>svcname</i> parameter is located in the database manager configuration file on the server.) This value must not be in use by any other applications, and must be unique within the <code>services</code> file.</p> <p>Ensure that you do not specify a port number that is being used by the Fast Communication Manager (FCM).</p> <p>Contact your database administrator for the values used to configure the server.</p>	<p>server1</p> <p>3700/tcp</p>	
Node name (<i>node_name</i>)	A local alias, or nickname, that describes the node to which you are trying to connect. You can choose any name you want; however, all node name values within your local node directory must be unique.	db2node	

The following steps configure this protocol on the client. Replace the sample values with your worksheet values.

Chapter 10. Configuring Client-to-Server Communications Using the Command Line Processor **101**

The client must know the address of the server to which it is attempting to establish communications. If a name server does not exist on your network, you may directly specify a hostname that maps to the IP address (*ip_address*) of the server in the local hosts file. See Table 11 for the location of the hosts file for your particular platform.



If you are planning on supporting a UNIX client that is using Network Information Services (NIS), and you are not using a name server on your network, you must update the hosts file located on your NIS master server.

Table 11. Location of the Local Hosts and Services Files

Platform	Location
OS/2	Specified by the <i>etc</i> environment variable. Enter the set etc command to determine the location of your local hosts or services files.
Windows NT	Located in the winnt\system32\drivers\etc directory.
Windows 9x	Located in the windows directory.
UNIX	Located in the /etc directory.

Edit the client's hosts file and add an entry for the server's hostname. For example:

```
9.21.15.235    serverhost    # host address for serverhost
```

where:

9.21.15.235 represents the *ip_address*

serverhost represents the *hostname*

represents a comment describing the entry



If the server is not in the same domain as the client, you must provide a fully qualified domain name such as *serverhost.vnet.ibm.com*, where *vnet.ibm.com* is the domain name.

Step B. Update the Services File



If you are planning to catalog a TCP/IP node using a port number (*port_number*), skip this step and go to "Step C. Catalog a TCP/IP Node" on page 103.

Using a local text editor, add the Connection Service name and port number to the client's services file for TCP/IP support. For example:

```
server1 3700/tcp # DB2 connection service port
```

where:

server1 represents the Connection Service name

3700 represents the Connection port number
 tcp represents the communication protocol that you are using
 # represents a comment describing the entry

The port number used on the client must match the port number used on the server.



If you are planning on supporting a UNIX client that uses Network Information Services (NIS), you must update the services file located on your NIS master server.

The file called services is located in the same directory as the local hosts file that you may have edited in “A. Resolve the Server’s Host Address” on page 101.

Ensure that you did not specify a port number that is being used by the Fast Communication Manager (FCM) or any other process.

See Table 11 on page 102 for the location of the services file for your particular platform.

Step C. Catalog a TCP/IP Node

You must add an entry to the client’s node directory to describe the remote node. This entry specifies the chosen alias (*node_name*), the *hostname* (or *ip_address*), and the *svcname* (or *port_number*) that the client will use to access the remote server.

To catalog a TCP/IP node, 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 137.



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 116.

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

Step 2. If you are using a UNIX client, set up the instance environment and invoke the DB2 command line processor. Run the start-up script as follows:

```
. INSTHOME/sql1lib/db2profile    (for Bash, Bourne or Korn shell)
source INSTHOME/sql1lib/db2cshrc (for C shell)
```

where *INSTHOME* is the home directory of the instance.

Step 3. Catalog the node by entering the following commands:

```
catalog tcpip node node_name remote [hostname|ip_address] \  
server [svccname|port_number]  
terminate
```

For example, to catalog the remote server *serverhost* on the node called *db2node*, using the service name *server1*, enter the following:

```
catalog tcpip node db2node remote serverhost server server1  
terminate
```

To catalog a remote server with the IP address *9.21.15.235* on the node called *db2node*, using the port number *3700*, enter the following:

```
catalog tcpip node db2node remote 9.21.15.235 server 3700  
terminate
```



If you need to change values that were set with the **catalog node** command, perform the following steps:

Step 1. Run the **uncatalog node** command in the command line processor as follows:

```
uncatalog node node_name
```

Step 2. Recatalog the node with the values that you want to use.

Step D. Catalog the Database

Before a client application can access a remote database, the database must be cataloged on the server node and on any client nodes that will connect to it. When you create a database, it is automatically cataloged on the server with the database alias (*database_alias*) the same as the database name (*database_name*). The information in the database directory, along with the information in the node directory, is used on the client to establish a connection to the remote database.

To catalog a database on the client, 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 137.



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 116.

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

Step 2. Fill in the *Your Value* column in the following worksheet.

Table 12. Worksheet: Parameter Values for Cataloging Databases

Parameter	Description	Sample Value	Your Value
Database name (<i>database_name</i>)	The database alias (<i>database_alias</i>) of the <i>remote</i> database. When you create a database, it is automatically cataloged on the server with the database alias (<i>database_alias</i>) the same as the database name (<i>database_name</i>).	sample	
Database alias (<i>database_alias</i>)	An arbitrary local nickname for the remote database, on the client. If you do not provide one, the default is the same as the database name (<i>database_name</i>). This is the name that you use when connecting to a database from a client.	tor1	
Node name (<i>node_name</i>)	The name of the node directory entry that describes where the database resides. Use the same value for node name (<i>node_name</i>) that you used to catalog the node in the previous step.	db2node	

Step 3. If you are using a UNIX client, set up the instance environment and invoke the DB2 command line processor. Run the start-up script as follows:

```
. INSTHOME/sql/lib/db2profile    (for Bash, Bourne or Korn shell)
source INSTHOME/sql/lib/db2cshrc (for C shell)
```

where *INSTHOME* represents the home directory of the instance.

Step 4. Catalog the database by entering the following commands:

```
catalog database database_name as database_alias at node node_name
terminate
```

For example, to catalog a remote database called *sample* so that it has the alias *tor1*, on the node *db2node*, enter the following commands:

```
catalog database sample as tor1 at node db2node
terminate
```



If you need to change values that were set with the **catalog database** command, perform the following steps:

Step a. Run the **uncatalog database** command as follows:

```
uncatalog database database_alias
```

Step b. Recatalog the database with the value that you want to use.

Step 3. Test the Client-to-Server Connection

When you have finished configuring the client for communications, perform the following steps to test the connection:



You will need to connect to a remote database to test the connection.

Step 1. Start the database manager by entering the **db2start** command on the server (if it was not automatically started at boot time).

Step 2. Enter the following command to connect the client to the remote database:

```
connect to database_alias user userid using password
```

The values for *userid* and *password* must be valid for the system on which they are authenticated. By default, authentication takes place on the server for a DB2 server and on the host or AS/400 machine for a DB2 Connect server.

If the connection is successful, you will get a message showing the name of the database to which you have connected. You are now able to retrieve data from that database. For example, to retrieve a list of all the table names listed in the system catalog table, enter the following SQL command in the Command Center or command line processor:

```
"select tablename from syscat.tables"
```

When you are finished using the database connection, enter the **connect reset** command to end the database connection.



You are now ready to start using DB2. For more advanced topics, refer to the *Administration Guide* and the *Installation and Configuration Supplement*.

Troubleshooting the Client-to-Server Connection

If the connection fails, check the following items:

At the *server*:

- ___ 1. The *db2comm* registry value includes the value *tcpip*.



Check the settings for the *db2comm* registry value by entering the **db2set DB2COMM** command. For more information, refer to the *Administration Guide*.

- ___ 2. The services file was updated correctly.
- ___ 3. The service name (*svcname*) parameter was updated correctly in the database manager configuration file.
- ___ 4. The security service was started. Enter the **net start db2ntsecserver** command (for Windows NT servers only).
- ___ 5. The database was created and cataloged properly.
- ___ 6. The database manager was stopped and started again (enter the **db2stop** and **db2start** commands on the server).
- ___ 7. Ensure that you did not specify a port number that is being used by the Fast Communication Manager (FCM) or any other process.



If there are problems starting a protocol's connection managers, a warning message appears and the error messages are logged in the *db2diag.log* file.

Additionally, any failure encountered while executing the **db2start** command is logged to a time-stamped file in the log subdirectory within the instance directory. This file contains the **db2start** results from each database partition server in the instance.

For more information on the *db2diag.log* file, refer to the *Troubleshooting Guide*.

At the *client*:

- ___ 1. If used, the services and hosts files were updated correctly.
- ___ 2. The node was cataloged with the correct hostname (*hostname*) or IP address (*ip_address*).
- ___ 3. The port number must match, or the service name must map to, the port number used on the server.
- ___ 4. The node name (*node_name*) that was specified in the database directory points to the correct entry in the node directory.
- ___ 5. The database was cataloged properly, using the *server's* database alias (*database_alias*) that was cataloged when the database was created on the server, as the database name (*database_name*) on the client.

If the connection still fails after you verify these items, refer to the *Troubleshooting Guide*.

Part 5. 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 the Software Registration Tool”.
 - “Starting the Client Configuration Assistant”.
 - “Starting the Control Center” on page 112.
 - “Entering Commands Using the Command Center” on page 113.
 - “Entering Commands Using the Command Line Processor” on page 114.
 - “Working with the System Administrative Group” on page 116.
-

Starting the Software Registration Tool

Start the Software Registration Tool as follows:

OS/2 Click on **OS/2 Warp** and select **DB2 for OS/2->Registration**

Windows 9x or Windows NT

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

UNIX

DB2 software registration is handled automatically if you installed your DB2 product from CD-ROM using the DB2 Installer program. If you installed DB2 using your UNIX operating system's native installation tools, you must enter the **db2licm** command to register DB2. See the *Command Reference* for further information.

Starting the Client Configuration Assistant

Start the Client Configuration Assistant (CCA) as follows:

OS/2 Click on **OS/2 Warp**, and select **DB2 for OS/2->Client Configuration Assistant**

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 OS/2

- Enter the **db2cc** command at a command prompt.
or
- Click on the **OS/2 Warp** button and select **DB2 for OS/2->Control Center**.

Note: You must have the prerequisite Java Runtime Environment (JRE) level to launch the Control Center using this command. For more information, refer to the Control Center readme, which can be found in the `x:\sql11b\cc\prime` directory, where `x` is drive where you installed your DB2 product.

On UNIX

- Enter the **db2cc** command at a command prompt.

Note: You must have the prerequisite Java Runtime Environment (JRE) level to launch the Control Center using this command. For more information, refer to the Control Center README, which can be found in the `INSTHOME/sql11b/cc/prime` directory, where `INSTHOME` is the home directory of the instance owner.

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:\sql11b\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.

Note: You must have the prerequisite Java Runtime Environment (JRE) level to launch the Command Center using this 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**.

On OS/2

- Click on the **OS/2 Warp** button and select **DB2 for OS/2 ->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 then **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:

OS/2 Click on **OS/2 Warp**, and select **DB2 for OS/2->Command Line Processor**, or open any command window.

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.

UNIX

Open any operating system command window.

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:

OS/2 Click on **OS/2 Warp**, and select **DB2 for OS/2->Command Line Processor** or enter the **db2** command.

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.

UNIX Enter the **db2** command from the command line processor

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:

UNIX Any valid DB2 username that belongs to the primary group of the instance owner's user ID.

Appendix B. Migrating from Previous Versions and Releases

This section describes how to migrate previous versions of DB2 to the Version 6 format. It also includes information on how to update a Version 6 single-partition database system to the Version 6 multipartition 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.



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.

Migrating from Previous Versions of DB2

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 then to move data from the previous database system to the new one. Although your data is not changed during the process, the database is different from the original one. The differences depend on the type of database system from which you are migrating.



This section contains the following information:

- “Migrating Instances” on page 120.
- “Migrating Databases” on page 129.
- “Optional Post Migration Actions” on page 130.

To avoid possible problems migrating to the DB2 Version 6 format, read all of the information in this section *before* starting the migration process.

Migrating Instances

This procedure describes how to migrate DB2 instances that were created using a previous version of DB2.



To update a Version 6 single-partition database system to a Version 6 partitioned database system, you must update the instance using the **db2iupdt** command. For more information on the **db2iupdt** command, refer to the *Command Reference*.

Each DB2 instance must be migrated separately. To successfully migrate a DB2 instance, perform the following steps:

- Step 1. Prepare the DB2 instance for migration.
- Step 2. Verify that the databases can be migrated. There are also migration considerations you should take into account if you are using the Version 2 user exit program.
- Step 3. Migrate the DB2 instance.

If you want to migrate several instances, you must repeat these steps for each instance.

Step 1. Prepare the DB2 Instance for Migration

Before you can migrate a DB2 instance, all applications using any databases owned by this instance must be completed. To prepare a DB2 instance for migration, perform the following steps:

- Step 1. Log in as the DB2 instance owner.
- Step 2. Ensure that there are no applications using any databases owned by this DB2 instance. To get a list of all applications owned by the instance, enter the **db2 list applications** command.

You can end a session by entering the **db2 terminate** command. It is not recommended to force termination of applications using the **db2 force applications all** command, since some applications may have unexpected behavior when terminated using this command. See the *Command Reference* for usage and details of this command.
- Step 3. When all applications are complete, stop all database server processes owned by the DB2 instance by entering the **db2stop** command.
- Step 4. Stop all command line processor sessions by entering the **db2 terminate** command in each session that was running the command line processor.
- Step 5. Enter the **db2_kill** command to clean up any remaining DB2 resources.

Step 6. Make a backup image of all databases. Refer to the *Administration Guide* for your DB2 product for information on making a backup copy of a database and to the *Command Reference* for the syntax of the backup command.

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

Step 7. Log off.

The DB2 instance is now ready for migration.

Step 2. Verify that Databases Can Be Migrated

DB2 provides the **db2ckmig** migration command which is used to verify whether all cataloged databases can be migrated.



To ensure that you can migrate the instance to DB2 Enterprise - Extended Edition, you should run the **db2ckmig** command. If instance migration failed, you must correct errors reported by this command. You can choose to run the **db2ckmig** command again to verify that the errors have been corrected and then migrate the instance.

For detailed information about the **db2ckmig** command, refer to the *Command Reference*.

To verify that all cataloged databases can be migrated, perform the following steps:

Step 1. Log in as the instance owner.

Step 2. Enter the following command:

```
DB2DIR/bin/db2ckmig -h -a 0 -l INSTHOME/migration.log
```

where <i>DB2DIR</i>	= /usr/lpp/db2_06_01	on AIX
	= /opt/IBMdbs/V6.1	on Solaris

and INSTHOME is the home directory of the instance and migration.log is the name for the output file.

Step 3. Check the log file. The log file displays the errors that occur when you run the **db2ckmig** command. If it shows any errors, see Table 13 on page 122 for suggested corrective actions. In the partitioned database system, errors may be returned at the database partition level.

Step 4. Check that the migration log file is empty before continuing with the instance migration.

Step 5. Backup the database after making corrections. See Table 13 on page 122 for more information on correcting error messages.

All local databases now have the same authentication type as the instance where they reside; the authentication type in the database directory is ignored by DB2 Version 6 servers. If a warning is logged due to a conflicting authentication type, and you want a database to retain its previous authentication type, then you can do one of the following:

- Change the authentication type of the instance to the previous one.
- Move the database to another instance that has the required authentication type.



Before changing the authentication type of the instance, you should make sure that the new authentication type will be appropriate for all databases residing there. Be certain to consider the security implications of the different authentication types.

If there are databases that you do not want to migrate, you can uncatalog them (along with all aliases). The **db2ckmig** command does not perform any verification of uncataloged databases.

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

Table 13. 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 13. 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 13. 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. For more information, refer to the <i>Administration Guide</i>.</p>

Table 13. 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 13. 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 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 127.

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*.

If you are using the Version 2.x user exit program, you should consider the following before migrating instances:

- If the Version 2.x **db2uexit** program is installed in the `INSTHOME/sqllib/adm` directory before migration, it will remain in this directory after migration. The DB2 Version 6 **db2uext2** program will be also installed in this directory. Its function is to invoke **db2uexit** using the Version 2 interface. This allows the old user exit program to be used on DB2 Version 6.
- If **db2uexit** is installed in a directory other than `INSTHOME/sqllib/adm`, it will not be installed after migration. For example, if **db2uexit** was in the `INSTHOME/sqllib/bin` directory, after migration the **db2uexit** file will not be in the `INSTHOME/sqllib/bin` directory.
- If you want to continue using the old user exit after migration, you must copy **db2uexit** to the `INSTHOME/sqllib/adm` directory. Then, copy **db2uext2.v2** from the `DB2DIR/misc` directory to the `INSTHOME/sqllib/adm` directory and rename it to **db2uext2**. Enter the following command to copy the file:

```
cp DB2DIR/misc/db2uext2.v2 INSTHOME/sqllib/adm/db2uext2
```

where <i>DB2DIR</i>	= <code>/usr/lpp/db2_06_01</code>	on AIX
	= <code>/opt/IBMdbs/V6.1</code>	on Solaris

Note: You must ensure that **db2uext2** is owned by the instance owner and is executable by the owner.

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 `INSTHOME/sqllib/adm` directory.

Installing DB2 Version 6

After you have successfully completed the pre-installation checks, you can now start installing DB2 Version 6 using either the interactive or distributed method. For installation procedures, see the following sections:

- “Chapter 3. Installing and Configuring DB2 Universal Database on AIX” on page 29 or “Chapter 4. Installing and Configuring DB2 Universal Database on Solaris” on page 51 for interactive installation.
- *Installation and Configuration Supplement* for distributed installation.

Step 3. Migrate the DB2 Instance



Only local cataloged databases that reside in the DB2 instance are checked for migration. Uncataloged databases may be unusable after the instance has been migrated. Refer to the *Administration Guide* for further information.

After an instance is ready for migration, use the **db2imigr** command to migrate the instance as follows:

1. Log in as user with root authority.



If the *library_path* environment variable is set to */usr/lib* on AIX or */opt/lib* on Solaris, and there is a link in */usr/lib* or */opt/lib* to the Version 6 libdb2 shared library, this can cause an error when using the **db2imigr** command. To fix the error, you should reset the *library_path* environment variable so that it does not reference the libraries in those paths by entering the following command:

```
unset library_path
```

where *library_path* represents:

- *LIBPATH* on AIX
- *LD_LIBRARY_PATH* on Solaris

After migrating the DB2 instance, you should reset *LIBPATH* to its original setting.

2. Run the **db2imigr** command as follows:

```
DB2DIR/instance/db2imigr [-d] [-a AuthType] [-u fencedID] InstName
```

where <i>DB2DIR</i>	= <i>/usr/lpp/db2_06_01</i>	on AIX
	= <i>/opt/IBMdbs2/V6.1</i>	on Solaris

and where:

-d Sets the debug mode that you can use for problem determination. This parameter is optional.

-a *AuthType*

Specifies the authentication type for the instance. Valid authentication types are (SERVER), (CLIENT), and (DCS). If the *-a* parameter is not specified, the authentication type defaults to (SERVER), if a DB2 server is installed. Otherwise, the *AuthType* is set to (CLIENT). This parameter is optional.

Notes:

- a. The authentication type of the instance applies to all databases owned by the instance.
- b. While authentication type (DCE) is an optional parameter, it is not valid to choose (DCE) for this command.

-u fencedID

Is the user under which the fenced user-defined functions (UDFs) and stored procedures will execute. This parameter is optional only when a DB2 Run-Time Client is installed. It is required for all other DB2 products.

InstName

Is the login name of the instance owner.



Because the INSTHOME directory is NFS mounted on all machines, you only have to run the **db2imigr** command on one machine to migrate the entire instance.

3. If there are any errors in verifying that all databases can be migrated, see Table 13 on page 122 and take the suggested corrective actions. Then, reenter the **db2imigr** command.



If you are migrating a DB2 Version 2.1 or Version 5 instance, created on AIX, and the instance uses the environment variable *DB2SORT* set to a keyword *SMARTSORT*, you must set the registry value *db2sort* after the instance is migrated to Version 6. Set the *db2sort* registry value to the run time library for the sort command as follows:

```
db2set DB2SORT="/usr/lib/libsort.a"
```

Migrating Databases

To migrate databases owned by an instance, perform the following steps:

- Step 1. Log on with a user ID that has SYSADM authority on the instance.
- Step 2. Ensure that the databases you want to migrate are cataloged.
- Step 3. Migrate the database. Refer to the *Command Reference* for the database migration syntax command.

Trouble-Shooting Database Migration Errors

The SQL1704N warning message, **Database migration failed** with reason code 6, is received because of a failure to relocate the database. To correct this error, you should check the db2diag.log file and do the following if you receive the following messages.

Table 14. Correcting SQL1704N warning messages

Error Message	Cause/Action
Mount unmounted filesystem(s) failed. Manual remount is required before retrying migration. Please check fname .	A problem was encountered when unmounting some of the filesystems. Attempts to clean up these unmounted filesystems also failed. Check the fname file and remount the listed filesystems before retrying the migration.
Unmount mounted filesystem(s) failed. No database relocation is done.	Attempts to unmount some filesystem(s) failed. Remount the unmounted filesystems and try the database migration again.
Remount unmounted filesystem(s) to new mount points failed. Restart of database migration is needed.	Attempts to remount filesystems to new mount points failed. Retry database migration.

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 116 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.



After you have finished migrating your previous versions of DB2 to the Version 6 format, you can now install DB2 Version 6 clients. See “Part 3. Installing DB2 Clients” on page 65 for more information.

Appendix C. National Language Support (NLS)

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

Language and Codeset Support for UNIX Operating Systems

DB2 supports many code sets and locales without translating the messages for the corresponding languages. Supporting a locale means that you can create and use a database in that locale, but you may have to view all panels and messages in a different language, if translated messages are not available in DB2. For a complete list of locales supported, refer to the *Administration Guide*.

If you want to operate in a different language environment, do the following:


1. Ensure that the appropriate message option for the desired language has been installed.
2. Set the *LANG* environment variable to the desired locale.

For example, to use fr_FR messages on DB2 for AIX, you must have the fr_FR message option installed and must set *LANG* to fr_FR.

The selected message catalog filesets are placed in the /usr/lpp/db2_06_01/msg/%L directory on the target workstation, where %L is equal to the locale name of the message catalog.

Code Page and Language Support for OS/2 and Windows Operating Environments

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 15 on page 136 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 15. Languages and Code Pages

Country Code	Language
bg	Bulgarian
br	Brazilian Portuguese
cn	Simplified Chinese (PRC)
cz	Czech
de	German
dk	Danish
en	English
es	Spanish
fi	Finnish
fr	French
gr	Greek
hu	Hungarian
il	Hebrew
it	Italian
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 138
- “Username, User ID, Group Name, and Instance Name Rules” on page 139
- “Password Rules” on page 141
- “DB2SYSTEM Naming Rules” on page 141

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 137.

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.

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 137.

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 137.

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 137.

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 137.

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



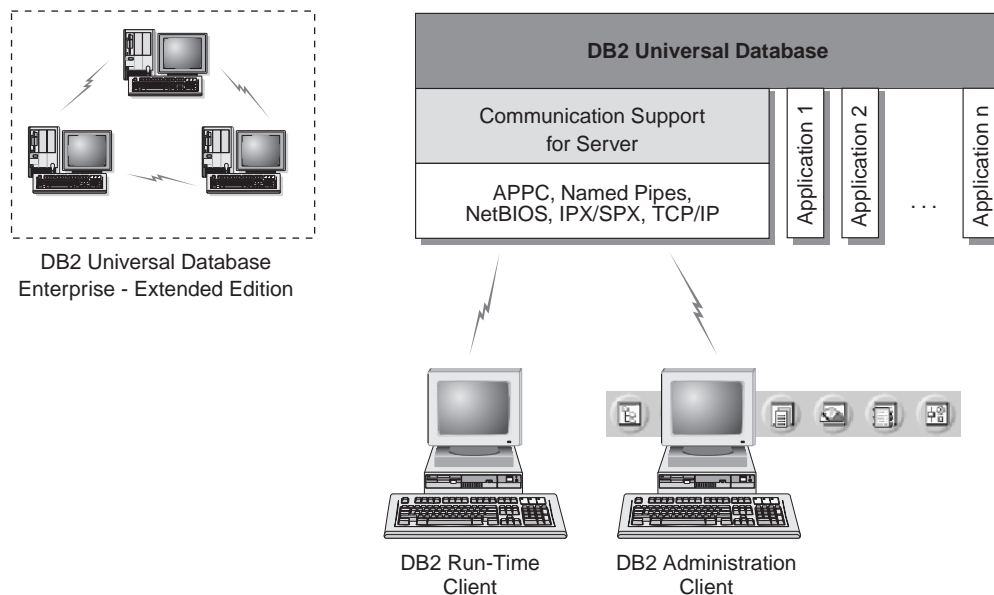
DB2 clients for the following releases and platforms are available for download from the web:

- 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 149 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 150.

Communicating Between DB2 Universal Database Servers

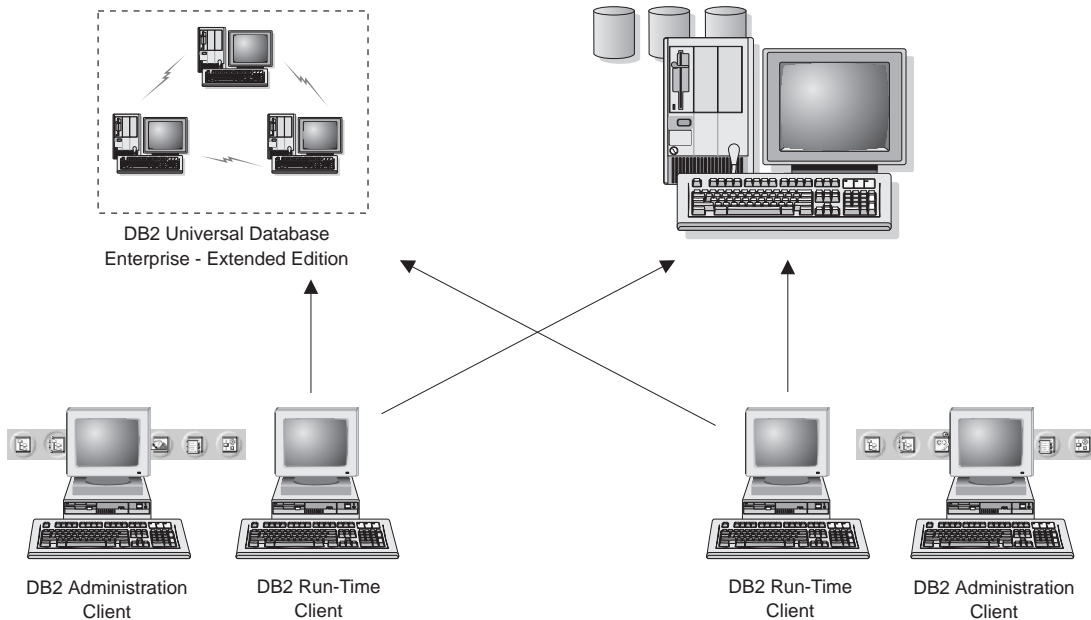


Figure 5. Accessing Data on Multiple Servers

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.

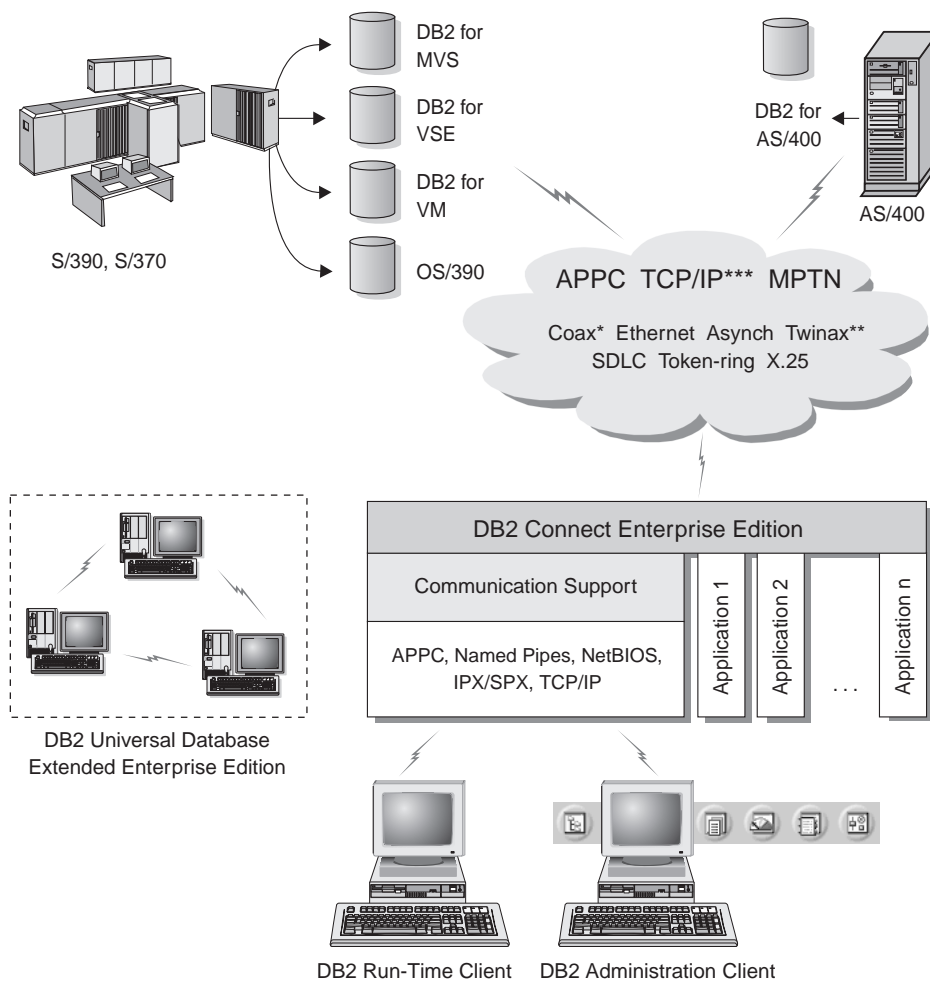
Although DB2 Connect is often installed on an intermediate server machine to connect DB2 clients to a host or AS/400 database, it is also installed on machines where multiple local users want to access the host or AS/400 servers directly. For example, DB2 Connect may be installed on a large machine with 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 153 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 155 for an example of a Java-enabled browser accessing data from remote DB2 databases.

Accessing DB2 Data Using JDBC

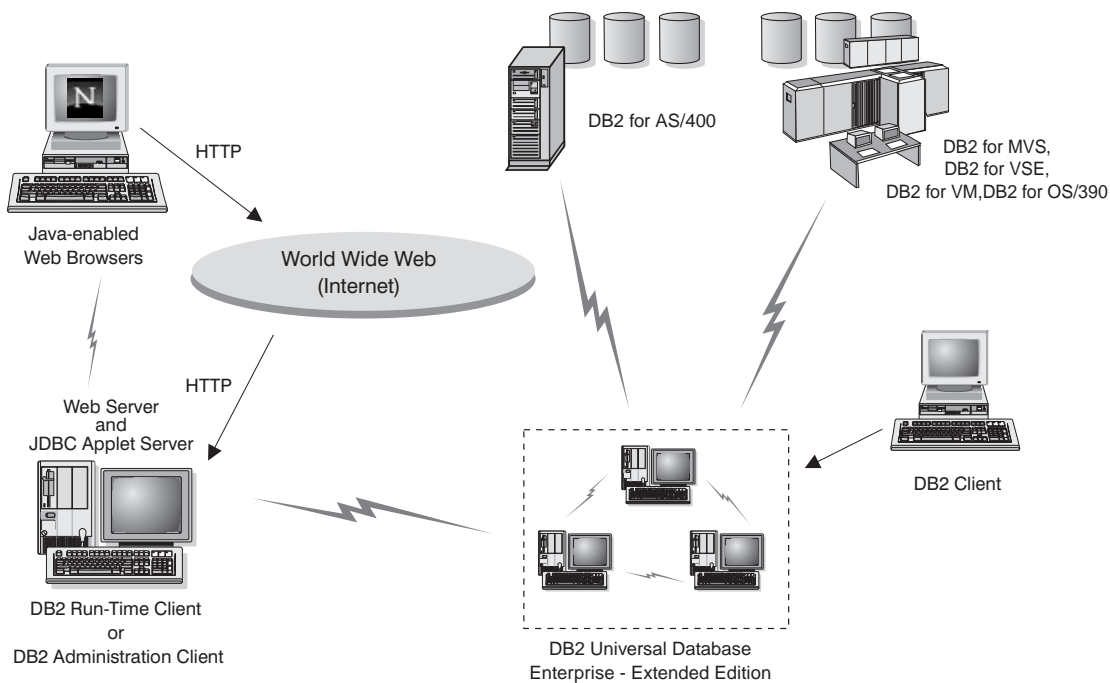


Figure 7. Accessing Data Stored on DB2 Using JDBC

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.

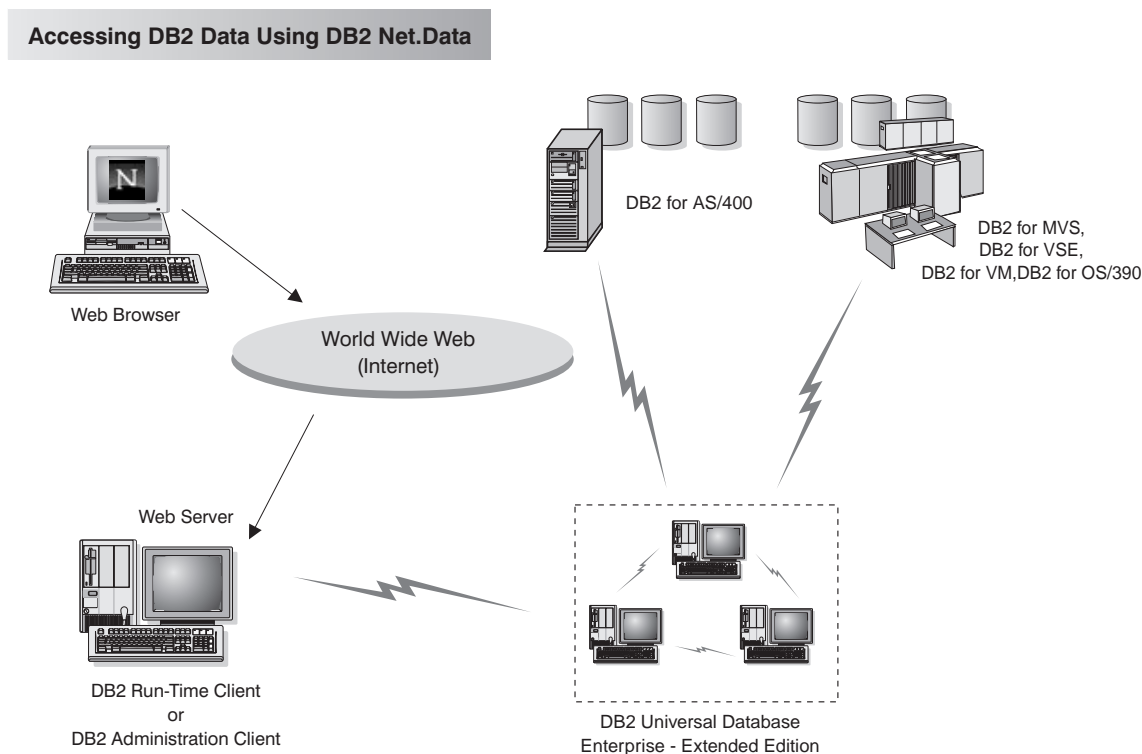


Figure 8. Accessing Internet Data Stored on DB2 Using Net.Data

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.

In a partitioned database system, you must also have a *listener* daemon that runs on each machine in the instance. This daemon is called `db2cc1st`, and is used by all instances that are on the machine. The Control Center uses the listener daemon to retrieve status, connection, and snapshot information from each database partition server. The daemon is not associated with a specific instance; rather, it functions as a global server for the machine.

The listener daemon requires a predefined named port. The named port must be called `db2ccmsrv` and must be defined in the `/etc/services` file on every machine. The named port can be assigned any unused port number, but the same number must be used for all machines.

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.

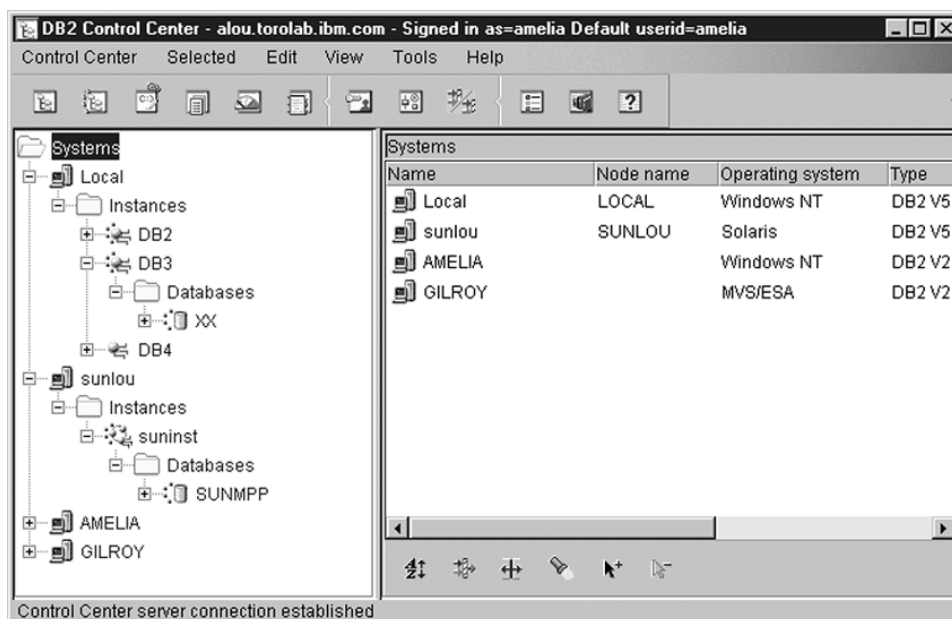


Figure 9. Control Center Main Window

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 165 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” 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 162 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.

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, see "Chapter 10. Configuring Client-to-Server Communications Using the Command Line Processor" on page 99.

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.

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 176 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 176.

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 175.
Search	See “Searching Online Information” on page 178.
Print	See “Printing the PostScript Books” on page 178.
Order	See “Ordering the Printed Books” on page 179.

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 176 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 176 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 168. 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 168. 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. Removing DB2 Products

This section shows you how to remove DB2 from your UNIX system.

To remove DB2, perform the following steps:

- Step 1. Stop the Administration Server.
- Step 2. Stop all DB2 Instances.
- Step 3. Remove the Administration Server.
- Step 4. Remove DB2 Instances. This step is optional.
- Step 5. Remove the DB2 products.

Step 1. Stop the Administration Server

You must stop the Administration Server before you remove DB2 products. To stop the Administration Server, you need to perform the following steps:

- Step 1. Log in as user with root authority.
- Step 2. Stop the Administration Server.

To stop the Administration Server, perform the following steps:

Step a. Log in as the Administration Server owner.

Step b. Run the start up script as follows:

```
. INSTHOME/sqllib/db2profile (for Bash, Bourne or Korn shell)
source INSTHOME/sqllib/db2cshrc (for C shell)
```

where *INSTHOME* is the home directory of the instance.

Step c. Stop the Administration Server by entering the following command:

```
db2admin stop
```

Step 2. Stop All DB2 Instances

You must stop all DB2 Instances before you remove DB2. To stop a DB2 instance, perform the following steps:

- Step 1. To obtain a list of the names of all DB2 instances on your system, enter the following command:

```
DB2DIR/bin/db2ilist
```

where <i>DB2DIR</i>	= /usr/lpp/db2_06_01	on AIX
	= /opt/IBMdb2/V6.1	on Solaris

Step 2. Stop the instance.

To stop a database instance, perform the following steps:

Step a. Log in as the instance owner.

Step b. Run the start up script as follows:

```
. INSTHOME/sqllib/db2profile    (for Bourne or Korn shell)
source INSTHOME/sqllib/db2cshrc (for C shell)
```

where *INSTHOME* is the home directory of the instance that you want to use.

Step c. Stop the DB2 database manager by entering the **db2stop** command from a command line.

Step 3. Repeat steps 2 and 3 for each instance.

Step 3. Remove the Administration Server

You must remove the Administration Server before you remove DB2. To remove the Administration Server, you need to perform the following steps:

Step 1. Log in as user with root authority.

Step 2. Obtain the name of the Administration Server by entering the following command:

```
DB2DIR/bin/dasilist
```

Step 3. Log in as the Administration Server owner.

Step 4. Run the start up script as follows:

```
. INSTHOME/sqllib/db2profile (for Bash, Bourne, or Korn shell)
source INSTHOME/sqllib/db2cshrc (for C shell)
```

where *INSTHOME* is the home directory of the instance.

Step 5. Back up the files in the *ASHOME/sqllib* directory, if needed, where *ASHOME* is the home directory of the Administration Server.

Step 6. Log off.

Step 7. Log in as *root* and remove the Administration Server by entering the following command:

```
DB2DIR/instance/dasidrop ASName
```

where *ASName* represents the name of the administration instance being removed.

The **dasidrop** command removes the */sqllib* directory under the home directory of the Administration Server.

Step 4. Remove DB2 Instances

You can optionally remove some or all of the DB2 Version 6 instances on your system. Once an instance is removed, all the DB2 databases owned by the instance, if any, will not be usable. Remove DB2 instances only if you are not planning to use DB2 Version 6 products, or if you do not want to migrate existing instances to a later version of DB2.

To remove an instance, perform the following steps:

Step 1. Log in as the instance owner.

Step 2. Stop the instance by issuing the **db2stop** command.

Step 3. Back up files in the `INSTHOME/sql1lib` directory, if needed, where `INSTHOME` is the home directory of the instance owner.

You might want to save the database manager configuration file, `db2sysm`, the `db2nodes.cfg` file, or user defined function or fenced stored procedure applications in `INSTHOME/sql1lib/function`, where `INSTHOME` is the home directory of the instance owner.

Step 4. Log off as the instance owner.

Step 5. Log in as user with root authority.

Step 6. Remove the instance by entering the following command:

```
DB2DIR/instance/db2idrop InstName
```

where <i>DB2DIR</i>	= /usr/lpp/db2_06_01	on AIX
	= /opt/IBMdb2/V6.1	on Solaris

and where *InstName* is the login name of the instance.



Because `INSTHOME` is mounted, you only have to run the **db2idrop** command once to remove the instance.

Step 7. Optionally, as a user with root authority, remove the instance owner's user ID and group (if used only for that instance). Do not remove these if you are planning to re-create the instance.

This step is optional since the instance owner and the instance owner group may be used for other purposes.

The **db2idrop** command removes the instance entry from the list of instances and removes the `INSTHOME/sql1lib` directory, where *INSTHOME* is the home directory of the instance.

Step 5. Remove DB2 Products



You must stop all outstanding DB2 processes before removing the DB2 products.

The following steps describe how to remove DB2 products on UNIX operating systems.

Remove DB2 Products on AIX Systems

You can remove the DB2 products on AIX using the **SMIT** interface. To remove a DB2 product, perform the following steps:

- Step 1. Log in as user with root authority.
- Step 2. Type **smit install_remove** to proceed directly to the Remove Software Products screen.
- Step 3. Press **F4** to display a list of the software to remove. Press **F7** at some or all of the entries that have a prefix of **db2_06_01**.
For a detailed description of the software packages, refer to the *Installation and Configuration Supplement*.
- Step 4. Press Enter to start removing the DB2 products.



You can also remove all DB2 Version 6 products on AIX using the **installp** command with the de-install option:

```
installp -u db2_06_01
```

Remove DB2 Products on Linux Systems

You can remove the DB2 products on Linux using the **db2_deinstall** command as follows:

- Step 1. Log in as user with root authority.
- Step 2. Enter the **db2_deinstall -n** command. This command is found on the DB2 Version 6 product CD.

Remove DB2 Products on Solaris Systems

You can remove the DB2 products on the Solaris operating system using the **db2_deinstall** command as follows:

- Step 1. Log in as user with root authority.
- Step 2. Enter the **db2_deinstall -n** command. This command is found on the DB2 Version 6 product CD.

Appendix H. 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 95
 - using 94
- accessing data
 - through DB2 Connect 151
 - through Net.Data or JDBC 154
- accessing multiple servers 99, 149
- accessing servers
 - overview 99
 - TCP/IP 99
- Administering Satellites Guide and Reference 172
- Administration Guide 168
- Administration Server
 - overview 163
 - remove 182
 - start 181
 - stop 181
- Administrative API Reference 169
- ADSM
 - software requirements 24
- Alert Center 160
- APPC
 - software requirements 24, 26, 27
 - support on AIX
 - SNA Server 24
 - support on Solaris
 - SunLink SNA 24, 26
 - supported platforms 24
- APPC, CPl-C and SNA Sense Codes 169
- Application Building Guide 169
- application development
 - using Net.Data or JDBC 154
- Application Development Guide 169
- application programs
 - coordinator node 4
 - migrating to DB2 Enterprise - Extended Edition 4
- authority
 - levels 129

C

- capacity
 - disk storage 7
- catalog node 9

- cataloging
 - databases 104
 - TCP/IP node 103, 104
- CD-ROM
 - installing DB2 Software Developer's Kit 73
 - installing DB2 Universal Database 73
- CLI Guide and Reference 170
- Client Configuration Assistant
 - administering clients
 - communications 88
 - changing privileges 119
 - client profiles 89
 - exporting 96
 - using 95
 - overview 162
 - searching the network 88
 - using access profiles 94
- client profile
 - exporting 96
- clients
 - configuring 99
 - installing 67
 - operating systems
 - supported 148
 - overview 148
- column
 - long-field, partitioning key considerations 10
- Command Center
 - entering DB2 commands 113
 - entering SQL statements 113
 - overview 157
- Command Reference 170
- commands
 - alter table 124
 - db2admin 181
 - db2ckmig 121
 - db2idrop 183
 - db2imigr 121, 128, 129
 - db2licd 121
 - db2start 131
 - db2stop 120, 182, 183
 - db2uiddl 131
 - drop function 124
 - drop trigger 124
 - drop view 124
 - mount 80
- commands (*continued*)
 - rlogin 124
 - smit_install_remove 184
 - SYSADM 129
- communication protocols
 - APPC 24, 25, 26, 27
 - IPX/SPX 24, 25, 26, 27
 - Named Pipes 26, 27
 - NetBIOS 24, 26, 27
 - TCP/IP 24, 25, 26, 27, 99
 - configuring 99
- communications
 - configuring the client 88, 99
 - managing 99
 - Client Configuration Assistant 88
 - Control Center 161
 - fast communication manager 13
 - supported protocols
 - TCP/IP 99
- concurrency
 - improving with logical nodes 11
- configuration parameters
 - setting DB2 119
 - SYSADM_GROUP 119
- configuring
 - TCP/IP 99
- configuring client communications
 - setting configuration parameters 99, 132
 - using the Client Configuration Assistant 88
 - using the command line processor 99
- configuring communications
 - overview 99
 - using Client Configuration Assistant 88
- configuring server communications
 - setting DB2COMM 95
- connecting to the Internet using Net.Data 154
- Connectivity Supplement 170
- Control Center
 - components 158
 - listener daemon
 - requirements 157
 - overview 157

- coordinator node
 - description 4
- cost-based query optimizer
 - description 4

D

- Data Movement Utilities Guide and Reference 170

- database
 - authentication type 122, 128
 - definition 3
 - migration 129
 - naming rules 137
 - partition 3

- database access
 - two-phase commit 150
- database administration tools
 - Control Center 157
 - overview 157

- database alias
 - naming rules 137

- database manager
 - definition 3
 - stopping 182

- database objects
 - naming rules 138

- database partition server
 - definition 3

- DB2 client
 - accessing databases 149
 - installing 67
 - Windows 9x 69
 - Windows NT 69
 - installing on OS/2 73
 - overview 148, 149
 - platforms supported 148
 - software requirements 24

- DB2 Connect
 - overview 146, 151

- DB2 Connect Enterprise Edition for OS/2 and Windows NT Quick Beginnings 173

- DB2 Connect Enterprise Edition for UNIX Quick Beginnings 173

- DB2 Connect Personal Edition Quick Beginnings 170

- DB2 Connect User's Guide 170

- DB2 Data Links Manager for AIX Quick Beginnings 173

- DB2 Data Links Manager for Windows NT Quick Beginnings 173

- DB2 Enterprise - Extended Edition
 - configuration planning 18
 - hardware configuration 6

- DB2 Enterprise - Extended Edition for UNIX Quick Beginnings 172
- DB2 Enterprise - Extended Edition for Windows NT Quick Beginnings 172

- DB2 Enterprise Edition
 - memory requirements 18

- DB2 library
 - books 168
 - Information Center 176
 - language identifier for books 174
 - late-breaking information 175
 - online help 166
 - ordering printed books 179
 - printing PostScript books 178
 - searching online
 - information 178
 - setting up document server 177
 - SmartGuides 165
 - structure of 165
 - viewing online information 175

- DB2 Personal Edition Quick Beginnings 172

- DB2 Query Patroller Administration Guide 173

- DB2 Query Patroller Installation Guide 173

- DB2 Query Patroller User's Guide 173

- db2 terminate command 120

- DB2 Universal Database
 - components included
 - Client Configuration Assistant 162
 - Control Center 157, 158
 - DB2 client 148
 - DB2 Performance Monitor 162
 - DB2 Snapshot Monitor 154
 - Visual Explain 161, 162
 - overview 158
 - platforms supported 157
 - software requirements 22

- db2cclst listener daemon 157

- db2ckmig command 121

- db2fcmdm daemon

- description 13

- db2idrop command 183

- db2imigr command 121, 129

- db2stop command 182, 183

- db2uiddl command 131

- developing applications
 - using Net.Data or JDBC 154

- disk requirements

- client 18

- server 18

- Distributed Computing Environment
 - software requirements 24, 25, 26

E

- Enterprise - Extended Edition

- capacity 7

- storage 7

- Enterprise - Extended Edition configuration
 - logical nodes 11
 - security for an instance 11
 - shared nothing definition 3

- environment variables 12

- db2profile script 75

- error handling

- log 122

- migration 121

- execs

- db2uexit.exe 127

- db2uext2.exe 127

F

- fast communication manager (FCM)

- db2fcmdm daemon 13

- description 13

- files

- db2system 183

- listing for an instance 12

- fixed disks

- hardware requirements 18

G

- GB (gigabyte) 7

- gigabyte (GB) 7

- Glossary 170

H

- hardware requirements

- fixed disk 18

I

- IBMCATGROUP nodegroup 9

- IBMDEFAULTGROUP nodegroup 9

- IBMTEMPGROUP 9

- IBMTEMPGROUP nodegroup 8

- installation

- CID using SystemView LAN 71, 74

- clients

- AIX 29, 51, 52, 75

- HP-UX 75

- OS/2 73

- Solaris 75

- installation (*continued*)
 - Windows 9x 71
 - Windows NT 69
- DB2 server on UNIX systems
 - removing DB2 products 181
 - updating kernel
 - components 54, 75, 76, 77
- disk requirements
 - client 18
 - server 18
- errors 71, 74
- log 71, 74
- memory requirements
 - client 17
 - server 17
- mounting the CD-ROM 77
 - AIX 78, 79
 - HP-UX 79
 - Solaris 55, 80
- remote 81
- servers
 - AIX 29
 - OS/2 73
- Installation and Configuration Supplement 170
- instance owner 12
- instances
 - definition 11
 - listing files in 12
 - multiple on one processor 12
 - naming restrictions 139
 - removing 183
 - SPM instance creation 183
 - stopping 182
- inter-operator parallelism 5
- IPX/SPX
 - software requirements 26
 - verifying the connection 106

J

- Java Support 154
- Journal 160

K

- kernel configuration parameters
 - updating 75

L

- LANG environment variable 135
- listener daemon (db2cclst) 157
- listing
 - files in an instance 12
- logical node
 - improving concurrency 11
 - purpose 10
 - supported configurations 11

- long-field columns
 - partitioning key
 - considerations 10

M

- managing connections
 - client
 - overview 99, 162
 - using the Client Configuration Assistant 88, 162
 - using the command line processor 99
 - server
 - overview 162
 - using the Client Configuration Assistant 88
 - using the Command Line Processor 99
- managing databases using the Control Center 158
- managing server communications
 - overview 161
- memory requirements
 - client 17
 - estimating 17
 - recommended 17
 - server 17
- Message Reference 171
- Microsoft SNA Server
 - version required 27
- migration
 - data 121
 - database 128, 129
 - DB2 Version 2.x 119, 120
 - db2chkmig 121
 - instance 120, 121, 128, 129
 - of applications 4
 - post-installation 120
 - post-migration
 - rebinding packages 132
 - unique index conversion 130
 - updating database and database manager configuration 132
 - updating statistics 131
 - previous releases 119, 120
 - user exit program 126, 127

N

- naming rules
 - database 137
 - database alias 137
 - database objects 138
 - general 137
 - groups 139

- naming rules (*continued*)
 - instance names 137
 - password 143
 - userid 139
 - username 69, 139
 - workstation name (nname) 139
- National Language Support (NLS)
 - code set 135
 - codepage support 135
 - determining codepage 137
- Net.Data
 - overview 154
- node
 - definition 3
 - FCM daemon 13
 - in the same instance 11
- nodegroup
 - IBMCATGROUP 9
 - IBMDEFAULTGROUP 9
 - multipartition nodegroup 7

O

- operating systems
 - client 148

P

- parameters
 - setting kernel configuration 75
 - SYSADM_GROUP 119
- partition
 - definition 3
- partitioning key
 - definition 10
- partitioning map
 - overview 10
- password
 - naming rules 143
- Performance Monitor
 - using 161
- planning
 - DB2 configuration 17
 - DB2 Connect configuration 17
- pools
 - agent 11
- privileges
 - required 119
- product
 - descriptions 146
 - overview 146
- programs
 - db2chkmig 121
- protocols
 - TCP/IP 99

Q

- query
 - cost-based optimizer 4
- Quick Beginnings for OS/2 172
- Quick Beginnings for UNIX 172
- Quick Beginnings for Windows NT 172

R

- removing
 - administration server 182
 - product options 181, 184
 - AIX 184
 - HP-UX 184
- Replication Guide and Reference 171
- restrictions
 - instance name 139

S

- scaling the database system
 - definition 10
- script
 - db2cshrc 182
 - db2profile 182
- Script Center 160
- security
 - for an instance 11
- setting configuration parameters 99
- setting up client communications
 - using the Client Configuration Assistant 88
 - using the command line processor 99
- setting up document server 177
- shared-nothing configuration
 - definition 3
 - storage 7
- Software Developer's Kit
 - overview 163
- software requirements
 - communication protocols 22
 - DB2 client 22, 24, 26, 27
 - DB2 Connect 22
 - DB2 Software Developer's Kit 22, 24, 25, 26, 27
 - DB2 Universal Database 22
 - Net.Data 24, 26, 27
- SQL
 - access plans
 - viewing using Visual Explain 162
- SQL Getting Started 171
- SQL Reference 171

- starting
 - database manager 182
- symmetric multiprocessor architecture 10
- SYSADM
 - privileges
 - controlling 119
- SYSADM_GROUP parameter 119
- system configuration
 - with DB2 Connect 152
 - with DB2 Universal Database 149
- System Monitor Guide and Reference 171

T

- table
 - in multinode nodegroup 10
 - in single-node nodegroup 10
 - maximum size 7
- TB (terabyte) 7
- TCP/IP
 - client 99
 - configuring 99
 - preventing socket collisions 100
 - setting up client-to-server 99
 - software requirements 24, 25, 26
 - troubleshooting 100
- terabyte (TB) 7
- Tools Setting 160
- transparent parallelism 4
- Troubleshooting Guide 171

U

- username
 - naming rules 139
- using the command line processor (CLP)
 - cataloging a database 104
 - cataloging a node 103
 - connecting to a database 106

V

- verifying the connection
 - using the command line processor
 - TCP/IP 99
- Visual Explain
 - overview 162

W

- well-known address
 - description 13
- What's New 171

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: CT6D1NA



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

GC09-2832-00



CT6D1NA

