

IBM[®] DB2[®] Universal Database
Enterprise - Extended Edition
for UNIX[®]



Quick Beginnings

Version 7

IBM[®] DB2[®] Universal Database
Enterprise - Extended Edition
for UNIX[®]



Quick Beginnings

Version 7

Before using this information and the product it supports, be sure to read the general information under "Appendix G. Notices" on page 259.

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

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

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

© **Copyright International Business Machines Corporation 1993, 2000. All rights reserved.**

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

Contents

Welcome to DB2 Universal Database! . . .	vii
Conventions	vii

Part 1. About 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	8
Multiple Logical Nodes	10
Instances	11
Fast Communications Manager	12
High Availability	13

Part 2. Installing DB2 Servers. . . . 15

Chapter 2. Planning for Installation	17
Memory Requirements	17
Disk Requirements	18
DB2 Servers	18
DB2 Clients	18
Software Requirements	19
Server Product Requirements	19
Client Product Requirements	21
Possible Client-to-Server Connectivity Scenarios	26
Migrating from Previous Versions of DB2	27
Preparing Databases and Instances for Migration	27
NetQuestion Search System	29
Your next step	29
Chapter 3. Installing and Configuring DB2 Universal Database on AIX	31
Before You Begin	31
Performing the Installation	37
Post Installation Steps	38
Step 1. Create a DB2 Instance	39

Step 2. Create the Administration Server	40
Step 3. Update the Node Configuration File	40
Step 4. Enable the Fast Communications Manager	43
Step 5. Modify the Environment to run DB2 Enterprise - Extended Edition	45
Step 6. Enable the Execution of Remote Commands	47
Step 7. Create Links for DB2 Files (Optional).	48
Step 8. Start the Administration Server	48
Step 9. Create the SAMPLE Database (Optional).	50
Step 10. Install the License Key	52

Chapter 4. Installing and Configuring DB2 Universal Database on HP-UX 53

Before You Begin	53
Performing the Installation	57
Post Installation Steps.	59
Step 1. Create a DB2 Instance	59
Step 2. Create the Administration Server	60
Step 3. Update the Node Configuration File	61
Step 4. Enable the Fast Communications Manager	63
Step 5. Enable the Execution of Remote Commands	65
Step 6. Create Links for DB2 Files (Optional).	65
Step 7. Start the Administration Server	66
Step 8. Create the SAMPLE Database (Optional).	67
Step 9. Install the License Key	68

Chapter 5. Installing and Configuring DB2 Universal Database on Linux. 71

Before You Begin	71
Performing the Installation	75
Post Installation Steps.	76
Step 1. Create a DB2 Instance	77
Step 2. Create the Administration Server	78
Step 3. Update the Node Configuration File	79

Step 4. Enable the Fast Communications Manager	81	Step 9. Create the SAMPLE Database (Optional)	121
Step 5. Check Available Paging Space	83	Step 10. Install the License Key	122
Step 6. Enable the Execution of Remote Commands	83	Chapter 8. DB2 Post-installation Migration Tasks	125
Step 7. Create Links for DB2 Files (Optional).	84	Migrating Instances	125
Step 8. Start the Administration Server	84	Possible Migration Error Messages and User Responses	126
Step 9. Create the SAMPLE Database (Optional).	85	Migrating Databases	127
Step 10. Install the License Key	86	Optional Post-Migration Tasks	128
Chapter 6. Installing and Configuring DB2 Universal Database for NUMA-Q	89	Chapter 9. Your Next Step	131
Before You Begin	89	DB2 Universal Database Administration Guide.	131
Performing the Installation	92	DB2 Universal Database Administration Guide: Planning	131
Post Installation Steps.	93	DB2 Universal Database Administration Guide: Implementation	132
Step 1. Create a DB2 Instance	94	DB2 Universal Database Administration Guide: Performance	132
Step 2. Create the Administration Server	95	Task Information in other DB2 Documentation	132
Step 3. Update the Node Configuration File	95	<hr/>	
Step 4. Enable the Fast Communications Manager	99	Part 3. Installing DB2 Clients	135
Step 5. Enable the Execution of Remote Commands	100	Chapter 10. Installing DB2 Clients	137
Step 6. Create Links for DB2 Files (Optional)	101	DB2 Run-Time Client	137
Step 7. Start the Administration Server	101	DB2 Administration Client.	138
Step 8. Create the SAMPLE Database (Optional)	102	DB2 Application Development Client	138
Step 9. Install the License Key	103	Distributed Installation	138
Chapter 7. Installing and Configuring DB2 Universal Database on Solaris.	105	DB2 Thin Client	139
Before You Begin	105	Chapter 11. Installing DB2 Clients on Windows 32-Bit Operating Systems	141
Performing the Installation.	111	Before You Begin Installing	141
Post Installation Steps	112	Installing Without Administrator Authority	141
Step 1. Create a DB2 Instance	112	Installation Steps	142
Step 2. Create the Administration Server	114	Chapter 12. Installing DB2 Clients on OS/2 Operating Systems.	145
Step 3. Update the Node Configuration File	114	Before You Begin Installing	145
Step 4. Enable the Fast Communications Manager	117	Installation Steps	145
Step 5. Check Available Paging Space	118	Chapter 13. Installing DB2 Clients on UNIX Operating Systems.	149
Step 6. Enable the Execution of Remote Commands	119	Before You Begin	149
Step 7. Create Links for DB2 Files (Optional)	119	About the db2setup utility.	149
Step 8. Start the Administration Server	120		

Updating Kernel Configuration Parameters	150
HP-UX Kernel Configuration Parameters	150
NUMA-Q/PTX Kernel Configuration Parameters	151
Solaris Kernel Configuration Parameters	152
Installing Your DB2 Client	153
Your next step	154
Chapter 14. Control Center Installation and Configuration	155
Application versus Applet	155
Machine Configurations	156
Supported Java Virtual Machines for the Control Center	157
Setting Up and Working with the Control Center	158
Control Center Services Setup (Applet Mode only)	158
Working with the Control Center	160
Functional Considerations	162
Installation Tips for Control Center Help on UNIX Operating Systems	162
Configuring TCP/IP on OS/2	163
Enabling Local Loopback	163
Enabling Localhost	163
Verifying TCP/IP Configuration on OS/2	164
Troubleshooting Information	164
Administering DB2 for OS/390 and DB2 Connect Enterprise Edition servers with the Control Center	165
Preparing DB2 for OS/390 Servers for the Control Center	166
Working with the Control Center	166
Other Sources of Information	166

Part 4. Configuring DB2 Communications. 169

Chapter 15. Configuring Client-to-Server Communications Using the Client Configuration Assistant	171
LDAP Directory Support Considerations	171
Before You Begin	171
Configuration Steps	172
Adding a Database Using a Profile	173
Adding a Database Using Discovery	174
Adding a Database Manually	176
Creating and Using Profiles	178
Server Profiles	178

Client Profiles	179
-----------------	-----

Chapter 16. Configuring Client-to-Server Communications Using the Command Line Processor	183
Configuring TCP/IP on the Client	183
Step 1. Identify and Record Parameter Values	184
Step 2. Configure the Client	185
Step 3. Test the Client-to-Server Connection	190

Part 5. Appendixes 193

Appendix A. Basic Task Knowledge	195
Starting the Client Configuration Assistant	195
Starting the DB2 Control Center	195
Entering Commands Using the Command Center	196
Entering Commands Using the Command Line Processor	197
DB2 Command Window	198
Interactive Input Mode	198
Working with the System Administrative Group	199
Working with the Business Intelligence Functions	199
Mounting CD-ROMs on UNIX Operating Systems	200
Mounting the CD-ROM on AIX	200
Mounting the CD-ROM on HP-UX	201
Mounting the CD-ROM on Linux	201
Mounting the CD-ROM on PTX	201
Mounting the CD-ROM on Solaris	202
Setting the Number of Licensed Processors	202
Upgrading DB2 from Try and Buy Mode	203

Appendix B. National Language Support (NLS)	205
Language and Codeset Support for UNIX Operating Systems	205
Code Page and Language Support for OS/2 and Windows Operating Environments	205

Appendix C. Naming Rules	207
General Naming Rules	207
Database, Database Alias, and Catalog Node Name Rules	207
Object Name Rules	208

Username, User ID, Group Name, and Instance Name Rules	209
Workstation Name (nname) Rules	209
DB2SYSTEM Naming Rules	211
Password Rules	211

Appendix D. About DB2 Universal Database for UNIX, Windows, and OS/2 . 213

DB2 Products	213
DB2 Everywhere	213
DB2 Universal Database	214
Host Databases	216
DB2 Connect	216
Related Products	217
DB2 Relational Connect	217
DB2 Warehouse Manager	217
DB2 OLAP Server	218
Intelligent Miner	218
DB2 Spatial Extender	218
DB2 Net Search Extender	218
DB2 Data Links Manager	219
Tivoli Enterprise	219
Working with DB2 Data.	219
Accessing DB2 Data from Remote Clients	219
Accessing Multiple DB2 Servers	220
Accessing Host or AS/400 DB2 Data from the Desktop using DB2 Connect Enterprise Edition.	221
Accessing DB2 Data from the Web using Java	225
Accessing DB2 Data from the Web using Net.Data	226
Accessing DB2 Data from Host and AS/400 Client Machines	228
Administering Instances and Databases with the DB2 Administration Tools.	228
Managing Instances and Database Objects using the Control Center	228
Managing Communications on the Server	232

Monitoring Databases using DB2 Performance Monitor	232
Viewing SQL Access Plans using Visual Explain	232
Managing Connections to Databases using the Client Configuration Assistant	233
Managing Warehouses using the Data Warehouse Center	234
Understanding the Administration Server	234
Developing Applications using the DB2 Application Development Client	235
Running Your Own Applications.	236

Appendix E. Using the DB2 Library . . . 237

DB2 PDF Files and Printed Books	237
DB2 Information	237
Printing the PDF Books	246
Ordering the Printed Books	247
DB2 Online Documentation	248
Accessing Online Help	248
Viewing Information Online	250
Using DB2 Wizards	252
Setting Up a Document Server	253
Searching Information Online.	254

Appendix F. Removing DB2 Products . . . 255

Stop the Administration Server	255
Stop all DB2 instances	255
Remove the Administration Server	256
Remove DB2 instances (Optional)	257
Remove DB2 Products	257

Appendix G. Notices 259

Trademarks	262
----------------------	-----

Index 265

Contacting IBM 269

Product Information	269
-------------------------------	-----

Welcome to DB2 Universal Database!

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

This *Quick Beginnings* book will guide you through the planning, installation, migration (if necessary), and setup of a partitioned database system. 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 or the Command Line Processor.



Conventions

This book uses the following highlighting conventions:

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



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



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

For a complete description of the DB2 library, see “Appendix E. Using the DB2 Library” on page 237.



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

Part 1. About 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 using system resources, including CPU, memory, disk space, and communications. In a partitioned database system, a single database manager and the collection of data and system resources that it manages are referred to collectively as a *database partition server (node)*. A *partitioned database system* is the collection of all the database partition servers that you create to handle data requests.

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

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

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

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

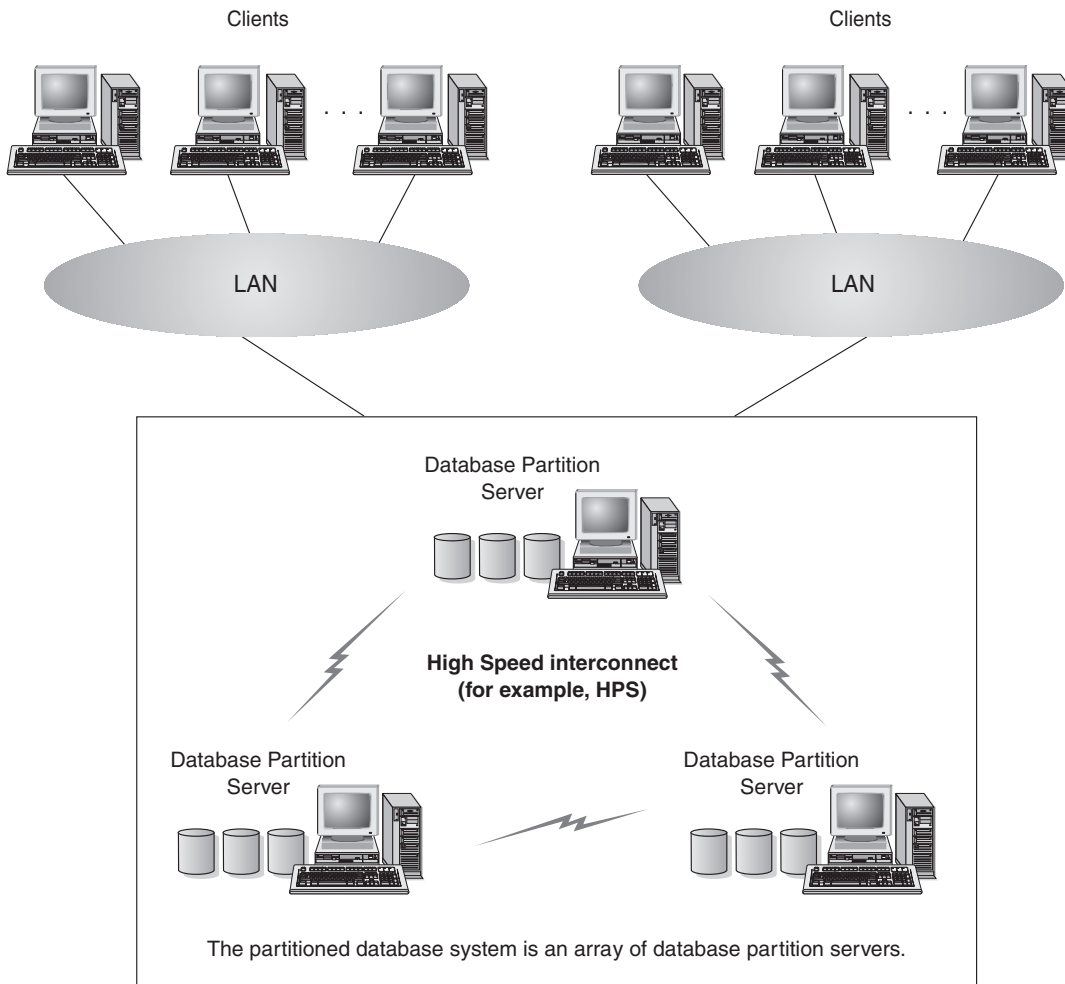


Figure 1. DB2 Enterprise - Extended Edition Hardware Configuration

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, High Performance Switch (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 determined 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 System/6000 Scalable POWER Parallel Systems (RS/6000 SP) environment that is running AIX, the number of database partition servers is only limited by the possible size of an AIX RISC System/6000 SP system.

For an HP-UX 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. For example, 24 database partition servers could run on a cluster of 4 K580 Enterprise Servers with 6 CPUs each.

In a Linux Environment, the number of database partition servers is limited by the size of the machines and the number of machines clustered together. For example, in a configuration that has two database partition servers per machine, 40 database partition servers could be run on a clustered system of 20 Linux machines.

In a PTX environment, the number of database partition servers is limited by the number of quads in a machine. We recommend that you run one database partition server per NUMA-Q quad. For example, five multiple logical nodes on a five quad system, with each logical node having four processors.

In a Solaris** Operating 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 512 GB (gigabytes) per database partition. For example, in a database that has 128 partitions, the maximum size of one table is approximately 64 TB (terabytes).

Nodegroups and Data Partitioning

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

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

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

The IBMDEFAULTGROUP nodegroup contains all the database partitions for the database. When you create a database, a database partition is created at each database partition server (node) 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, nodegroup 2 is a single-partition nodegroup, and nodegroup 3 is a multipartition nodegroup.




Figure 2. Nodegroups in a Database

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

You can drop database partitions from a nodegroup, or if 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 more than one database partition server assigned to each machine. If these database partition servers (nodes) participate in the *same* instance, this is referred to as a Multiple Logical Node (MLN) configuration.

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

For more information about setting up logical nodes, 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, which shows two separate instances. Instance 1 contains six database partition servers and Instance 2 contains eight database partition servers. (Multiple database partition servers are indicated when more than one line is shown between a database partition server and the instance directory.) The two instances overlap, but this is due to the assignment of two database partition servers to each of the three machines in the middle of the figure.

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

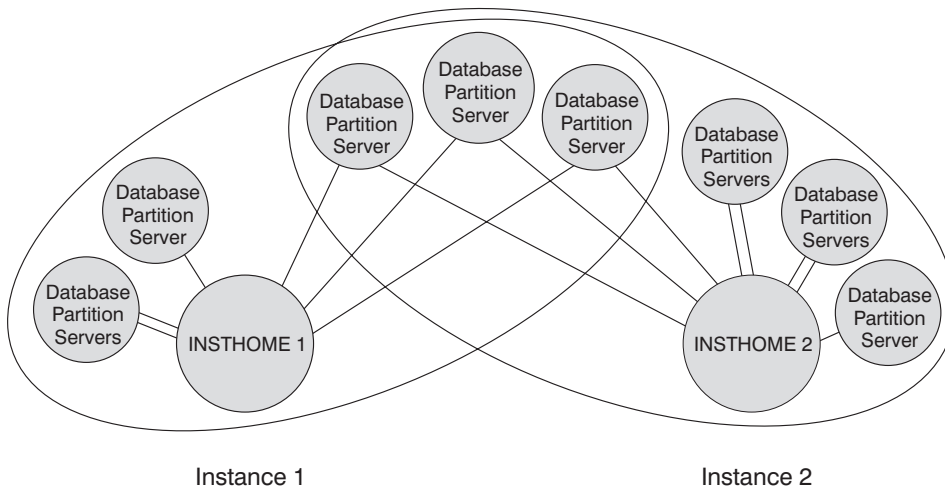


Figure 3. Two Instances

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 7
- 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 System Administrative (*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.

There is a one-to-one relationship 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 Communications Manager

The Fast Communications 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/sql11ib/db2nodes.cfg*, where *INSTHOME* is the home directory of the instance owner) and defines a well-known address to use for 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 `fcf_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 or software failure. HACMP provides increased availability through a cluster of processors which share resources such as disks or network access.

On Solaris systems, you implement failover support using Sun Cluster 2.2. Sun Cluster 2.2 performs both failure detection and the restarting of resources in a clustered environment, as well as failover support for physical disks and IP addresses.

At this time, DB2 failover support for HP-UX, Linux and the PTX operating system is a manual process requiring you to restart the failing node manually on another node that has access to the failing node's disk.

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

Part 2. Installing DB2 Servers

Chapter 2. Planning for Installation

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

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

- "Memory Requirements".
- "Disk Requirements" on page 18.
- "Software Requirements" on page 19.
- "Possible Client-to-Server Connectivity Scenarios" on page 26.
- "Migrating from Previous Versions of DB2" on page 27.



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 31
- "Chapter 4. Installing and Configuring DB2 Universal Database on HP-UX" on page 53
- "Chapter 5. Installing and Configuring DB2 Universal Database on Linux" on page 71
- "Chapter 6. Installing and Configuring DB2 Universal Database for NUMA-Q" on page 89
- "Chapter 7. Installing and Configuring DB2 Universal Database on Solaris" on page 105

For information on the DB2 family of products, see "Appendix D. About DB2 Universal Database for UNIX, Windows, and OS/2" on page 213.

Memory Requirements

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

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

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

Disk Requirements

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

DB2 Servers

A default installation of DB2 requires a minimum of 250 to 300 MB of disk space. This amount includes the online production documentation and tools. On NUMA-Q systems, a default installation requires a minimum of 80 MB.

DB2 Clients

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

Table 1. Disk Requirements for Client Components

Client Components	Recommended Minimum Disk (MB)
OS/2	
DB2 Run-Time Client	30 MB
DB2 Application Development Client	125 MB, not including the Java Development Kit (JDK)
DB2 Administration Client	95 MB
UNIX platforms	
DB2 Run-Time Client	30 to 40 MB (70 MB for Silicon Graphics IRIX)
DB2 Application Development Client	90 to 120 MB, not including the JDK (40 MB for NUMA-Q)
DB2 Administration Client	80 to 110 MB

Table 1. Disk Requirements for Client Components (continued)

Client Components	Recommended Minimum Disk (MB)
Note: PTX/NUMA-Q and Silicon Graphics IRIX operating systems do not support the DB2 Administration Client.	
Windows 32-bit Operating Systems	
DB2 Run-Time Client	25 MB
DB2 Application Development Client	325 MB, including the JDK
DB2 Administration Client	125 MB

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

Software Requirements

This section outlines the software required to run DB2 products.

Server Product Requirements

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

On all platforms you will need a Java Runtime Environment (JRE) Version 1.1.8 to run DB2's Java-based tools, such as the Control Center. If you intend to run the Control Center as an applet, you will need a Java-enabled browser. See "Chapter 14. Control Center Installation and Configuration" on page 155 for more information.

Table 2. Software Requirements

Hardware/Software Requirements	Communications
	DB2 Universal Database for AIX

Table 2. Software Requirements (continued)

Hardware/Software Requirements	Communications
RISC System/6000 and 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 direct addressing. – AIX base operating system 4.3 or later (contains Novell Netware Services for AIX Version 4.1), which supports 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 V5.0.3 – Bull DPX/20 SNA/20 <p>Notes:</p> <ol style="list-style-type: none"> 1. If you plan to use the DB2 OLAP Starter Kit, you require AIX Version 4.3 or later. 2. For LDAP (Lightweight Directory Access Protocol) support, you require an IBM SecureWay Directory Client V3.1.1 running on AIX V4.3.1 or higher. 3. For Data Links Manager in DCE-DFS environments, you require DCE Version 3.1. 4. If you plan to use DCE (Distributed Computing Environment) and you are not using DB2 Data Links Manager, you require a DCE product that is provided with your AIX operating system. 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. 5. If you plan to use the Tivoli Storage Manager facilities back up and restore to your databases, you require the Tivoli Storage Manager Client Version 3 or later. 6. If you plan to use the Simple Network Management Protocol (SNMP) subagent, you require DPI 2.0 provided by IBM SystemView Agent.
DB2 Universal Database 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"> • TCP/IP is provided with the HP-UX base operating system. • For APPC connectivity, HP-UX Version 11.00 requires the following: <ul style="list-style-type: none"> – SNAplus2 Link R6.11.00.00 – SNAplus2 API R.6.11.00.00 <p>Notes:</p> <ol style="list-style-type: none"> 1. HP-UX only supports outbound client APPC requests. It does not provide support for inbound client APPC requests. 2. If you plan to use DCE (Distributed Computing Environment), you require a DCE product that is provided by the HP-UX Version 11 operating system. With DB2 Connect, you must install DCE Directory Services on the client and the host server. You do not need DCE installed on a DB2 Connect Enterprise Edition server. 3. If you plan to use the Tivoli Storage Manager facilities to back up and restore your databases, you require the Tivoli Storage Manager Client Version 3 or later.
DB2 Universal Database for Linux	

Table 2. Software Requirements (continued)

Hardware/Software Requirements	Communications
<ul style="list-style-type: none"> Linux kernel 2.2.12 or higher; <i>glibc</i> Version 2.1.2 or higher; <i>pdksh</i> package 5.2 or greater (required to run the DB2 command line processor); and <i>libstdc++</i> Version 2.9.0. <p>To install DB2, you will need <i>rpm</i> 3.0 or greater.</p>	<p>TCP/IP</p> <ul style="list-style-type: none"> For TCP/IP connectivity, no additional software is required. APPC connectivity is not supported in this version of DB2 UDB for Linux. To run the Control Center, or any DB2 Java utility you must install the IBM Developer Kit for Java March 22, 2000 release or later.
DB2 Universal Database for NUMA-Q	
<ul style="list-style-type: none"> NUMA-Q system running PTX Version 4.5 or later. <i>ptx</i>/EFS v1.4.0 with <i>templog</i> is required. 	<p>TCP/IP</p> <ul style="list-style-type: none"> For TCP/IP connectivity, no additional software is required.
DB2 Universal Database for the Solaris Operating Environment	
<p>Solaris SPARC-based computer and the following:</p> <ul style="list-style-type: none"> Solaris Version 2.6 or later. <p>The following patches are required for Solaris version 2.6:</p> <ul style="list-style-type: none"> 105181-17 or higher 105210-25 or higher 105568-12 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 APPC connectivity, you require SunLink SNA 9.1 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"> If you plan to use DCE (Distributed Computing Environment), you require Transarc DCE Version 2.0 or later. 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. If you plan to use the Tivoli Storage Manager facilities for backup and restore of your databases, you require the Tivoli Storage Manager Client Version 3 or later.

Client Product Requirements

Table 3 on page 22 lists the software requirements needed for a DB2 Administration Client, DB2 Run-Time Client, or a DB2 Application Development Client.

On all platforms you will need a Java Runtime Environment (JRE) Version 1.1.8 to run the DB2 tools, such as the Control Center. If you intend to run the Control Center as an applet on Windows 32-bit or OS/2 systems, you will need a Java-enabled browser. See “Chapter 14. Control Center Installation and Configuration” on page 155 for more information.

Table 3. 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 Application Development Client for AIX 	<p>RISC System/6000 and the following:</p> <ul style="list-style-type: none"> AIX Version 4.2 or later For OLAP Starter Kit, AIX Version 4.3 or later. For LDAP (Lightweight Directory Access Protocol) support, you require an IBM SecureWay Directory Client V3.1.1 running on AIX V4.3.1 or later. For Warehouse Agent, bos.iconv.ucs.com and bos.iconv.ucs.pc at AIX Version 4.2 or later. For Data Links Manager in DCE-DFS environments, you require DCE Version 3.1. <p>Note: When the DB2 Application Development Client is installed, the JDK 1.1.8 is only installed if no other version of the JDK is detected.</p>	<p>APPC or TCP/IP</p> <ul style="list-style-type: none"> For APPC connectivity, you require IBM eNetwork Communications Server Version 5.0.3 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) and you are not using DB2 Data Links Manager, you require a DCE product that is provided by the AIX base operating system.</p>
<ul style="list-style-type: none"> DB2 Run-Time Client for HP-UX DB2 Administration Client for HP-UX DB2 Application Development Client 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>Note: When the DB2 Application Development Client is installed, the JDK is not installed. Contact your operating system vendor for the latest version of the JDK.</p>	<p>APPC or TCP/IP</p> <ul style="list-style-type: none"> TCP/IP is provided with the HP-UX base operating system. 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>

Table 3. Software Requirements for Clients (continued)

Component	Hardware/Software Requirements	Communications
<ul style="list-style-type: none"> • DB2 Run-Time Client for Linux • DB2 Administration Client for Linux • DB2 Application Development Client for Linux 	<ul style="list-style-type: none"> • Linux kernel 2.2.12 or higher; • <i>glibc</i> Version 2.1.2 or higher; • <i>pdksh</i> package (required to run the DB2 command line processor); and • <i>libstdc++</i> Version 2.9.0. <p>To install DB2, you will need <i>rpm</i>.</p> <p>Note: When the DB2 Application Development Client is installed, the JDK is not installed. Contact your operating system vendor for the latest version of the JDK.</p>	<p>TCP/IP</p> <ul style="list-style-type: none"> • The Linux base operating system provides TCP/IP connectivity, if selected during installation.
<ul style="list-style-type: none"> • DB2 Run-Time Client for OS/2 • DB2 Administration Client for OS/2 • DB2 Application Development Client for OS/2 	<ul style="list-style-type: none"> • OS/2 Warp Version 4 • OS/2 Warp Server Version 4 • OS/2 Warp Server Advanced V4 • OS/2 Warp Server Advanced V4 with SMP Feature • OS/2 Warp Server for e-business <p>Note: When the DB2 Application Development Client is installed, the JDK is not installed. You can install the latest version of the JDK from your product CD-ROM.</p>	<p>APPC, IPX/SPX, NetBIOS, or TCP/IP</p> <ul style="list-style-type: none"> • For APPC connectivity, you require IBM eNetwork Communications Server for OS/2 Warp Version 5 or IBM eNetwork Personal Communications for OS/2 Warp Version 4.2. • For IPX/SPX connectivity, you require the Novell NetWare client for OS/2 Version 2.10 or later. IPX/SPX can only be used to connect to databases. It cannot be used to connect to host or AS/400 databases. • The OS/2 base operating system provides NetBIOS and TCP/IP connectivity, if selected during installation. • The OS/2 base operating system provides Named Pipes (Local) connectivity. Named Pipes is supported in DOS and WIN-OS/2 sessions. <p>Notes:</p> <ol style="list-style-type: none"> 1. Net.Data requires a Web server such as WebSphere. 2. For DCE Cell Directory Services Support (CDS) for DB2 Clients for OS/2, you must install an IBM Distributed Computing Environment Cell Directory Service client on each client workstation. 3. If you are planning to use Tivoli Storage Manager, PTF 3 for Tivoli Storage Manager Version 3 is required for an OS/2 client.

Table 3. Software Requirements for Clients (continued)

Component	Hardware/Software Requirements	Communications
<ul style="list-style-type: none"> • DB2 Run-Time Client for NUMA-Q • DB2 Application Development Client for NUMA-Q 	<ul style="list-style-type: none"> • NUMA-Q system running PTX Version 4.5 or later. • ptx/EFS v1.4.0 with templog is required. <p>Note: When the DB2 Application Development Client is installed, the JDK is not installed. Contact your operating system vendor for the latest version of the JDK.</p>	<p>TCP/IP</p> <ul style="list-style-type: none"> • For TCP/IP connectivity, no additional software is required.
<ul style="list-style-type: none"> • DB2 Run-Time Client for Silicon Graphics IRIX • DB2 Application Development Client 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 <p>The following patches are required for Versions 6.2 and 6.3:</p> <ul style="list-style-type: none"> – 2791.0 – 3778.0 <p>Note: When the DB2 Application Development Client is installed, the JDK is not installed. Contact your operating system vendor for the latest version of the JDK.</p>	<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 Application Development Client for Solaris 	<p>Solaris SPARC-based computer and the following:</p> <ul style="list-style-type: none"> • Solaris Version 2.6 or later. <p>The following patches are required for Solaris version 2.6:</p> <ul style="list-style-type: none"> • 105568 – 12 or higher • 105210 – 25 or higher • 105181 – 17 or higher <p>Note: When the DB2 Application Development Client is installed, the JDK is not installed. Contact your operating system vendor for the latest version of the JDK.</p>	<p>APPC or TCP/IP</p> <ul style="list-style-type: none"> • For APPC connectivity, you require SunLink SNA 9.1 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. • If you plan to use DCE (Distributed Computing Environment), you require Transarc DCE Version 2.0 or later.

Table 3. Software Requirements for Clients (continued)

Component	Hardware/Software Requirements	Communications
<ul style="list-style-type: none"> • DB2 Run-Time Client for Windows 9x • DB2 Administration Client for Windows 9x • DB2 Application Development Client for Windows 9x 	<ul style="list-style-type: none"> • Windows 95 4.00.950 or later • Windows 98 <p data-bbox="360 348 615 453">Note: When the DB2 Application Development Client is installed, the JDK 1.1.8 is installed.</p>	<p data-bbox="653 236 1085 256">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 and Windows 2000 servers. • If you plan to use LDAP (Lightweight Directory Access Protocol), you require either a Microsoft LDAP client or an IBM SecureWay LDAP client V3.1.1. For more information, refer to the <i>Administration Guide</i>. • If you plan to use the Tivoli Storage Manager facilities for backup and restore of your databases, you require the Tivoli Storage Manager Client Version 3 or later. • If you have the IBM Antivirus program installed on your operating system, it must be disabled or uninstalled to complete a DB2 installation.
<ul style="list-style-type: none"> • DB2 Run-Time Client for Windows • DB2 Administration Client for Windows • DB2 Application Development Client for Windows 	<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) • Windows 2000 <p data-bbox="360 947 615 1052">Note: When the DB2 Application Development Client is installed, the JDK 1.1.8 is installed.</p>	<p data-bbox="653 722 1157 743">APPC, IPX/SPX, Named Pipes, NetBIOS, or TCP/IP</p> <ul style="list-style-type: none"> • The Windows NT and Windows 2000 base operating systems provide 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 V5.01 or later. – Windows 2000: IBM eNetwork Personal Communications for Windows Version 4.3 CSD2 or later – Windows NT: IBM eNetwork Personal Communications for Windows Version 4.2 or later – Microsoft SNA Server Version 3 Service Pack 3 or later – Wall Data Rumba • If you plan to use DCE (Distributed Computing Environment) and connect to a DB2 for OS/390 V5.1 database, that database must be enabled for DCE support using OS/390 DCE Base Services Version 3. • If you plan to use LDAP (Lightweight Directory Access Protocol), you require either a Microsoft LDAP client or an IBM SecureWay LDAP client V3.1.1. For more information, refer to the <i>Administration Guide</i>. • If you plan to use the Tivoli Storage Manager facilities for backup and restore of your databases, you require the Tivoli Storage Manager Client Version 3 or later. • If you have the IBM Antivirus program installed on your operating system, it must be disabled or uninstalled to complete a DB2 installation.

Possible Client-to-Server Connectivity Scenarios

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

Table 4. Possible Client-to-Server Connectivity Scenarios

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

1. Direct Addressing
2. File Server Addressing

Migrating from Previous Versions of DB2

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

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

If you are migrating a DB2 installation on a workstation that is running on a supported Linux distribution, you should be aware of the following:

- The migration of DB2 Version 5.2 Beta is not supported.
- DB2 Version 7 can only run with `glibc` version 2.1.2 or higher, as mentioned in “Software Requirements” on page 19. If your Linux distribution has a previous version of this library, you will not be able to use your DB2 product.

To prepare your workstation for installation, you have to upgrade your existing Linux distribution to a more recent level that includes the required (or later) version of the `glibc` library. You also have the option to update only the local version of the `glibc` library. We recommend that you update your distribution’s installation to a more recent version that includes the required level for the `glibc` library. Updating just the `glibc` library without updating your distribution’s installation version is not recommended as errors during this process can affect your operating system.

When you migrate from a Version 5.x or Version 6 database system to a DB2 Enterprise - Extended Edition Version 7 database system, depending on the installation option you choose, the installation program will migrate all instances to the Version 7 multipartition or single-partition format. If you are migrating a database with populated tables, refer to the *Administration Guide* for descriptions of how to add database partition servers to a system and redistribute data across database partitions.

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

Preparing Databases and Instances for Migration

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

You must be logged in as the instance owner to complete these steps.

To prepare your databases for migration:

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

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

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

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

```
db2 list database directory
```
3. Make a backup copy of all Version 5.x databases. You do not have to back up Version 6 databases. Refer to the *Administration Guide* for your DB2 product for information on making a backup copy of a database and to the *Command Reference* for the syntax of the backup command.
4. When all applications are complete and you have backed up your databases, stop all database server processes owned by the DB2 instance by entering the **db2stop** command.
5. Stop the DB2 license daemon by entering the **db2licd -end** command.
6. Stop all command line processor sessions by entering the **db2 terminate** command in each session that was running the command line processor.
7. Ensure that the **db2profile** (bash, Bourne, or Korn shells) or **db2cshrc** (C shell) instance environment setup script, under the instance's `INSTHOME/sql1lib/` directory, is of the proper shell syntax.

If necessary, ensure that each export statement is separated into different lines. For example:

```
DB2INSTANCE=db2inst1      //bash, Bourne, or Korn shells  
export DB2INSTANCE  
  
set DB2INSTANCE=db2inst1  //C shell
```

To update a DB2 Enterprise Edition Version 7 database system to a DB2 Enterprise – Extended Edition Version 7 system, you must update your instance using the **db2iupdt** command. For more information, refer to the *Administration Guide*.

NetQuestion Search System

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

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

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

Your next step

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

- “Part 2. Installing DB2 Servers” on page 15 for interactive installation.
- The *Installation and Configuration Supplement* for distributed installation.

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 7. See “Migrating from Previous Versions of DB2” on page 27 for information.

This section describes how to install and configure an AIX-based partitioned database system. If you want to install a DB2 client, go to “Chapter 10. Installing DB2 Clients” on page 137. 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 environment variable. To update this list, enter the following command:

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

where `path` is the location where `eeelist.txt` was created, and `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:

```
dsh -a 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:

```
dsh -a 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 is 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 **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 31.



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 **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 users' 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 **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 C. Naming Rules" on page 207.

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 or NIS+, groups and users must be created on the NIS server before running `db2setup`.

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
```



If you are using NIS or NIS+, secondary groups must be created for the DB2 instance owner and the Administration Server on the NIS server. You must then add the primary group of the instance owner to the secondary Administration Server group. Likewise, you must add the primary Administration Server group to the secondary group for the instance owner.

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
  core=-1 data=491519 stack=32767 rss=-1 fsize=-1 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_07_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_07_01.xlic
```

Step 7. The `db2_07_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. To install any non-English DB2 messages, 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 **OK**.
- Step f. Click **List** 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 to this method to install any other DB2 components that were not installed by default.

- Step h. Click **OK**.

If you installed a DB2 Product Library (HTML) fileset, you must run the `/usr/lpp/db2_07_01/doc/db2insthtml locale` command to uncompress and untar the HTML documentation, where *locale* represents the language identifier of the installed fileset. If you installed any non-English DB2 Product Library filesets, run this command for the English fileset first, then re-run this command for every installed non-English fileset locale.

- Step 8. Enter the `lspp -l | grep db2_07_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.



DB2 provides two environment profiles in the `sqllib` directory:

- `db2profile` for `sh`, `bash`, and `ksh`
- `db2cshrc` for `csh`

In Version 7, these environment profiles can no longer be modified.

To modify the DB2 environment in Version 7, first create, then modify the following profiles: `sqllib/userprofile` and `sqllib/usercshrc`. These environment profiles will be called after the DB2 environment is initialized.

Step 1. Create a DB2 Instance



If you are using NIS or NIS+, before you create an instance, there must be an entry for the instance in the `etc/services` file. For example, if you want to create an instance for the user `db2inst1`, you require an entry similar to the following:

```
DB2_db2inst1    60000/tcp
```

This section describes how to create an instance using the DB2 Installer program. We recommend that you use 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 the RS/6000 SP workstation where you installed DB2 Enterprise - Extended Edition (DB2 EEE).
- Step 2. Enter the `/usr/lpp/db2_07_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 31.
- 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 31.
- 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 the RS/6000 SP workstation where you installed DB2 Enterprise - Extended Edition (DB2 EEE).
- Step 2. Enter the **/usr/lpp/db2_07_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.
- 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 31.
- 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/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
```

`nodenum`, `hostname`, `logical port`, and `netname` are defined as follows:

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

To scale your partitioned database system, you add an entry for each database partition server to the `db2nodes.cfg` file. The *nodenum* value 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 a Multiple Logical Node (MLN) and wish to keep the nodes logically grouped in this file.

This entry is required.

hostname The TCP/IP hostname of the database partition server for use by the FCM.

This entry is required.

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.

In this case, if you specify an entry for the *netname* field (see below), the entry *must* be set to 0.

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

Furthermore, if you specify a *logical port* entry for one database partition server, you must specify a *logical port* for each database partition server listed in your `db2nodes.cfg` file.

This field is only optional if 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 communication.

If an entry is specified for this field, all communication 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 are not using a clustered environment and want to have four database partition servers on one physical workstation called `workstation1`, update the `db2nodes.cfg` file as follows:

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

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

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

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

```
4 workstation1 0
6 workstation1 1
8 workstation1 2
9 workstation2 0
```

If you want your partitioned database system to contain two physical workstations, called `workstation1` and `workstation2` (with `workstation2` running two database partition servers), and use a high speed interconnect called `switch1` and `switch2`, 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:

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 199.
2. Ensure that 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` file is locked when the instance is running and can only be edited when the instance is stopped.

3. Edit the `db2nodes.cfg` file and add an entry for each database partition server that will participate in your partitioned database system.
4. Enter the `INSTHOME/sqllib/adm/db2start` command, where `INSTHOME` is the home directory of the instance owner.
5. Log out.

For more information on the `db2nodes.cfg` file, see 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 communication between the database partition servers that participate in your partitioned database system. Communication 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 `db2setup` utility, 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 for any instance that was created using the **db2setup** utility.

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

```
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 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
```

where:

- *instance_name* is the name of the multipartition instance.
- *port_number* is the port number of 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 four database partitions, you will need to specify four consecutive ports as your port range. In our example, the entry in the `/etc/services` file is:

```
DB2_db2inst1      60000/tcp
DB2_db2inst1_END  60003/tcp
```



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 60003/tcp # that is running 4 MLNs.
```

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

1. Log on as root to the workstation where you created a DB2 instance.
2. Open the `/etc/services` file and verify that the port or port range entries that you have reserved for FCM communication are there.
3. Log out.

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.
5. Log out.

For more information on FCM communication, see 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      = 65536
sb_max       = 1310720
rfc1323      = 1
tcp_sendspace = 221184
tcp_recvspace = 221184
udp_sendspace = 65536
udp_recvspace = 655360
ipqmaxlen    = 250
somaxconn    = 1024
```



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 -o 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 221184 on every workstation, enter the following command:

```
dsh -a no -o tcp_sendspace=221184
```



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 *rpoolsize* for *css0* to the following values:

```
spoolsize    16777216
rpoolsize    16777216
```



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=16777216
dsh -a /usr/lpp/ssp/css/chgcss -l css0 -a rpoolsize=16777216
```

If you are not using the `/tftpboot/tuning.cst` file to tune your system, you can use the `/usr/lpp/db2_07_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:

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


```
cp /usr/lpp/db2_07_01/misc/rc.local.sample /etc/rc.local
chown root:sys /etc/rc.local
chmod 744 /etc/rc.local
```
- Review the `/etc/rc.local` file and update it if necessary.
- 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 **mkkitab** command to add an entry to the `/etc/inittab` file. To add this entry, enter the following command:

```
mkkitab "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:

```
l sitab 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:

```
l sps -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 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* file or the */etc/hosts.equiv* file, see your Linux 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 system.

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_07_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_07_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

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 system in your partitioned database system.

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/sqllib/db2nodes.cfg
db2_all "DASINSTHOME/sqllib/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/sqllib/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)

To verify that DB2 Enterprise - Extended Edition (DB2 EEE) is installed and configured correctly, create the **SAMPLE** database on your system, and access data from the database. For the purposes of testing the installation and configuration of DB2 EEE, you should only create the **SAMPLE** database for your instance after you have defined all of the database partition servers that will participate in your partitioned database system. If you want to add a database partition server to an instance where a database has already been created, there are other activities that you need to perform. For more information, see the *Administration Guide*.



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

For more information on how to create a file system, see your Linux documentation.

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

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 199.
2. Enter the following command to create the SAMPLE database:

```
INSTHOME/sql1lib/bin/db2sampl 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/db2sampl /database
```

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

3. Start the database manager by entering the **db2start** command.
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 196 or “Entering Commands Using the Command Line Processor” on page 197.



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 `node1ock` file, located in the `/var/iform` directory.
3. Update your DB2 product license by entering the following command:

```
/usr/lpp/db2_07_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 `db2udbee.lic`.

For example, if the CD-ROM is mounted in the `/cdrom` directory and the name of the license file is `db2udbee.lic`, you would enter the following command:

```
/usr/lpp/db2_07_01/cfg/db2licm /cdrom/db2/license/db2udbee.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 DB2 Control Center” on page 195.
- 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 HP-UX

This section describes how to install and configure a Hewlett-Packard based partitioned database system. If you want to install a DB2 client, go to “Chapter 10. Installing DB2 Clients” on page 137. 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 install 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 7. See “Migrating from Previous Versions of DB2” on page 27 for information.

Before You Begin



The instructions in this section assume that you are installing DB2 Enterprise - Extended Edition (DB2 EEE) for HP-UX on one system and running Multiple Logical Nodes (MLNs) to create your partitioned database system. If you are installing DB2 EEE on a cluster, you will have to run the DB2 Installer program on each physical machine that you want to participate in 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/fstab` file so that this file system is mounted each time the system is rebooted.



If you are installing DB2 EEE on a cluster of HP-UX systems, you must add an entry to the `/etc/exports` 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.



If you are using NIS or NIS+, groups and users must be created on the NIS server before running `db2setup`.

The usernames you create must conform to both your operating system's naming rules, and those of DB2. You can have the DB2 Installer create these usernames, or you can create them manually. For more information on naming rules, see "Appendix C. Naming Rules" on page 207.

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
```



If you are using NIS or NIS+, secondary groups must be created for the DB2 instance owner and the Administration Server on the NIS server. You must then add the primary group of the instance owner to the secondary Administration Server group. Likewise, you must add the primary Administration Server group, to the secondary group for the instance owner.

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:



The HP-UX operating system automatically reboots after changing the kernel configuration parameter values.

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

Step b. Update the kernel configuration parameters as follows:

Table 5. HP-UX Kernel Configuration Parameters (Recommended Values)

Kernel Parameter	Physical Memory			
	64MB - 128MB	128MB - 256MB	256MB – 512MB	512MB+
maxuprc	256	384	512	1500
maxfiles	256	256	256	256
nproc	512	768	1024	2048
nflocks	2048	4096	8192	8192
ninode	512	1024	2048	2048
nfile	(4 * ninode)	(4 * ninode)	(4 * ninode)	(4 * ninode)
msgseg	8192	16384	32767 (1)	32767 (1)
msgmb	65535 (2)	65535 (2)	65535 (2)	65535 (2)
msgmax	65535 (2)	65535 (2)	65535 (2)	65535 (2)
msgtql	256	512	1024	2048
msgmap	130	258	258	2050
msgmni	128	256	256	1024
msgssz	16	16	16	16
semnmi	128	256	512	2048
semmap	130	258	514	2050
semnms	256	512	1024	4096
semnmu	256	512	1024	1024
shmmax	67108864	134217728 (3)	268435456 (3)	268435456 (3)
shmseg	16	16	16	16
shmmni	300	300	300	1000

Notes:

- 1) The msgseg parameter must be set no higher than 32767.
- 2) The msgmb and msgmax parameters must be set to at least 65 535.
- 3) The shmmax parameter should be set to 134 217 728 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 shmmax to 184968806 (196*1024*1024*0.9).
- 4) To maintain the interdependency among kernel parameters, change parameters in the same sequence in which they appear in the preceding table.

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.

Performing the Installation

This section describes how to install DB2 EEE on an HP-UX system or an HP-UX 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 EEE, perform the following steps:

- Step 1. Log on to the system as a user with root authority.
- Step 2. Insert the appropriate CD-ROM into the drive.
- Step 3. Mount the CD-ROM. Because DB2 Version 7 for HP-UX contains several files with long file names, the mount command may fail. Perform the following steps to avoid any mounting problems when mount your CD-ROM:
 - a. In the */etc* directory, add the following line to the *pfs_fstab* file:


```
/dev/dsk/c0t2d0 mount_point pfs-rrip ro,hard
```

where *mount_point* is the mount point of the CD-ROM.
 - b. Start the *pfs* daemon by entering the following commands (if they are not already running):


```
/usr/sbin/pfs_mountd &  
/usr/sbin/pfsd 4 &
```

- c. Enter the following commands to create a directory called `/cdrom` and mount the CD-ROM on this directory:

```
mkdir /cdrom
/usr/sbin/pfs_mount /cdrom
```

where `/cdrom` represents the mount point of the CD-ROM.

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 4. 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.

- Step 5. Enter the `./db2setup` command to start the DB2 install program. After a few moments, the **Install DB2 V7** window opens.

- Step 6. From the product list on the *Install DB2 V7* 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 select or deselect optional components for a DB2 product that you want to install, select the **Customize** option. To go back to a previous window 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/V7.1` directory.

Post Installation Steps

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



DB2 provides two environment profiles in the `sqllib` directory:

- `db2profile` for `sh`, `bash`, and `ksh`
- `db2cshrc` for `csh`

In Version 7, these environment profiles can no longer be modified.

To modify the DB2 environment in Version 7, first create, then modify the following profiles: `sqllib/userprofile` and `sqllib/usercshrc`. These environment profiles will be called after the DB2 environment is initialized.

Step 1. Create a DB2 Instance



If you are using NIS or NIS+, before you create an instance, there must be an entry for the instance in the `etc/services` file. For example, if you want to create an instance for the user `db2inst1`, you require an entry similar to the following:

```
DB2_db2inst1    60000/tcp
```

This section describes how to create an instance using the DB2 Installer program. We recommend that you use 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 on a machine where you installed DB2 Enterprise - Extended Edition (DB2 EEE).
- Step 2.** Enter the `/opt/IBMdb2/V7.1/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.
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 31.
- 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 31.
- 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 a machine where you installed DB2 Enterprise - Extended Edition (DB2 EEE).
- Step 2. Enter the `/opt/IBMdb2/V7.1/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 31.

- 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
```

`nodenum`, `hostname`, `logical port`, and `netname` are defined as follows:

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

To scale your partitioned database system, you add an entry for each database partition server to the `db2nodes.cfg` file. The *nodenum* value 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 a Multiple Logical Node (MLN) and wish to keep the nodes logically grouped in this file.

This entry is required.
- hostname** The TCP/IP hostname of the database partition server for use by the FCM.

This entry is required.
- 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.

In this case, if you specify an entry for the *netname* field (see below), the entry *must* be set to 0.

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

Furthermore, if you specify a *logical port* entry for one database partition server, you must specify a *logical port* for each database partition server listed in your `db2nodes.cfg` file.

This field is only optional if 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 communication.

If an entry is specified for this field, all communication 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 are not using a clustered environment and want to have four database partition servers on one physical workstation called `workstation1`, update the `db2nodes.cfg` file as follows:

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

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

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

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


```

4      workstation1  0
6      workstation1  1
8      workstation1  2
9      workstation2  0

```

If you want your partitioned database system to contain two physical workstations, called `workstation1` and `workstation2` (with `workstation2` running two database partition servers), and use a high speed interconnect called `switch1` and `switch2`, 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:

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 199.
2. Ensure that 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` file is locked when the instance is running and can only be edited when the instance is stopped.

3. Edit the `db2nodes.cfg` file and add an entry for each database partition server that will participate in your partitioned database system.
4. Enter the `INSTHOME/sqllib/adm/db2start` command, where `INSTHOME` is the home directory of the instance owner.
5. Log out.

For more information on the `db2nodes.cfg` file, see the *Administration Guide*.

Step 4. Enable the Fast Communications Manager

This section describes how to enable communication between the database partition servers that participate in your partitioned database system. Communication 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 `db2setup` utility, 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 for any instance that was created using the **db2setup** utility.

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

```
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 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
```

where:

- *instance_name* is the name of the multipartition instance.
- *port_number* is the port number of 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 four database partitions, you will need to specify four consecutive ports as your port range. In our example, the entry in the `/etc/services` file is:

```
DB2_db2inst1        60000/tcp
DB2_db2inst1_END    60003/tcp
```



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    60003/tcp # that is running 4 MLNs.
```

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

1. Log on as root to the workstation where you created a DB2 instance.
2. Open the `/etc/services` file and verify that the port or port range entries that you have reserved for FCM communication are there.
3. Log out.
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.
5. Log out.

For more information on FCM communication, see the *Administration Guide*.

Step 5. 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 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` file or the `/etc/hosts.equiv` file, see your Linux documentation.

Step 6. 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 system.

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 each workstation as a user with root authority.
- Step 2. Enter the following command to create the links for the DB2 files:
`/opt/IBMdb2/V7.1/cfg/db2ln`
- Step 3. Log out.

Step 7. Start 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 system in your partitioned database system.

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 following command to the `/etc/inittab` file.

```
c1:234:once:DASINSTHOME/sql1lib/bin/db2cc1st
```

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

Step 8. Create the SAMPLE Database (Optional)

To verify that DB2 Enterprise - Extended Edition (DB2 EEE) is installed and configured correctly, create the SAMPLE database on your system, and access data from the database. For the purposes of testing the installation and configuration of DB2 EEE, you should only create the SAMPLE database for your instance after you have defined all of the database partition servers that will participate in your partitioned database system. If you want to add a database partition server to an instance where a database has already been created, there are other activities that you need to perform. For more information, see the *Administration Guide*.



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

For more information on how to create a file system, see your Linux documentation.

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

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 199.
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 cataloged with the database alias SAMPLE when it is created.

3. Start the database manager by entering the **db2start** command.
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 196 or “Entering Commands Using the Command Line Processor” on page 197.



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 9. 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 `node1ock` file, located in the `/var/lum` directory.
3. Update your DB2 product license by entering the following command:

```
/opt/IBMd2/V7.1/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 `db2udbee.lic`.

For example, if the CD-ROM is mounted in the `/cdrom` directory and the name of the license file is `db2udbee.lic`, you would enter the following command:

```
/opt/IBMd2/V7.1/cfg/db2licm /cdrom/db2/license/db2udbee.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 DB2 Control Center” on page 195.
- 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 5. Installing and Configuring DB2 Universal Database on Linux

This section describes how to install and configure a Linux-based partitioned database system. If you want to install a DB2 Run-Time Client or a DB2 Administration Client, go to “Chapter 10. Installing DB2 Clients” on page 137. 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*.

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



If you are migrating from a single-partition database system or a previous release of this product, you must complete certain procedures before installing DB2 Universal Database Enterprise - Extended Edition Version 7. See “Migrating from Previous Versions of DB2” on page 27 for information.

Before You Begin



To install DB2 EEE on a cluster, you must run the DB2 Installer program on each physical machine that you want to participate in 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. In order to run DB2 Enterprise – Extended Edition, you should update your kernel configuration parameters and recompile the kernel:
 - a. Log on as root.

- b. Update the kernel configuration parameters as follows:

Table 6. Linux Kernel Configuration Parameters (Recommended Values)

Kernel Parameter	Physical Memory	
	64MB - 256MB	256MB - or higher
SEMMNI (1)	256 or 512	512 to 1024
_SHM_ID_BITS (2)	8-9	9

- To increase the number of semaphore arrays, edit `/usr/src/linux/include/linux/sem.h` and change the `#define` for `SEMMNI` from 128 to 256, 512 or 1024 according to your physical memory.
 - To increase the number of shared memory segment identifiers, edit `/usr/include/asm/shmparam.h` and change the `#define` for `_SHM_ID_BITS` from 7 to 8 or 9, but not higher than 9.
- c. Recompile and install the new kernel. See your Linux documentation for information.
3. Ensure that you have a file system that is available to all machines that will participate in your partitioned database system. This file system will be used as the instance home directory.

For configurations that use more than one machine for a single database instance, NFS (Network File System) is used to share this file system. Typically, one machine in a cluster is used to export the file system using NFS, and the remaining machines in the cluster mount the NFS file system from this machine. The machine that exports the file system has the file system mounted locally. To create this file system, perform the following steps:

- a. On one machine, select a disk partition or create one using **fdisk**.
- b. Using a utility like **mkfs**, create a file system on this partition. The file system should be large enough to contain the necessary DB2 files as well as non-DB2 files.
- c. Locally mount the file system you have just created and add an entry to the `/etc/fstab` file so that this file system is mounted each time the system is rebooted. For example:


```
/dev/hda1 /nfshome ext2 defaults 1 2
```
- d. To automatically export an NFS file system on Linux at boot time, add an entry to the `/etc/exports` file. Be sure to include all of the host names participating in the cluster as well as all of the names that a machine might be known as. Also, ensure that each machine in the cluster has root authority on the exported file system by using the "root" option.



The `/etc/exportfs` is an ASCII file which contains the following type of information:

```
/nfshome machine1_name (rw) machine2_name (rw)
```

To export the NFS directory, run

```
/usr/sbin/exportfs -a
```

- e. On each of the remaining machines in the cluster, add an entry to the `/etc/fstab` file to NFS mount the file system automatically at boot time. As in the following example, when you specify the mount point options, ensure that the file system is mounted at boot time, is read-write, is mounted hard, includes the `bg` (background) option, and that `setuid` programs can be run properly.

```
fusion-en:/nfshome /nfshome nfs - rw,time0=300,retrans=5,  
wsize=1400,rsize=1400,hard,intr,bg,suid,rw,no1ock
```

`fusion-en` is the machine name.

- f. Once you have added a similar entry to the `/etc/fstab` file on each machine (except for the machine acting as the NFS server), NFS mount the exported file system on each of the remaining machines in the cluster by entering the following command:

```
mount /nfshome
```

If the `mount` command fails, use the **showmount** command to check the status of the NFS server. For example:

```
showmount -e fusion-en
```

This **showmount** command should list the file systems which are exported from the machine named `fusion-en`. If this command fails, the NFS server may not have been started. Run the following command as root on the NFS server to start the server manually:

```
/etc/rc.d/init.d/nfs restart
```

Assuming the present run level is 3, you can have this command run automatically at boot time by renaming `K20nfs` to `S20nfs` under the following directory: `/etc/rc.d/rc3.d`.

- g. Ensure that the following steps were successful:
- 1) On a single machine in the cluster, you have created a file system to be used as the instance and home directory.
 - 2) If you have a configuration that uses more than one machine for a single database instance, you have exported this file system using NFS.
 - 3) You have mounted the exported file system on each of the remaining machines in the cluster.

4. Create three separate groups and user accounts:

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



If you are using NIS or NIS+, groups and users must be created on the NIS server before running **db2setup**.

You can have the DB2 Installer create these user names, or you can create them manually. If you want to have the DB2 Installer create the users and groups, and you are not using NIS, skip this step, and proceed to “Performing the Installation” on page 75.

The user names you create must conform to both the Linux naming rules, and those of DB2. For more information on naming rules, see “Appendix C. Naming Rules” on page 207.

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

- a. Log on to the system as a user with root authority.
- b. Create a primary group for the instance owner (for example, `db2iadm1`), the user who will execute UDFs or stored procedures (for example, `db2fadm1`), and the Administration Server (for example, `db2asgrp`):

```
groupadd -g 999 db2iadm1
groupadd -g 998 db2fadm1
groupadd -g 997 db2asgrp
```

The group Ids you provided must be unique.



If you are using NIS or NIS+, secondary groups must be created for the DB2 instance owner and the Administration Server on the NIS server. You must then add the primary group of the instance owner to the secondary Administration Server group. Likewise, you must add the primary Administration Server group to the secondary group for the instance owner.

- c. Create a user who belongs to each group that you created in the previous step:

```
useradd -u 1004 -g db2adm1 -m -d /nfshome/db2inst1
db2inst1

useradd -u 10044 -g db2fadm1 -m -d /nfshome/db2fenc1
db2fenc1
```

```
useradd -u 10044 -g db2asgrp -m -d /nfshome/db2as
db2as
```

The user Ids you provide must be unique.

d. Set an initial password for each user that you created:

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

e. Log out.

Performing the Installation



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.
- Sometimes display problems can occur when running the **db2setup** utility. To avoid most potential display problems, install DB2 through a virtual console session which is a terminal window outside of the graphical interface that most Linux distributions are installed with.
- 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 EEE on a Linux system or cluster, perform the following steps:

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



If you are deploying DB2 EEE on a Linux cluster, you must install DB2 EEE on each machine.

2. Enter the following command to mount the CD-ROM:

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

where */cdrom* represents the mount point of the CD-ROM.

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

4. Change to the directory where the CD-ROM is mounted by entering the **cd /cdrom** command where **cdrom** is the mount point of your product CD-ROM.
5. Enter the **./db2setup** command to start the DB2 install program. After a few moments, the **Install DB2 V7** window opens.
6. From the product list on the *Install DB2 V7* screen, select the **DB2 Enterprise - Extended Edition** option and press the **Enter** key.
Press the **Tab** key to change the highlighted option and press the **Enter** key to select or deselect an option.
To select or deselect optional components for a DB2 product that you want to install, select the **Customize** option. To go back to a previous window at any time, select the **Cancel** option.



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

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.
8. Log out.

When the installation completes, the software is installed in the `/usr/IBMDB2/V7.1` directory.

Post Installation Steps

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



DB2 provides two environment profiles in the `sqllib` directory:

- `db2profile` for `sh`, `bash`, and `ksh`
- `db2cshrc` for `csh`

In Version 7, these environment profiles can no longer be modified.

To modify the DB2 environment in Version 7, first create, then modify the following profiles: `sqllib/userprofile` and `sqllib/usercshrc`. These environment profiles are called after the DB2 environment is initialized.

Step 1. Create a DB2 Instance



If you are using NIS or NIS+, before you create an instance, there must be an entry for the instance in the `/etc/services` file. For example, if you want to create an instance for the user `db2inst1`, you require an entry similar to the following:

```
DB2_db2inst1    60000/tcp
```

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 configures the instance automatically for communications.

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

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

1. Log on as a user with root authority on a machine where you installed DB2 Enterprise - Extended Edition (DB2 EEE).
2. Enter the `/usr/IBMd2/V7.1/install/db2setup` command to start the DB2 Installer program. The DB2 Installer window opens.
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.
4. Select the **Create a DB2 Instance** option and press the **Enter** key.
5. Fill in all of the fields for the username that you created for the instance owner in “Before You Begin” on page 31.
6. Select the **OK** option and press the **Enter** key.
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 31.
8. Select the **OK** option and press the **Enter** key. You are returned to the Create DB2 Services window.
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.
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.
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. The Administration Server is required if you plan to use the Control Center graphical administration tool to administer your partitioned database system.

Multiple Administration Servers can be created, but there can only be one Administration Server for each machine. If several users will be administering your system using the Control Center, you may consider creating multiple administration servers to distribute the network traffic. In a partitioned database system that has one Administration Server, the Administration Server often exists on the instance-owning node.

The **db2admin** commands allow you to start, stop, and configure the Administration Server. Refer to the *Command Reference* for more information about the **db2admin** commands.

We recommend that you use the DB2 Installer program to create the Administration Server because it configures the Administration Server automatically for communications. For information on how to create an instance using the **dasicrt** command, and configure it for communications, see 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 a machine where you installed DB2 Enterprise - Extended Edition (DB2 EEE).
- Step 2. Enter the **/usr/IBMDB2/V7.1/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.
- 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 31.
- 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/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
```

`nodenum`, `hostname`, `logical port`, and `netname` are defined as follows:

- nodenum** A unique number, between 0 and 999, that identifies a database partition server in a partitioned database system.
- To scale your partitioned database system, you add an entry for each database partition server to the `db2nodes.cfg` file. The *nodenum* value 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 a Multiple Logical Node (MLN) and wish to keep the nodes logically grouped in this file.
- This entry is required.
- hostname** The TCP/IP hostname of the database partition server for use by the FCM.
- This entry is required.
- 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.
- In this case, if you specify an entry for the *netname* field (see below), the entry *must* be set to 0.
- If you are using MLNs, the *logical port* value you specify *must* start at 0 and continue in ascending order with no gaps (for example, 0,1,2).

Furthermore, if you specify a *logical port* entry for one database partition server, you must specify a *logical port* for each database partition server listed in your `db2nodes.cfg` file.

This field is only optional if 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 communication.

If an entry is specified for this field, all communication 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 are not using a clustered environment and want to have four database partition servers on one physical workstation called `workstation1`, update the `db2nodes.cfg` file as follows:

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

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

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

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

```
4 workstation1 0
6 workstation1 1
8 workstation1 2
9 workstation2 0
```

If you want your partitioned database system to contain two physical workstations, called `workstation1` and `workstation2` (with `workstation2` running two database partition servers), and use a high speed interconnect called `switch1` and `switch2`, 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:

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 199.
2. Ensure that 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` file is locked when the instance is running and can only be edited when the instance is stopped.

3. Edit the `db2nodes.cfg` file and add an entry for each database partition server that will participate in your partitioned database system.
4. Enter the `INSTHOME/sqllib/adm/db2start` command, where `INSTHOME` is the home directory of the instance owner.
5. Log out.

For more information on the `db2nodes.cfg` file, see the *Administration Guide*.

Step 4. Enable the Fast Communications Manager

This section describes how to enable communication between the database partition servers that participate in your partitioned database system. Communication 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 `db2setup` utility, 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 for any instance that was created using the **db2setup** utility.

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

```
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 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
```

where:

- *instance_name* is the name of the multipartition instance.
- *port_number* is the port number of 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 four database partitions, you will need to specify four consecutive ports as your port range. In our example, the entry in the `/etc/services` file is:

```
DB2_db2inst1      60000/tcp
DB2_db2inst1_END  60003/tcp
```



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  60003/tcp # that is running 4 MLNs.
```

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

1. Log on as root to the workstation where you created a DB2 instance.
2. Open the `/etc/services` file and verify that the port or port range entries that you have reserved for FCM communication are there.
3. Log out.

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.
5. Log out.

For more information on FCM communication, see the *Administration Guide*.

Step 5. Check Available Paging Space

1. Log on to a database partition server as a user with root authority.
2. 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 DB2 EEE system may not operate properly. To check for available paging space, enter the following command:

```
free -b
```

This command returns output similar to the following:

	total	used	free	shared	buffers	cached
Mem:	262819840	254001152	8818688	60784640	49319936	109191168
-/+ buffers/cache:		95490048	167329792			
Swap:	526376960	4898816	521478144			

We recommend that the paging space available be equal to two-and-a-half to three times the amount of physical memory installed on your workstation, the minimum being 1 GB.

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 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` file or the `/etc/hosts.equiv` file, see your Linux 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 system.

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:

1. Log on to each workstation as a user with root authority.
2. Enter the following command to create the links for the DB2 files:

```
/usr/IBMd2/V7.1/cfg/db2ln
```

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

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 system in your partitioned database system.

To enable your partitioned database system so that it can be administered by the CCA or the Control Center, perform the following steps on each node where an Administration Server was created.

- Step 1. Log on to the workstation with the user name that you created for the Administration Server (for example, `db2as`).

- Step 2. 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 3. Enter the following command to start the `db2cc1st` daemon:

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

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

- Step 4. 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 5. 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/sqllib/bin/db2cc1st
```

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

Step 9. Create the SAMPLE Database (Optional)

To verify that DB2 Enterprise - Extended Edition (DB2 EEE) is installed and configured correctly, create the SAMPLE database on your system, and access data from the database. For the purposes of testing the installation and configuration of DB2 EEE, you should only create the SAMPLE database for your instance after you have defined all of the database partition servers that will participate in your partitioned database system. If you want to add a database partition server to an instance where a database has already been created, there are other activities that you need to perform. For more information, see the *Administration Guide*.



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

For more information on how to create a file system, see your Linux documentation.

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

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 199.
2. Enter the following command to create the SAMPLE database:

```
INSTHOME/sqllib/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/db2samp1 /database
```

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

3. Start the database manager by entering the **db2start** command.
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 196 or “Entering Commands Using the Command Line Processor” on page 197.



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 node1ock file, located in the /var/lum directory.
3. Update your DB2 product license by entering the following command:

```
/usr/IBMDB2/V7.1/adm/db2licm license_filename
```

where *license_filename* is the full path name and file name for the license file that corresponds to the product you have purchased.

The name of the license file for this product is `db2udbee.lic`.

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

```
/usr/IBMd2/V7.1/adm/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:

1. Start the Control Center. For more information, see “Starting the DB2 Control Center” on page 195.
2. Click on the **License Center** icon from the Control Center toolbar. The License Center opens.
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 6. Installing and Configuring DB2 Universal Database for NUMA-Q

This section describes how to install and configure a NUMA-Q/PTX-based partitioned database system. If you want to install a DB2 Client, go to “Chapter 10. Installing DB2 Clients” on page 137. For information on how to deploy this product using a distributed installation method, 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 install program, as well as create an instance and the Administration Server using the DB2 Installer program.

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. 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/vfstab` file so that this file system is mounted each time the system is rebooted.
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.



If you are using NIS or NIS+, groups and users must be created on the NIS server before running `db2setup`.

The usernames you create must conform to both your operating system's naming rules, and those of DB2. You can have the DB2 Installer create these usernames, or you can create them manually. For more information on naming rules, see "Appendix C. Naming Rules" on page 207.

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
```



If you are using NIS or NIS+, secondary groups must be created for the DB2 instance owner and the Administration Server on the NIS server. You must then add the primary group of the instance owner to the secondary Administration Server group. Likewise, you must add the primary Administration Server group, to the secondary group for the instance owner.

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:



For changes to the kernel configuration parameters to take effect, you will be required reboot the system.

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

Step b. Update the kernel configuration parameters as follows:

Table 7. PTX Kernel Configuration Parameters (Recommended Values)

Kernel Parameter	Physical Memory
	512MB+
msgmap	514
msgmax(1)	65535
msgmnb(1)	65535
msgmni	512
msgssz	64
msgtql	1024
msgseg(2)	32767
semnmi	1024
semmap	514
semmnu	2048
semmns	2048
semume	80
shmmax(3)	2147483647
shmseg	16
shmmni	300
shm_lock_ok	1
shm_lock_uid	-1

Notes:

- 1) The msgmax and msgmnb must be set to 65535 or larger.
- 2) The msgseg parameter must be set no higher than 32767.
- 3) The shmmax must be set to 2147483647 or larger.

To modify your PTX kernel configuration parameters, perform the following steps:

- Step 1. Log on as a user with root authority.
- Step 2. Enter the menu command.
- Step 3. Press A to select the **System Administration** option.
- Step 4. Press C to select the **Kernel Configuration** option.
- Step 5. Press Ctrl+F in the **Change Kernel Configuration Disk** form. If you want to build the new kernel in a disk other than the root disk, enter the disk and press Ctrl+F.
- Step 6. In the Compile, Configure, or Remove a Kernel window, select the kernel configuration type with which your current kernel is built and press K.
- Step 7. In the **Configure a kernel with site specific parameters** form, go down one page (press Ctrl+D), press A for **All** in the **Visibility level for parameter changes** and then press Ctrl+F.
- Step 8. In the Configure Files With Adjustable Parameters window select ALL (press Ctrl+T) and press Ctrl+F.

- Step 9. In the Tunable Parameters window, use arrow keys to navigate. Press `Ctrl+T` to select the parameter you want to change and press `Ctrl+F`.
- Step 10. In the Detail of Parameter Expression(s) window, press `s` to set the new value.
- Step 11. In the **Add site specific 'set' parameter** form, enter the new value and press `Ctrl+F`.
- Step 12. Repeat steps 9 through 11 to change the values of all other parameters that you want to change.
- Step 13. Once you are done changing all the parameters, press `Ctrl+E` from the Tunable Parameters window.
- Step 14. Compile the kernel.
- Step 15. Press `Ctrl+X` to exit the menu.
- Step 16. Reboot the system so that the changes can take effect.

Performing the Installation

This section describes how to install DB2 EEE on a NUMA-Q/PTX system.



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 EEE, perform the following steps:

- Step 1. Log on to the system as a user with root authority.
- Step 2. Insert the appropriate CD-ROM into the drive.
- Step 3. Mount the CD-ROM by entering the following commands:

```
mkdir /cdrom
mount -r -F cdfs /dev/dsk/cd0 /cdrom
```

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

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 4. 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.

Step 5. Enter the **./db2setup** command to start the DB2 install program. After a few moments, the **Install DB2 V7** window opens.

Step 6. From the product list on the *Install DB2 V7* 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 select or deselect optional components for a DB2 product that you want to install, select the **Customize** option. To go back to a previous window 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/V7.1* directory.

Post Installation Steps

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



DB2 provides two environment profiles in the *sqllib* directory:

- *db2profile* for *sh*, *bash*, and *ksh*
- *db2cshrc* for *csh*

In Version 7, these environment profiles can no longer be modified.

To modify the DB2 environment in Version 7, first create, then modify the following profiles: *sqllib/userprofile* and *sqllib/usercshrc*. These environment profiles will be called after the DB2 environment is initialized.

Step 1. Create a DB2 Instance



If you are using NIS or NIS+, before you create an instance, there must be an entry for the instance in the `etc/services` file. For example, if you want to create an instance for the user `db2inst1`, you require an entry similar to the following:

```
DB2_db2inst1    60000/tcp
```

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 on a machine where you installed DB2 Enterprise - Extended Edition (DB2 EEE).
- Step 2. Enter the `/opt/IBMDb2/V7.1/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.
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 31.
- 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 31.
- 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 **dasict** 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 a machine where you installed DB2 Enterprise - Extended Edition (DB2 EEE).
- Step 2. Enter the `/opt/IBMdb2/V7.1/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 31.
- 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/sqlllib/` 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 different for PTX when compared to its format for other UNIX-based platforms. In PTX, the `db2nodes.cfg` can optionally take into account the quad number that the database partition server runs on. The format is:

```
nodenum    hostname    logical port    quad    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 the database partition server 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 *quad* or *netname* fields, you must enter a number for the *logical port* field.

If you are not using MLNs and you specify an entry for the *quad* or *netname* fields (see below), 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).

Furthermore, if you specify a *logical port* entry for one database partition server, you must specify a *logical port* for each database partition server listed in your `db2nodes.cfg` file.

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

quad

This field is used to specify a particular quad where the database partition server will run, including all processes and shared memory. You may want to direct database partition servers to run on specific quads because other applications may already be running on a subset of quads, for performance reasons, or to address data skew without redistribution, etc. For example, if application A was running on quads 2 and 3, you may want all DB2 processing to run on quads 0 and 1.

If you specify a quad number for one database partition server, you must specify a quad number for each database partition server listed in your `db2nodes.cfg` file.

If there is no entry for this field, the default is to assign a database partition server to run on an available quad in a round-robin fashion. For example, if you defined 3 database partition servers (nodes 0,1,2) in your `db2nodes.cfg` file and did not specify a quad number, DB2 would assign node 0 to run on quad 0, node 1 to run on quad 1, and finally node 2 to run on quad 2.

We recommend that you specify this field.

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) are 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 system called `workstation1`, the `db2nodes.cfg` file was updated as follows:

```
0          workstation1  0
```

We recommend that you run 1 multiple logical node per quad. Typically, in a PTX environment, you would run four database partition servers on a 4 quad

system, note however that some systems are bigger than 4 quads. If this system was called `workstation1`, you would update the `db2nodes.cfg` file as follows:

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

If you wanted to equally divide DB2 processing on 2 quads only, you would update the `db2nodes.cfg` file as follows:

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

In this example, database partition servers 0 and 1 will run on quad 0 and database partition servers 2 and 3 will run on quad 1.

If you want your partitioned database system to contain two physical 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 199.
- 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. Enter the `INSTHOME/sqllib/adm/db2start` command, where `INSTHOME` is the home directory of the instance owner.
- Step 5. Log out.

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

Step 4. Enable the Fast Communications Manager

This section describes how to enable communication between the database partition servers that participate in your partitioned database system. Communication 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 `db2setup` utility, 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 for any instance that was created using the `db2setup` utility.

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

```
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 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
```

where:

- `instance_name` is the name of the multipartition instance.

- *port_number* is the port number of 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 four database partitions, you will need to specify four consecutive ports as your port range. In our example, the entry in the `/etc/services` file is:

```
DB2_db2inst1      60000/tcp
DB2_db2inst1_END 60003/tcp
```



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 60003/tcp # that is running 4 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.

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

Step 5. 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 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` file or the `/etc/hosts.equiv` file, see your Linux documentation.

Step 6. 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 system.

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 each workstation as a user with root authority.

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

```
/opt/IBMDB2/V7.1/cfg/db2ln
```

Step 3. Log out.

Step 7. Start 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 system in your partitioned database system.

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:

db2ps | 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/sqllib/bin/db2cc1st
```

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

Step 8. Create the SAMPLE Database (Optional)

To verify that DB2 Enterprise - Extended Edition (DB2 EEE) is installed and configured correctly, create the SAMPLE database on your system, and access data from the database. For the purposes of testing the installation and configuration of DB2 EEE, you should only create the SAMPLE database for your instance after you have defined all of the database partition servers that will participate in your partitioned database system. If you want to add a database partition server to an instance where a database has already been created, there are other activities that you need to perform. For more information, see the *Administration Guide*.



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

For more information on how to create a file system, see your Linux documentation.

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

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 199.
2. Enter the following command to create the SAMPLE database:

```
INSTHOME/sqllib/bin/db2sampl 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 cataloged with the database alias SAMPLE when it is created.

3. Start the database manager by entering the **db2start** command.
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 196 or “Entering Commands Using the Command Line Processor” on page 197.



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 9. 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 `node1ock` file, located in the `/var/lum` directory.
3. Update your DB2 product license by entering the following command:

```
/opt/IBMdb2/V7.1/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:

```
/opt/IBMd2/V7.1/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 DB2 Control Center” on page 195.
- 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 7. 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 10. Installing DB2 Clients” on page 137. 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 7. See “Migrating from Previous Versions of DB2” on page 27 for information.

Before You Begin



If you are installing DB2 EEE on a cluster, and you are not running cluster software, you will have to run the DB2 Installer program on each physical machine that you want to participate in your partitioned database system.

If you are running SunCluster 2.2, you can use `ctelnet` or `cconsole` to perform your installation. For more information, see your product documentation.

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. Ensure that you have a file system that is available to all machines that will participate in your partitioned database system. This file system will be used as the instance home directory.

For configurations that use more than one machine for a single database instance, NFS (Network File System) is used to share this file system. Typically, one machine in a cluster will be used to export the file system via NFS, and the remaining machines in the cluster will mount the NFS file system from this machine. The machine that exports the file system will have the file system mounted locally. To create this file system, perform the following steps:



There are a number of ways to create a local file system on a Solaris system. If you want to use a product such as Veritas to create the file system, refer to the product's documentation.

- Step a. On one machine, select a disk partition or configure one using the **format** command. When using the **format** command, ensure that the disk partitions being used do not overlap. Overlapping partitions can cause data corruption or file system failures.
- Step b. Using a utility like **newfs** or **mkfs**, create a file system on this partition. The file system should be large enough to contain the necessary DB2 files as well as other non-DB2 files. A minimum of 300 MB is recommended.
- Step c. Locally mount the file system you have just created and add an entry to the `/etc/vfstab` file so that this file system is mounted each time the system is rebooted. For example:

```
dev/dsk/c1t0d2s2 /dev/rdisk/c1t0d2s2 /nfshome ufs 2 yes -
```



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

- Step d. To automatically export an NFS file system on Solaris at boot time, add an entry to the `/etc/dfs/dfstab` file. Be sure to include all of the host names participating in the cluster as well as all of the names that a machine might be known as. Also, ensure that each machine in the cluster has root authority on the exported file system by using the "root" option.

In the following example, an entry for a two machine cluster is added to the `/etc/dfs/dfstab` file. The system called `plasma` is given permission to mount the file system `/nfshome`, which will be used as the file system for the instance home directory.

```
share -F nfs -o \
  rw=plasma.torolab.ibm.com:plasma-tr.torolab.ibm.com,\
  root=plasma.torolab.ibm.com:plasma-tr.torolab.ibm.com \
  -d "homes" /nfshome
```

Note: -tr indicates that the system is using a token ring interface.



The `/etc/dfs/dfstab` file contains share commands which can also be run manually. The `/etc/dfs/dfstab` file can be run with any shell. For example:

```
sh /etc/dfs/dfstab
```

This command will export all file systems listed in the `/etc/dfs/dfstab` file.

Step e. On each of the remaining machines in the cluster, add an entry to the `/etc/vfstab` file to NFS mount the file system automatically at boot time. As in the following example, when you specify the mount point options, ensure that the file system is mounted at boot time, is read-write, is mounted hard, includes the `bg` (background) option, and that `setuid` programs can be run properly:

```
fusion-en:/nfshome - /nfshome nfs - yes rw,hard,intr,bg
```

Step f. Once you have added a similar entry to the `/etc/vfstab` file on each machine (except for the machine acting as the NFS server), NFS mount the exported file system on each of the remaining machines in the cluster by entering the following command:

```
mount /nfshome
```

If the mount command fails, use the `showmount` command to check the status of the NFS server. For example:

```
showmount -e fusion-en
```

This `showmount` command should list the file systems which are exported from the machine named `fusion-en`. If this command fails, the NFS server may not have been started. To start the server manually, run the following commands as root on the NFS server:

```
/usr/lib/nfs/mountd  
/usr/lib/nfs/nfsd -a 16
```

These commands are run automatically at boot time if there are any entries in the `/etc/dfs/dfstab` file. After starting the NFS server, export the NFS file system again by running the following command:

```
sh /etc/dfs/dfstab
```

Step g. Ensure that you have completed the following steps:

- 1) On a single machine in the cluster, you have created a file system to be used as the instance and home directory.

- 2) If you have a configuration that uses more than one machine for a single database instance, you have exported this file system via NFS.
 - 3) You have mounted the exported file system on each of the remaining machines in the 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



If you are using NIS or NIS+, groups and users must be created on the NIS server before running db2setup.

You can have the DB2 Installer create these user names, or you can create them manually. The user names you create must conform to both your operating system's naming rules, and those of DB2. For more information on naming rules, see "Appendix C. Naming Rules" on page 207.

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 primary 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
```



If you are using NIS or NIS+, secondary groups must be created for the DB2 instance owner and the Administration Server on the NIS server. You must then add the primary group of the instance owner to the secondary Administration Server group. Likewise, you must add the primary Administration Server group, to the secondary group for the instance owner.

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 /nfshome/db2inst1 -m db2inst1
useradd -g db2fadm1 -u 1003 -d /nfshome/db2fenc1 -m db2fenc1
useradd -g db2asgrp -u 1002 -d /nfshome/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 parameters 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 8. Solaris Kernel Configuration Parameters (Recommended Values)

Kernel Parameter	Physical Memory			
	64MB - 128MB	128MB - 256MB	256MB - 512MB	512MB+
msgsys:msginfo_msgmax	65535(1)	65535(1)	65535(1)	65535(1)
msgsys:msginfo_msgmnb	65535(1)	65535(1)	65535(1)	65535(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	8192	16384	32767(2)	32767(2)
shmsys:shminfo_shmmax	67108864	134217728(2)	268435456(3)	536870912(3)
shmsys:shminfo_shmseg	50	50	50	50
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_semmnu	256	512	1024	2048
semsys:seminfo_semume	50	50	50	50

Notes:

- 1) The msgsys:msginfo_msgmnb and msgsys:msginfo_msgmax parameters must be set to 65535 or larger.
- 2) The msgsys:msgseg parameter must be set no higher than 32767.
- 3) 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*1024*1024*0.9).

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/V7.1/cfg* directory, or on your DB2 product CD-ROM in the */db2/install/samples* directory. The names for these files are as follows:

kernel.param.64MB

for systems with 64MB–128MB 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 3 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 EEE, perform the following steps:

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

where */cdrom/unnamed_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/unnamed_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 install program. After a few moments, the **Install DB2 V7** window opens.

Step 6. From the product list on the *Install DB2 V7* 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 select or deselect optional components for a DB2 product that you want to install, select the **Customize** option. To go back to a previous window 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/V7.1` directory.

Post Installation Steps

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



DB2 provides two environment profiles in the `sqllib` directory:

- `db2profile` for `sh`, `bash`, and `ksh`
- `db2cshrc` for `csh`

In Version 7, these environment profiles can no longer be modified.

To modify the DB2 environment in Version 7, first create, then modify the following profiles: `sqllib/userprofile` and `sqllib/usercshrc`. These environment profiles will be called after the DB2 environment is initialized.

Step 1. Create a DB2 Instance



If you are using NIS or NIS+, before you create an instance, there must be an entry for the instance in the `etc/services` file. For example, if you want to create an instance for the user `db2inst1`, you require an entry similar to the following:

```
DB2_db2inst1    60000/tcp
```

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 on a machine where you installed DB2 Enterprise - Extended Edition (DB2 EEE).
- Step 2. Enter the `/opt/IBMdb2/V7.1/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.
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 31.
- 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 31.
- 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 a machine where you installed DB2 Enterprise - Extended Edition (DB2 EEE).
- Step 2. Enter the **/opt/IBMd2/V7.1/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 31.
- 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

nodenum, hostname, logical port, and netname are defined as follows:

- nodenum** A unique number, between 0 and 999, that identifies a database partition server in a partitioned database system.
- To scale your partitioned database system, you add an entry for each database partition server to the `db2nodes.cfg` file. The *nodenum* value 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 a Multiple Logical Node (MLN) and wish to keep the nodes logically grouped in this file.
- This entry is required.
- hostname** The TCP/IP hostname of the database partition server for use by the FCM.
- This entry is required.
- 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.
- In this case, if you specify an entry for the *netname* field (see below), the entry *must* be set to 0.
- If you are using MLNs, the *logical port* value you specify *must* start at 0 and continue in ascending order with no gaps (for example, 0,1,2).
- Furthermore, if you specify a *logical port* entry for one database partition server, you must specify a *logical port* for each database partition server listed in your `db2nodes.cfg` file.
- This field is only optional if 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 communication.
- If an entry is specified for this field, all communication 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 are not using a clustered environment and want to have four database partition servers on one physical workstation called `workstation1`, update the `db2nodes.cfg` file as follows:

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

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

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

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

```
4          workstation1  0
6          workstation1  1
8          workstation1  2
9          workstation2  0
```

If you want your partitioned database system to contain two physical workstations, called `workstation1` and `workstation2` (with `workstation2` running two database partition servers), and use a high speed interconnect called `switch1` and `switch2`, 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:

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 199.
2. Ensure that the DB2 instance is stopped by entering the `INSTHOME/sql/lib/adm/db2stop` command, where `INSTHOME` is the home directory of the instance owner.



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

3. Edit the `db2nodes.cfg` file and add an entry for each database partition server that will participate in your partitioned database system.
4. Enter the `INSTHOME/sqllib/adm/db2start` command, where `INSTHOME` is the home directory of the instance owner.
5. Log out.

For more information on the `db2nodes.cfg` file, see the *Administration Guide*.

Step 4. Enable the Fast Communications Manager

This section describes how to enable communication between the database partition servers that participate in your partitioned database system. Communication 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 **db2setup** utility, 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 for any instance that was created using the **db2setup** utility.

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

```
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 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
```

where:

- `instance_name` is the name of the multipartition instance.
- `port_number` is the port number of 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 four database partitions, you will need to specify four consecutive ports as your port range. In our example, the entry in the `/etc/services` file is:

```
DB2_db2inst1    60000/tcp
DB2_db2inst1_END 60003/tcp
```



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 60003/tcp # that is running 4 MLNs.
```

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

1. Log on as root to the workstation where you created a DB2 instance.
2. Open the `/etc/services` file and verify that the port or port range entries that you have reserved for FCM communication are there.
3. Log out.
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.
5. Log out.

For more information on FCM communication, see the *Administration Guide*.

Step 5. Check Available Paging Space

1. Log on to a database partition server as a user with root authority.
2. 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 DB2 EEE system may not operate properly. To check for available paging space, enter the following command:

```
swap -l
```


This command will return output similar to the following:

```
swapfile      dev  swaplo  blocks  free
/dev/dsk/c0t0d0s7  32,7  16      5120736 5070096
```

We recommend that the paging space available be equal to two and a half to three times the amount of physical memory installed on your workstation, the minimum being one GB.

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 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* file or the */etc/hosts.equiv* file, see your Linux 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 system.

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 each workstation as a user with root authority.

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

```
/opt/IBMdb2/V7.1/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 8. Start 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 system in your partitioned database system.

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 user name 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 following command to the `/etc/inittab` file.

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

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

Step 9. Create the SAMPLE Database (Optional)

To verify that DB2 Enterprise - Extended Edition (DB2 EEE) is installed and configured correctly, create the SAMPLE database on your system, and access data from the database. For the purposes of testing the installation and configuration of DB2 EEE, you should only create the SAMPLE database for your instance after you have defined all of the database partition servers that will participate in your partitioned database system. If you want to add a database partition server to an instance where a database has already been created, there are other activities that you need to perform. For more information, see the *Administration Guide*.



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

For more information on how to create a file system, see your Linux documentation.

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

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 199.
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 cataloged with the database alias SAMPLE when it is created.

3. Start the database manager by entering the **db2start** command.
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 196 or “Entering Commands Using the Command Line Processor” on page 197.



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 `node1ock` file, located in the `/var/lum` directory.
3. Update your DB2 product license by entering the following command:

```
/opt/IBMd2/V7.1/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 `db2udbee.lic`.

For example, if the CD-ROM is mounted in the `/cdrom` directory and the name of the license file is `db2udbee.lic`, you would enter the following command:

```
/opt/IBMd2/V7.1/cfg/db2licm /cdrom/db2/license/db2udbee.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 DB2 Control Center” on page 195.
- 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 8. DB2 Post-installation Migration Tasks

After verifying your DB2 installation, you should migrate your instances by running the **db2imigr** command. The **db2imigr** command checks that your instance can be migrated. It also calls the **db2ckmig** command which checks that the databases in the instance can be migrated. After migrating your instances, you must migrate each database in an instance by running the **migrate database** command.

You can skip this chapter if you are not migrating instances and databases.

The procedures in this chapter assume that you have performed the pre-installation migration tasks described in “Chapter 2. Planning for Installation” on page 17.

Migrating Instances

To migrate instances, including your Administration Server instance, to a format usable by DB2 Version 7, you must run the **db2imigr** command:

1. Log in as user with root authority.
2. Run the **db2imigr** command as follows:

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

where *DB2DIR* = /usr/lpp/db2_07_01 on AIX
= /opt/IBMcdb2/V7.1 on HP-UX, PTX, or Solaris

and where:

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

-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 “Possible Migration Error Messages and User Responses” and take the suggested corrective actions. Then, reenter the **db2imigr** command.

Note: If you are migrating a DB2 for AIX instance and the it uses the environment variable *DB2SORT* set to the keyword *SMARTSORT*, you must set the registry value *db2sort* after the instance is migrated to Version 7. Set the *db2sort* registry value to the run time library for the sort command as follows:

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

Possible Migration Error Messages and User Responses

If the **db2imigr** command fails, check the following list of error messages and take the appropriate 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.

Structured type and function have the same name

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

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

5. Import or load the data into the object.

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

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

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

Migrating Databases

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

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

When pre-Version 7 databases are migrated for use in a Version 7 partitioned database system, default nodegroups are created for them. For more information about default nodegroups and how they are used, see “Nodegroups and Data Partitioning” on page 8.

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

1. Log in with a user account that has SYSADM authority. See “Working with the System Administrative Group” on page 199 for more information.
2. Ensure that the databases you want to migrate are cataloged. To retrieve a list of all cataloged databases on your system, enter the **db2 list database directory** command. The syntax of this command is as follows:

DB2 LIST DATABASE DIRECTORY command

```
▶▶LIST DATABASE DIRECTORY ON path drive▶▶
```

The diagram shows the command syntax with brackets indicating the structure: LIST is followed by a bracketed field containing DATABASE and DB. This is followed by DIRECTORY, then ON, then a bracketed field containing path and drive. The command ends with a long arrow pointing to the right.

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

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

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

To convert unique indexes:

1. Log in with a user account that has SYSADM authority. See “Working with the System Administrative Group” on page 199 for more information.
2. Start the database manager by entering the **db2start** command.
3. Run the **db2uiddl** command against your migrated database. Refer to the *Command Reference* for the syntax of this command.

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

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

```
db2 -tvf filename
```

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

- **Update Statistics**

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

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

- **Rebind Packages**

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

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

- **Update database and database manager configuration**

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

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

- **Migrate Explain Tables**

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

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

where:

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

The explain tables belonging to the user ID that is running **db2exmig**, or that is used to connect to the database, will be migrated. The explain tables migration tool will rename the Version 5.x and Version 6 tables, create a new set of tables using the **EXPLAIN.DDL**, and copy the contents of the old tables to the new tables. Finally, it will drop the old tables. The **db2exmig** migration utility will preserve any user added columns on the explain tables.

Chapter 9. Your Next Step

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

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

- “Chapter 10. Installing DB2 Clients” on page 137
- “Chapter 11. Installing DB2 Clients on Windows 32-Bit Operating Systems” on page 141
- “Chapter 12. Installing DB2 Clients on OS/2 Operating Systems” on page 145
- “Chapter 13. Installing DB2 Clients on UNIX Operating Systems” on page 149
- “Chapter 14. Control Center Installation and Configuration” on page 155
- “Chapter 15. Configuring Client-to-Server Communications Using the Client Configuration Assistant” on page 171
- “Chapter 16. Configuring Client-to-Server Communications Using the Command Line Processor” on page 183

DB2 Universal Database Administration Guide

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

DB2 Universal Database Administration Guide: Planning

- Logical and physical database design
- Designing distributed databases
- Designing for transaction managers
- Designing for high availability
- High availability cluster multi processing, enhanced scalability, (HACMP) for AIX
- DB2 and high availability on Sun Cluster 2.2

DB2 Universal Database Administration Guide: Implementation

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

DB2 Universal Database Administration Guide: Performance

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

Task Information in other DB2 Documentation

Table 9. Task Information in other DB2 Documentation

If you want to ...	Refer to the ...
Install multiple DB2 clients without having to run the setup program at each DB2 client workstation	<i>DB2 Installation and Configuration Supplement</i>
Monitor a system's performance	<i>DB2 System Monitor Guide and Reference</i>
Load or export data from an existing DB2 database.	<i>DB2 Data Movement Utilities Guide and Reference</i>
Connect to Oracle data sources	<i>DB2 Installation and Configuration Supplement</i>
Perform replication activities	<i>DB2 Replication Guide and Reference</i>
Perform trouble shooting activities	<i>DB2 Troubleshooting Guide</i>

Table 9. Task Information in other DB2 Documentation (continued)

If you want to ...	Refer to the ...
Use text, image, audio, video, or spatial extenders	<i>DB2 Text Extender Administration and Programming, DB2 Image, Audio, and Video Extenders Administration and Programming, DB2 Spatial Extenders User's Guide and Reference</i>
Access data that resides on a host or AS/400 DB2 database	<i>DB2 Connect Enterprise Edition for UNIX Quick Beginnings</i>
Learn more about DB2 SQL	<i>DB2 SQL Reference, DB2 SQL Getting Started</i>

Refer to “Appendix E. Using the DB2 Library” on page 237 for more information about the DB2 product library.

Part 3. Installing DB2 Clients

Chapter 10. Installing DB2 Clients

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

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

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

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



To go to the section that provides client installation instructions:

- “Chapter 11. Installing DB2 Clients on Windows 32-Bit Operating Systems” on page 141
- “Chapter 12. Installing DB2 Clients on OS/2 Operating Systems” on page 145
- “Chapter 13. Installing DB2 Clients on UNIX Operating Systems” on page 149

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

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

DB2 Run-Time Client

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

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

DB2 Administration Client

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

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

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

DB2 Application Development Client

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

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

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

Distributed Installation

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

DB2 Thin Client

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

Chapter 11. Installing DB2 Clients on Windows 32-Bit Operating Systems

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

Before You Begin Installing

1. Ensure that your system meets all of the memory, hardware, and software requirements to install your DB2 client. For more information, see “Chapter 2. Planning for Installation” on page 17.
2. You require a user account to perform the installation.

Windows 9x

Any valid Windows 9x user.

Windows NT or Windows 2000

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

Installing Without Administrator Authority

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

- Control Center
- NetQuestion
- Integrated SNA Support

The following are some specific installation scenarios:

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

the same machine. In this scenario, the installation by the second user will fail, and return an error message that the user must be an administrator to install the product.

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

Installation Steps

To install a DB2 client, perform the following steps:

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



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

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


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

where:

- *x*: represents your CD-ROM drive
- *language* represents the country code for your language (for example, EN for English). Table 18 on page 206 lists the code for each available language.

- c. Click **OK**.

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



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

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



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

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

The setup program performs the following actions:

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



To configure your client to access remote servers, go to “Chapter 15. Configuring Client-to-Server Communications Using the Client Configuration Assistant” on page 171.

Chapter 12. Installing DB2 Clients on OS/2 Operating Systems

This section contains the information that you need to install a DB2 client on OS/2 operating systems. If you have a pre-Version 7 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.ibm.com/software/data/db2/db2tech/clientpak.html>

Before You Begin Installing

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 user ID USERID with password PASSWORD.

3. To verify that DB2 installed correctly, you will need to have a user account that belongs to the DB2 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 199. For more information on valid DB2 usernames, see “Appendix C. Naming Rules” on page 207.

Installation Steps

To install a DB2 client for OS/2, perform the following steps:

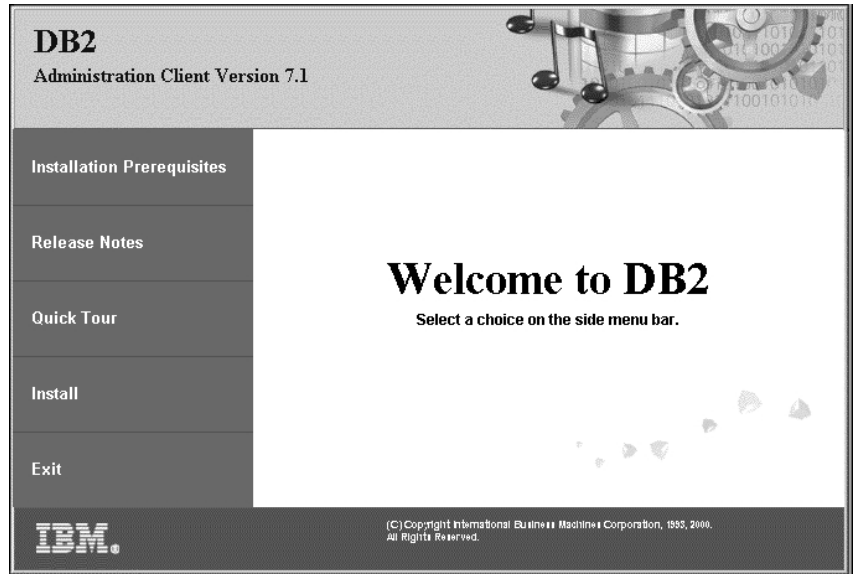
- Step 1. Insert the appropriate CD-ROM into the drive.

Step 2. Open an OS/2 command window, set the directory to your CD-ROM drive, and enter the following command:

```
x:\install
```

where *x* represents your CD-ROM drive.

Step 3. The DB2 Launchpad opens. It looks similar to the following:



Step 4. From this window, you can view the installation prerequisites and the release notes, you can take a Quick Tour to explore the features, capabilities, and benefits of DB2 Universal Database Version 7, or you can proceed directly to the installation.

Once you have initiated the installation, proceed by following the install program's prompts. Online help is available to guide you through the remaining steps. To invoke the online help, click **Help**, or press **F1**.



For information on errors encountered during installation, see the 11.1.log and 12.1.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 15. Configuring Client-to-Server Communications Using the Client Configuration Assistant" on page 171.

If you want to use ODBC applications with OS/2, you must ensure that the `\sql11ib\dll\odbc.dll` file appears as the first `odbc.dll` in the `LIBPATH` parameter of the `config.sys` file. The Install program does not automatically position the dynamic link library (dll) as of Version 7. If `odbc.dll` is not the first ODBC dll listed, you may have problems connecting to DB2 through ODBC applications.

Chapter 13. Installing DB2 Clients on UNIX Operating Systems

This section contains the information that you will need to install a DB2 client on a UNIX-based workstation.

Before You Begin

Before you begin installing a DB2 client using the `db2setup` utility, you need to 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.

About the `db2setup` utility

We recommend that you use the `db2setup` utility to install DB2 products on UNIX-based systems. This utility leads you through the installation process using a simple interface that includes online help. Default values are provided for all required installation parameters, but you can also enter your own values.

The `db2setup` utility can:

- Create or assign groups and user IDs.
- Create a DB2 instance.
- Install product messages.

Note: Documentation is available in HTML format on the product CD-ROM.

If you choose to install DB2 using another method, such as using your operating system’s native administration tools, you will have to perform these tasks manually. For more information on installing DB2 manually, see your server platform’s installation chapter.

The `db2setup` utility can generate a trace log to record errors during installation. To generate a trace log, run the `./db2setup -d` command. This generates a log in `/tmp/db2setup.trc`.

The `db2setup` utility works with Bash, Bourne, and Korn shells. Other shells are not supported.

Updating Kernel Configuration Parameters

This section applies only to DB2 clients running on HP-UX and NUMA-Q/PTX systems, and the Solaris** Operating Environment.

If you are installing a DB2 client on AIX, Linux, or SGI IRIX systems, go to “Installing Your DB2 Client” on page 153.

Before installing your DB2 client on a HP-UX, PTX or Solaris system, you may need to update your system’s kernel configuration parameters. We recommend that you set your system’s kernel configuration parameters to the values in the following sections:

- “HP-UX Kernel Configuration Parameters”
- “NUMA-Q/PTX Kernel Configuration Parameters” on page 151
- “Solaris Kernel Configuration Parameters” on page 152

HP-UX Kernel Configuration Parameters

Table 10 lists the recommended values for HP-UX kernel configuration parameters. These values are valid for DB2 clients running on HP-UX Version 11.

Note: You must reboot your machine after updating any kernel configuration parameters.

Table 10. 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 10.

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.

After you have updated your kernel parameters, go to “Installing Your DB2 Client” on page 153 to continue with the installation.

NUMA-Q/PTX Kernel Configuration Parameters

Table 11 lists the recommended values for NUMA-Q/PTX kernel configuration parameters for DB2 clients.

Note: You must reboot your machine after updating any kernel configuration parameters.

Table 11. NUMA-Q/PTX Kernel Configuration Parameters (Recommended Values)

Kernel Parameter	Recommended Value
msgmax	65535
msgmnb	65535
msgseg	8192
msgssz	16

To modify your PTX kernel configuration parameters, perform the following steps:

- Step 1. Log on as a user with root authority.
- Step 2. Enter the menu command.
- Step 3. Press A to select the **System Administration** option.
- Step 4. Press C to select the **Kernel Configuration** option.
- Step 5. Press Ctrl+F in the **Change Kernel Configuration Disk** form. If you want to build the new kernel in a disk other than the root disk, enter the disk and press Ctrl+F.
- Step 6. In the **Compile, Configure, or Remove a Kernel** window, select the kernel configuration type with which your current kernel is built and press K.
- Step 7. In the **Configure a kernel with site specific parameters** form, go down one page (press Ctrl+D), press A for **All** in the **Visibility level for parameter changes** and then press Ctrl+F.
- Step 8. In the **Configure Files With Adjustable Parameters** window select **ALL** (press Ctrl+T) and press Ctrl+F.

- Step 9. In the Tunable Parameters window, use arrow keys to navigate. Press Ctrl+T to select the parameter you want to change and press Ctrl+F.
- Step 10. In the Detail of Parameter Expression(s) window, press s to set the new value.
- Step 11. In the **Add site specific 'set' parameter** form, enter the new value and press Ctrl+F.
- Step 12. Repeat steps 9 through 11 on page 92 to change the values of all other parameters that you want to change.
- Step 13. Once you are done changing all the parameters, press Ctrl+E from the Tunable Parameters window.
- Step 14. Compile the kernel.
- Step 15. Press Ctrl+X to exit the menu.
- Step 16. Reboot the system so that the changes can take effect.

Notes:

1. The msgmax and msgmnb must be set to 65535 or larger.
2. The msgseg parameter must be set no higher than 32767.
3. The shmmax must be set to 2147483647 or larger.

After you have updated your kernel parameters, go to “Installing Your DB2 Client” on page 153 to continue with the installation.

Solaris Kernel Configuration Parameters

Table 12 lists the recommended values for Solaris kernel configuration parameters.

Note: You must reboot your machine after updating any kernel configuration parameters.

Table 12. 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 you have updated your kernel parameters, go to “Installing Your DB2 Client” to continue with the installation.

Installing Your DB2 Client

After updating your kernel configuration parameters and rebooting your system (if necessary), you can install your DB2 client.

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.

To install a DB2 client:

1. Log in as user with root authority.
2. Insert and mount the appropriate CD-ROM. For information on mounting CD-ROMs, see “Mounting CD-ROMs on UNIX Operating Systems” on page 200.
3. Change to the directory where the CD-ROM is mounted by entering the **cd /cdrom** command where **/cdrom** is the CD-ROM mount point.
4. Change to one of the following directories:

AIX /cdrom/db2/aix

HP-UX Version 11 /cdrom/db2/hpux11

Linux /cdrom/db2/linux

NUMA-Q/PTX /cdrom/db2/numaq

SGI/IRIX /cdrom/db2/sgi

Solaris /cdrom/unnamed_cdrom/db2/solaris

5. Enter the **./db2setup** command. After a few moments the Install DB2 V7 window opens.
6. Select the product that you want to install and select **OK**.

Press **Tab** to move between available options and fields. Press **Enter** to select or deselect an option. Selected options are denoted by an asterisk.

When you select to install a DB2 product, you can choose the product's **Customize** option to view and change the components that will be installed.

Select **OK** to continue the installation process or **Cancel** to go back to a previous window. Select **Help** for more information or assistance during the installation of any DB2 product.

When installation is complete DB2 software will be installed in the *DB2DIR* directory,

where <i>DB2DIR</i>	= /usr/lpp/db2_07_01	on AIX
	= /opt/IBMDB2/V7.1	on HP-UX, NUMA-Q/PTX, SGI IRIX, or Solaris
	= /usr/IBMDB2/V7.1	on Linux

You can use the **db2setup** program add additional products or components after your initial installation. To create or add additional DB2 products and components, enter the following command:

On AIX

```
/usr/lpp/db2_07_01/install/db2setup
```

On HP-UX, PTX, SGI IRIX, or Solaris

```
/opt/IBMDB2/V7.1/install/db2setup
```

On Linux

```
/usr/IBMDB2/V7.1/install/db2setup
```

Your next step

After installing your DB2 client, you should configure it to access a remote DB2 server. For more information see “Chapter 16. Configuring Client-to-Server Communications Using the Command Line Processor” on page 183.

Chapter 14. Control Center Installation and Configuration

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

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

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

Application versus Applet

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

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

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

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

On all other operating systems you must install the correct JRE level before running the Control Center. See Table 14 on page 157 for a list of correct JRE levels.

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

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

You must use a supported Java-enabled browser to run the Control Center as a Java applet. See Table 14 on page 157 for a list of supported browsers.

Machine Configurations

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

Table 13. Control Center Machine Configuration Scenarios

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

Figure 4 on page 157 summarizes the four basic Control Center machine configurations:




Figure 4. DB2 Control Center machine configurations

Supported Java Virtual Machines for the Control Center

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

Table 14. Supported Java Virtual Machines (JVMs) for the Control Center

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

Table 14. Supported Java Virtual Machines (JVMs) for the Control Center (continued)

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

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

Setting Up and Working with the Control Center

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

Control Center Services Setup (Applet Mode only)

If you will run the Control Center as an application, skip this section and go to “Running the Control Center as a Java Application” on page 160.

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

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

1. Start the Control Center JDBC Applet Server

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

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

The first time you start the Control Center JDBC Applet Server, it will create several node directory entries, together with various files for administration purposes. In Scenarios 1 and 3 in “Machine Configurations” on page 156, all of these administration files and directory entries will be created in the current DB2 instance.

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

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

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

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

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

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

```
net start DB2ControlCenterServer
```

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

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

2. Start the Windows NT or Windows 2000 Security Server

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

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

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

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

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

Working with the Control Center

You can run the Control Center as a Java application or as a Java applet. If your environment is configured similar to Scenarios 1 or 2 in Table 13 on page 156, you must run the Control Center as an application. If your environment is configured like Scenarios 3 or 4, you must run it as an applet.

Running the Control Center as a Java Application

To run the Control Center as a Java application, you must have the correct Java Runtime Environment (JRE) installed. See Table 14 on page 157 for the correct JRE level for your operating system.

1. To start the Control Center as an application:

On Windows 32-bit operating systems:

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

On OS/2:

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

On all supported platforms:

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

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

Running the Control Center as a Java Applet

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

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

`http://yourserver/temp`

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

To run Control Center as an applet on Windows 32-bit or OS/2 operating systems, you must run **db2classes.exe** on the machine where the DB2 JDBC Applet Server resides to unzip the required Java class files. On UNIX-based systems, you must uncompress and untar **db2classes.tar.Z** to make the required Java class files available.

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

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

Customizing Your Control Center HTML File

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

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

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

to

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

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

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

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

Configuring Your Web Server to Work with the Control Center

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

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

Functional Considerations

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

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

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

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

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

Installation Tips for Control Center Help on UNIX Operating Systems

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

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

- If you manually install the Control Center on UNIX-based workstations rather than using the db2setup utility, you must run the **db2insthtml** command to install the online documentation. For more information, see your server's installation chapter.

Configuring TCP/IP on OS/2

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

Enabling Local Loopback

To enable local loopback on your system:

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

Enabling Localhost

To enable localhost on your system:

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

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

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

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

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

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

Verifying TCP/IP Configuration on OS/2

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

Troubleshooting Information

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

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

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

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

- Verify that the correct JRE is installed. See Table 14 on page 157 for more information.

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

- Verify that you are running a supported browser. See Table 14 on page 157 for more information.
- Check your browser's Java console window for diagnostic and trace information for the Control Center.
- Make sure that the client browser does not have CLASSPATH set. To make sure CLASSPATH is not set, open a command window and enter **SET CLASSPATH=** then start your client browser from this command window. Also, note that if CLASSPATH is not set in a Windows NT or Windows 2000 environment, it may still get picked up from autoexec.bat from a Windows 9x installation on the same machine.
- Ensure that you are using the db2cc.htm file from the machine running the Control Center JDBC Applet Server.
- Remember that the Control Center works within the DB2 client's locale, and that the DB2 client is located at the Control Center JDBC Applet Server's location.

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

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

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

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

Database objects are presented in a consistent fashion for all DB2 family servers. This greatly reduces the amount of learning that is required for administrators who need to manage both DB2 for OS/390 and DB2 Universal Databases on Windows NT, Windows 2000, UNIX and OS/2 servers. While

the Control Center preserves consistency across servers it does not hide capabilities that are unique to each DB2 server. This gives database administrators the power to perform all aspects of their tasks.

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

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

Preparing DB2 for OS/390 Servers for the Control Center

The DB2 Control Center uses stored procedures to deliver many of its management functions. Therefore, for the Control Center to function properly each DB2 for OS/390 server that will be managed from the Control Center needs to have stored procedures enabled and the proper stored procedures installed on that server.

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

Working with the Control Center

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

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

Other Sources of Information

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

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

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

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

Part 4. Configuring DB2 Communications

Chapter 15. Configuring Client-to-Server Communications Using the Client Configuration Assistant

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

Notes:

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

LDAP Directory Support Considerations

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

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

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

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

Before You Begin

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

To complete the steps in this section, you should be familiar with how to start the CCA. For more information, see “Starting the Client Configuration Assistant” on page 195.

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

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

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

Configuration Steps

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

Step 1. Log on to the system with a valid DB2 user ID. For more information, see "Appendix C. Naming Rules" on page 207.



If you are adding a database to a system that has a DB2 server or DB2 Connect server product installed, log on to this system as a user with System Administrative (SYSADM) or System Controller (SYSCTRL) authority on the instance. For more information, see "Working with the System Administrative Group" on page 199.

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 195.

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

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

You can use one of the following configuration methods:

- "Adding a Database Using a Profile" on page 173.
- "Adding a Database Using Discovery" on page 174.
- "Adding a Database Manually" on page 176.

Adding a Database Using a Profile

A server profile contains information about server instances on a system, and databases within each server instance. For information on profiles, see “Creating and Using Profiles” on page 178.

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

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

Note: ODBC must be installed to perform this operation.

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

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

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

Adding a Database Using Discovery



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

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

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

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

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



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

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

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

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

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

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

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

Note: ODBC must be installed to perform this operation.

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

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

Adding a Database Manually

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

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

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

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

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

Click **Next**.

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

Step 5. Enter the database alias name of the remote database that you want to add in the **Database name** field and a local database alias name in the **Database alias** field.

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

Click **Next**.

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

Note: ODBC must be installed to perform this operation.

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

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

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



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

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

Creating and Using Profiles

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

Server Profiles

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



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

To create a server profile, perform the following steps:

- Step 1. Start the Control Center. For more information, refer to “Starting the DB2 Control Center” on page 195.
- Step 2. Select the system that you want to create a profile for and right click. If the system that you want to create a profile for is not shown, select the **Systems** icon, click the right mouse button, and select the **Add** option. Click the **Help** push button and follow the online help.
- Step 3. Select the **Export Server Profile** option.
- Step 4. Enter a path and filename for this profile and select **OK**.



You are ready to use this profile on your system. For more information on how to add a database to your system using a server profile, go to “Configuration Steps” on page 172.

Client Profiles

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

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

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

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

To create a client profile, perform the following steps:

- Step 1. Start the CCA. For more information, see “Starting the Client Configuration Assistant” on page 195.
- Step 2. Click **Export**. The Select Export Option window opens.
- Step 3. Select one of the following export options:

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

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



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

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

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

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

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

Step 8. Click **OK**.

To import a client profile, perform the following steps:

Step 1. Start the CCA. For more information, see “Starting the Client Configuration Assistant” on page 195.

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

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

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

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

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

Step 5. Click **OK**.



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

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

Note: ODBC must be installed to perform this operation.

- 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.
 - If you would like to create an ODBC data source file to share database access, select the **As a file data source** radio button and enter the path and file name for this file in the **File data source name** field.
 - Click the **Optimize for application** drop down box and select the application for which you want to tune the ODBC settings.
 - Click **Finish** to add the database that you selected. The Confirmation window opens.
- Step 9. Click the **Test Connection** push button to test the connection. The Connect to DB2 Database window opens.
- Step 10. In the Connect To DB2 Database window, enter a valid user ID and password for the remote database and click **OK**. If the connection is successful, a message confirming the connection appears.
- If the connection test failed, you will receive a help message. To change any settings you may have incorrectly specified, click the **Change** push button in the Confirmation window to return to the

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

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

Chapter 16. Configuring Client-to-Server Communications Using the Command Line Processor

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

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 OS/2 or Windows 32-bit client for communications. For more information see “Chapter 15. Configuring Client-to-Server Communications Using the Client Configuration Assistant” on page 171.

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

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

For instructions on entering DB2 commands, see “Entering Commands Using the Command Center” on page 196 or “Entering Commands Using the Command Line Processor” on page 197.

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 19 for the communication protocol requirements for your platform. See “Possible Client-to-Server Connectivity Scenarios” on page 26 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 15. TCP/IP Values Required at the Client

Parameter	Description	Sample Value	Your Value
Host Name	Use the <i>hostname</i> or <i>ip_address</i> of the remote server workstation.	serverhost	
• Hostname (<i>hostname</i>) or		or	
• IP address (<i>ip_address</i>)	To resolve this parameter: <ul style="list-style-type: none">• 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 hostname command.• On UNIX systems, you can also use the DB2/bin/hostlookup hostname command, where <i>DB2</i> is the directory where DB2 is installed.	9.21.15.235	

The client must know the IP 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 16 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 16. 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 or Windows 2000	Located in the <code>winnt\system32\drivers\etc</code> directory.
Windows 9x	Located in the <code>windows</code> directory.
UNIX	Located in the <code>/etc</code> 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 187.

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

tcp represents the communication protocol that you are using

represents a comment describing the entry

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 services file 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 185.

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

See Table 16 on page 186 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 C. Naming Rules” on page 207.



If you are adding a database to a system that has a DB2 server or DB2 Connect server product installed, log on to this system as a user with System Administrative (SYSADM) or System Controller (SYSCTRL) authority on the instance. For more information, see “Working with the System Administrative Group” on page 199.

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, 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:

```
db2 "catalog tcpip node node_name remote [hostname|ip_address]  
server [svcname|port_number]"  
db2 terminate
```

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

```
db2 catalog tcpip node db2node remote serverhost server server1  
db2 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:

```
db2 catalog tcpip node db2node remote 9.21.15.235 server 3700  
db2 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:

```
db2 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. By default, 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 C. Naming Rules” on page 207.



If you are adding a database to a system that has a DB2 server or DB2 Connect server product installed, log on to this system as a user with System Administrative (SYSADM) or System Controller (SYSCTRL) authority on the instance. For more information, see “Working with the System Administrative Group” on page 199.

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 17. 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>), unless specified otherwise.	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>). The database alias the name that you use when connecting to a database from a client.	tor1	
Authentication (<i>auth_value</i>)	The value of the authentication required by your enterprise. Please refer to the <i>DB2 Connect User's Guide</i> for more information on this parameter.	DCS This means that the userid and password supplied are validated at the host or AS/400 only.	
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, run the start-up script as follows:

```
. INSTHOME/sql11b/db2profile (for bash, Bourne or Korn shell)
source INSTHOME/sql11b/db2cshrc (for C shell)
```

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

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

```
db2 catalog database database_name as database_alias at node node_name
db2 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:

```
db2 catalog database sample as tor1 at node db2node
db2 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:

```
db2 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

After configuring the client for communications, 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. If you are using a UNIX client, 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* represents the home directory of the instance.

Step 3. Enter the following command on the client to connect the client to the remote database:

```
db2 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 CLP:

```
"select tabname from syscat.tables"
```

When you are finished using the database connection, enter the **command 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 (*svccname*) parameter was updated correctly in the database manager configuration file.
4. The database was created and cataloged properly.
5. The database manager was stopped and started again (enter the **db2stop** and **db2start** commands on the server).
6. Ensure that you did not specify a port number that is being used by the Fast Communications 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 located in the *INSTHOME/sql11ib/db2dump* directory.

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 matches, or the services name maps 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 (the *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 Client Configuration Assistant”.
 - “Starting the DB2 Control Center”.
 - “Entering Commands Using the Command Center” on page 196.
 - “Entering Commands Using the Command Line Processor” on page 197.
 - “Working with the System Administrative Group” on page 199.
 - “Working with the Business Intelligence Functions” on page 199.
 - “Mounting CD-ROMs on UNIX Operating Systems” on page 200.
 - “Setting the Number of Licensed Processors” on page 202.
 - “Upgrading DB2 from Try and Buy Mode” on page 203.
-

Starting the Client Configuration Assistant

Start the Client Configuration Assistant (CCA) as follows:

OS/2 Click on **OS/2 Warp**, and select **IBM DB2** → **Client Configuration Assistant**

Windows 32-bit operating systems
Click on **Start** and select **Programs** → **IBM DB2** → **Client Configuration Assistant**

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

Starting the DB2 Control Center

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

To run the Control Center as an application

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

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

To run the Control Center as an applet

You must have a Java-enabled browser and you must perform some additional configuration steps to run the Control Center as an applet. For detailed instructions on running the Control Center as an applet or application, see “Chapter 14. Control Center Installation and Configuration” on page 155.

Entering Commands Using the Command Center

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

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

From the Command Center you can:

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

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

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



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

```
list database directory
```

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

```
!dir
```

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

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

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

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

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

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



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

Entering Commands Using the Command Line Processor

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

DB2 Command Window

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

Interactive Input Mode

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

File Input Mode

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

DB2 Command Window

To invoke a DB2 Command Window, do the following:

OS/2 Open any OS/2 command window.

Windows 32-bit operating systems

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

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

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"
```

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

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

Interactive Input Mode

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

OS/2 Click on **OS/2 Warp** and select **IBM DB2** → **Command Line Processor** or enter the **db2** command.

Windows 32-bit operating systems

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

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

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
```

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

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

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

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

Working with the System Administrative Group

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

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

Working with the Business Intelligence Functions

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

Mounting CD-ROMs on UNIX Operating Systems

The following sections describe how to mount your DB2 product CD-ROM on UNIX-based operating systems.

Mounting the CD-ROM on AIX

To mount the CD-ROM on AIX using the System Management Interface Tool (SMIT), perform the following steps:

1. Log in as a user with root authority.
2. Insert the CD-ROM in the drive.
3. Create a CD-ROM mount point by entering the `mkdir -p /cdrom` command, where `cdrom` represents the CD-ROM mount point directory.
4. Allocate a CD-ROM file system using SMIT by entering the `smit storage` command.
5. After SMIT starts, select **File Systems** → **Add / Change / Show / Delete File Systems** → **CDROM File Systems** → **Add CDROM File System**.
6. In the Add a File System window:
 - Enter a device name for your CD-ROM file system in the **DEVICE Name** field. 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. In our example, we will use `/dev/cd0` as the device name.
 - Enter the CD-ROM mount point directory in the **MOUNT POINT** window. In our example, the mount point directory is `/cdrom`.
 - In the **Mount AUTOMATICALLY at system restart** field, select yes to enable automatic mounting of the file system.
 - Click **OK** to close the window, then click **Cancel** three times to exit SMIT.
7. Next, mount the CD-ROM file system by entering the `smit mountfs` command.
8. In the Mount a File System window:
 - Enter the device name for this CD-ROM file system in the **FILE SYSTEM name** field. In our example, the device name is `/dev/cd0`.
 - Enter the CD-ROM mount point in the **Directory over which to mount** field. In our example, the mount point is `/cdrom`.
 - Enter `cdrfs` in the **Type of Filesystem** field. To view the other kinds of file systems you can mount, click **List**.
 - In the **Mount as READ-ONLY system** field, select to yes.
 - Accept the remaining default values and click **OK** to close the window.

Your CD-ROM file system is now mounted. To view the contents of the CD-ROM, place the disk in the drive and enter the **cd /cdrom** command where **cdrom** is the CD-ROM mount point directory.

Mounting the CD-ROM on HP-UX

Because DB2 Version 7.1 for HP-UX contains several files with long file names, the mount command may fail. The following steps will enable you to mount successfully your DB2 for HP-UX product CD-ROM:

1. Log in as a user with root authority.
2. In the /etc directory, add the following line to the pfs_fstab file:

```
/dev/dsk/c0t2d0 mount_point pfs-rrip ro,hard
```

where *mount_point* represents the mount point of the CD-ROM.

3. Start the *pfs* daemon by entering the following commands (if they are not already running):

```
/usr/sbin/pfs_mountd &  
/usr/sbin/pfsd 4 &
```

4. Insert the CD-ROM in the drive and enter the following commands:

```
mkdir /cdrom  
/usr/sbin/pfs_mount /cdrom
```

where */cdrom* represents the mount point of the CD-ROM.

5. Log out.

Mounting the CD-ROM on Linux

To mount the CD-ROM on Linux:

1. Log in as a user with root authority.
2. Insert the CD-ROM in the drive and enter the following command:

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

where */cdrom* represents the mount point of the CD-ROM.

3. Log out.

Note that some window managers may automatically mount your CD-ROM for you. Consult your system documentation for more information.

Mounting the CD-ROM on PTX

To mount the CD-ROM on PTX:

1. Log in as a user with root authority.
2. Insert the CD-ROM in the drive and enter the following commands:

```
mkdir /cdrom  
mount -r -F cdfs /dev/dsk/cd0 /cdrom
```

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

3. Log out.

Mounting the CD-ROM on Solaris

To mount the CD-ROM on Solaris:

1. Log in as a user with root authority.
2. Insert the CD-ROM into the drive.
3. 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 and `/dev/dsk/c0t6d0s2` represents the CD-ROM drive device.

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
```

4. Log out.

Setting the Number of Licensed Processors

Note: This section applies only to DB2 Enterprise Edition, DB2 Enterprise-Extended Edition, and DB2 Warehouse Manager.

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

Note: An instance must be created on UNIX-based operating systems before performing these steps.

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

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

3. Obtain the product password by issuing the **db2licm -l** command. The DB2 products are listed as follows:

Enterprise-Extended Edition

DB2UDBEEE DB

Enterprise Edition

DB2UDBEE DB2

Warehouse Manager

DB2UDBWM DB2

Relational Connect

DB2RELC DB2

Spatial Extenders

DB2UDBGSE

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

Upgrading DB2 from Try and Buy Mode

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

On UNIX Operating Systems

To add a license using the command line:

1. Log in as a user with root authority.
2. The **db2licm** utility can be found in the following locations:
 - For AIX, `/usr/lpp/db2_07_01/adm/`
 - For LINUX, `/usr/IBMdb2/V7.1/adm/`
 - For HP-UX, PTX, Solaris, PTX, `/opt/IBMdb2/V7.1/adm/`
3. Issue the following command:

```
db2licm filename.lic
```

where *filename.lic* represents the name of the licensing file. For more information about the **db2licm** command, refer to the *Command Reference*.

On OS/2 and Windows 32-bit Operating Systems

To add a license using the command line:

1. Go to the `<install directory>\bin` directory, where *<install directory>* represents the directory where you installed the product.

2. To add a license, issue the following command:

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

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

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

Note: You cannot view specific license information using the **db2licm** utility until at least one instance is created.

Adding a license using the License Center

To add a license using the License Center:

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

Appendix B. 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, perform the following steps:

- Step 1. Ensure that the appropriate message option for the desired language has been installed.
- Step 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_07_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 18 on page 206 shows the languages into which the DB2 messages are translated. If the setup is run on a machine that is set up using a non-supported language, then English is the default unless the user has specified otherwise.

Table 18. 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 C. Naming Rules



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

- “General Naming Rules”
 - “Database, Database Alias, and Catalog Node Name Rules”
 - “Object Name Rules” on page 208
 - “Username, User ID, Group Name, and Instance Name Rules” on page 209
 - “DB2SYSTEM Naming Rules” on page 211
 - “Password Rules” on page 211
-

General Naming Rules

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

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

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

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

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

Database, Database Alias, and Catalog Node Name Rules

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

When naming a database, database alias, or catalog node name, see “General Naming Rules” on page 207. 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 207.

In addition, the name you specify:

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

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

For example, if you create a column with a + or – sign included in the name and you subsequently use that column in an index, you will experience

problems when you attempt to reorganize the table. To avoid potential problems with the use and operation of your database, *do not* violate these rules.

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

Usernames or *User IDs* are the identifiers assigned to individual users. When naming users, groups, or instances, see “General Naming Rules” on page 207.

In addition to the general naming rules:

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

Workstation Name (nname) Rules

A *workstation* name specifies the NetBIOS name for a database server or client that resides on the local workstation. This name is stored in the database manager configuration file. The workstation name is known as the *workstation nname*. When naming workstations, see “General Naming Rules” on page 207.

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 Windows Computer name and assigns it to *DB2SYSTEM*.

When creating a *DB2SYSTEM* name, see “General Naming Rules” on page 207.

In addition, the name you specify:

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

Password Rules

When determining passwords, consider the following rules:

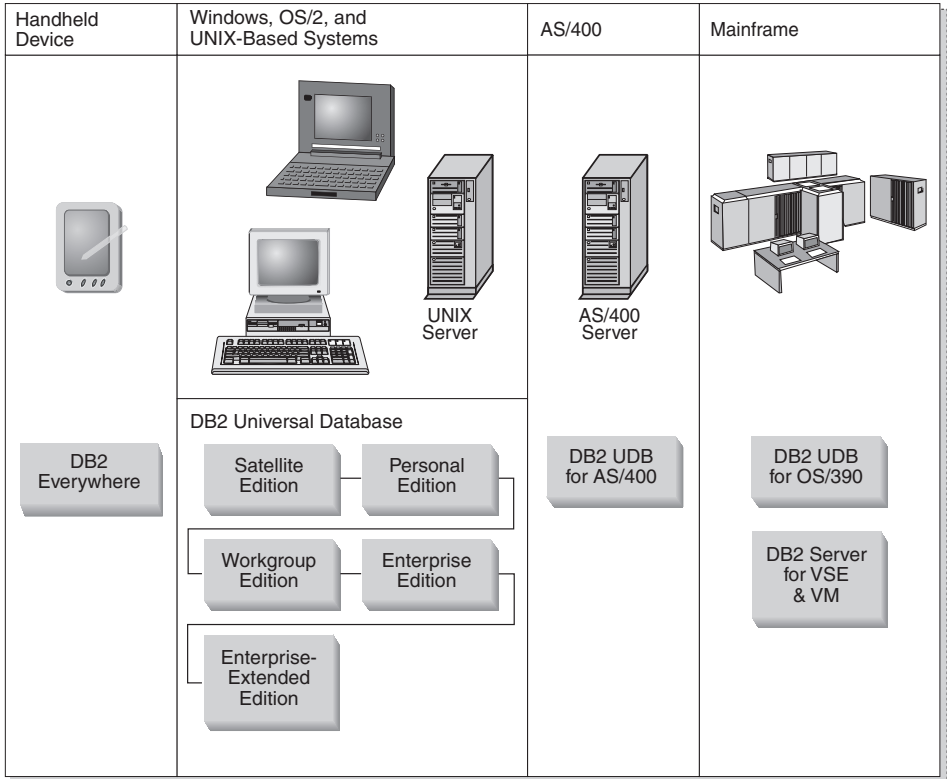
OS/2 A maximum of 14 characters.

UNIX A maximum of 8 characters.

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

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

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



DB2 Products

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

DB2 Everywhere

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

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

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

DB2 Universal Database

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

Table 19. DB2 Universal Database Platforms

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

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

Satellite Edition

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

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

Personal Edition

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

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

- An application development client
- An administration client

Workgroup Edition

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

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

Enterprise Edition

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

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

Enterprise - Extended Edition

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

- Support for clusters of servers

Developer Editions

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

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

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

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

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

- VisualAge for Java, Professional Edition


Host Databases

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

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

DB2 Connect

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



The following DB2 Connect products are available:

- Personal Edition
- Enterprise Edition
- Unlimited Edition

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

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

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

Related Products

The following products work together with DB2 Universal Database.

DB2 Relational Connect

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

DB2 Warehouse Manager

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

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

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

DB2 OLAP Server

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

With DB2 OLAP Server you can:

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

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

Intelligent Miner

The Intelligent Miner family consists of two products:

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

DB2 Spatial Extender

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

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

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

DB2 Net Search Extender

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

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

DB2 Data Links Manager

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

Tivoli Enterprise

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

Working with DB2 Data

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

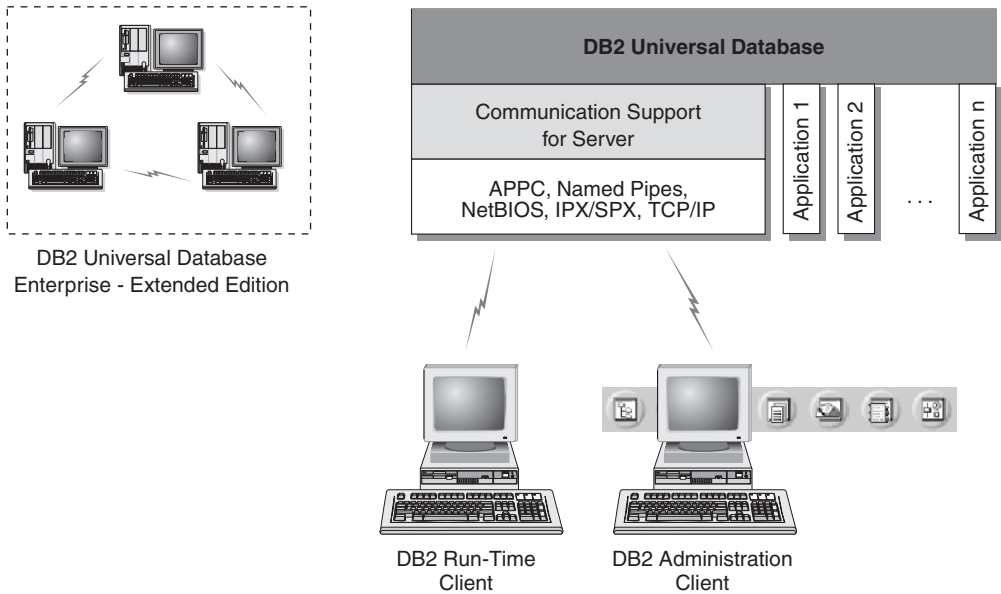
Accessing DB2 Data from Remote Clients

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

DB2 Version 7 clients are supported on:

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

Figure 5 on page 220 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.



Not all protocols are supported for all platforms.

Figure 5. DB2 Server with Local Applications and Remote Clients

Accessing Multiple DB2 Servers

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

For example, you can have a server on a Windows NT workstation that is connected to a LAN, and another server on a UNIX workstation connected to a LAN. As long as there is a connection between the two LANs, clients on either network can access either server. See Figure 6 on page 221.

Communicating Between DB2 Universal Database Servers

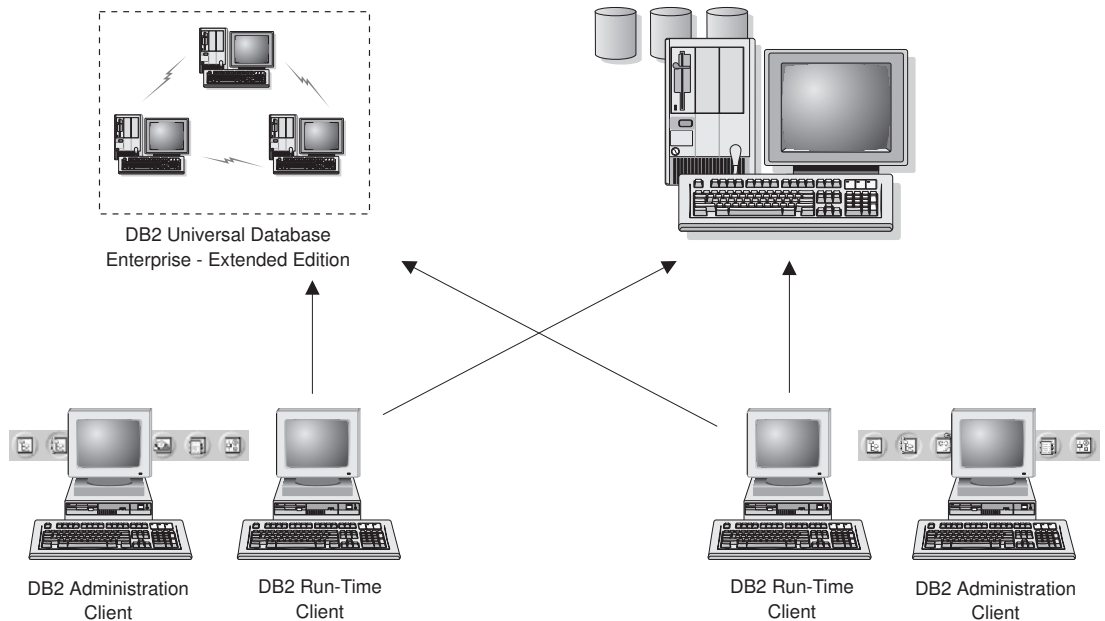


Figure 6. Accessing Data on Multiple Servers

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

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

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

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

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

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

DB2 Connect Enterprise Edition is most appropriate for environments where:

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

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

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

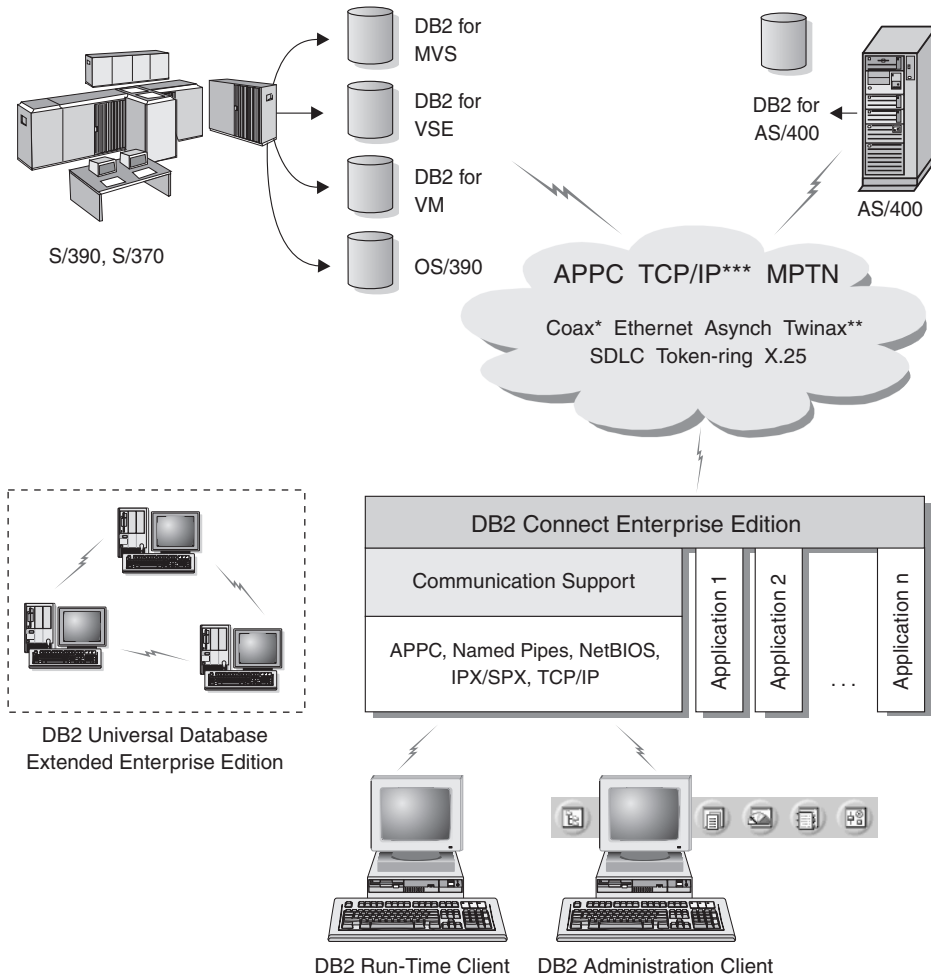
DB2 Connect may also be installed on a Web server, Transaction Processor (TP) monitor, or other 3-tier application server machines with multiple local SQL application processes and threads. In these cases, you can choose to install DB2 Connect on the same machine for simplicity, or on a separate machine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

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

to enterprise data. Figure 7 on page 224 illustrates IBM's solution for environments in which you want to use a DB2 client making an indirect connection to a host or AS/400 database server through DB2 Connect Enterprise Edition.

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

DB2 Connect Enterprise Edition



Not all protocols are supported for all platforms.

* For Host connections only

** For AS/400

*** TCP/IP connectivity requires DB2 for OS/390 V5R1, DB2 for AS/400 V4R2, or DB2 for VM V6.1

Figure 7. DB2 Connect Enterprise Edition

Accessing DB2 Data from the Web using Java

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

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

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

The DB2 JDBC Applet server and the DB2 client must reside on the same machine as the Web server. The DB2 JDBC Applet server calls the DB2 client to connect to local, remote, host, and AS/400 databases. When the applet requests a connection to a DB2 database, the JDBC client opens a TCP/IP connection to the DB2 JDBC Applet on the machine where the Web server is running. See Figure 8 on page 226 for an example of a Java-enabled browser accessing data from remote DB2 databases.

Accessing DB2 Data Using JDBC

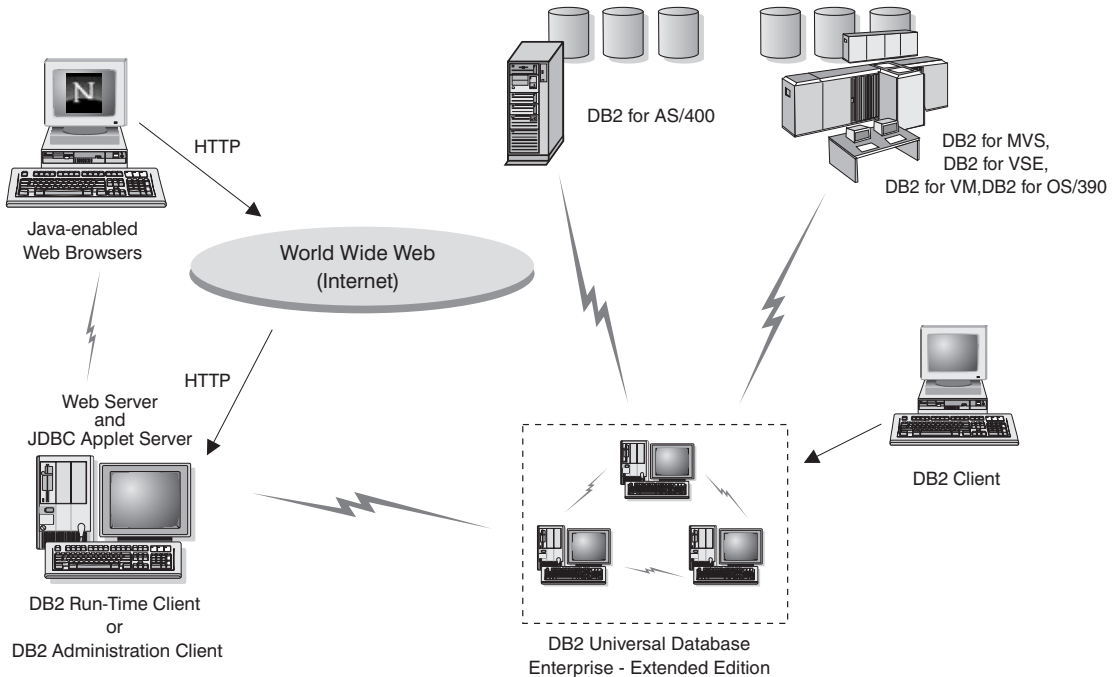


Figure 8. Accessing Data Stored on DB2 using JDBC

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

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

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

Accessing DB2 Data from the Web using Net.Data

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

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

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

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

Net.Data can be installed with a:

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

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




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

Accessing DB2 Data from Host and AS/400 Client Machines

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



The DRDA AS feature is available for:

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

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

Administering Instances and Databases with the DB2 Administration Tools

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

Managing Instances and Database Objects using the Control Center

The Control Center displays instances and database objects (such as table spaces, tables, and packages) and their relationships to each other. Using the Control Center, you can manage local and remote servers from a single point of control. See Figure 10 on page 229 for an example of the main Control Center window.

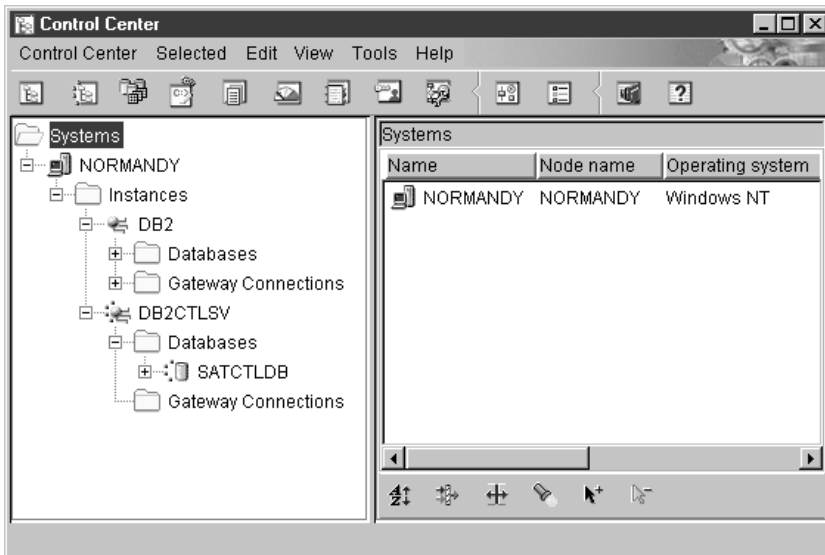


Figure 10. Control Center Main Window

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

From the Control Center, you can perform operations on database objects. These operations include:

- Create and drop a database
- Create, alter, and drop a table space or table
- Create, alter, and drop an index
- Create nodegroups
- Backup and recover a database partition or table space partition.
- Define the replication sources and subscriptions to replicate data between systems
- Monitor resources and events on a server.

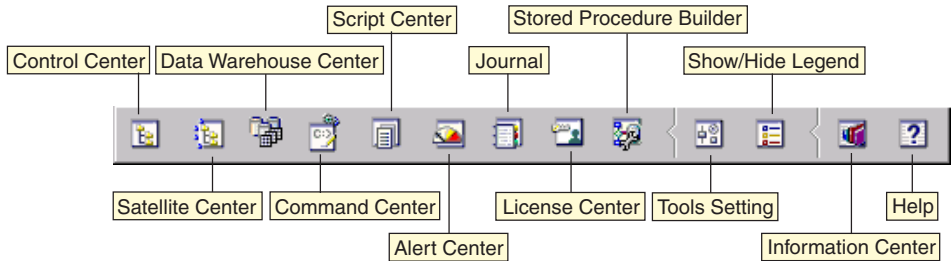
For more information on objects in a partitioned database system, see "Chapter 1. Introduction to DB2 Enterprise - Extended Edition" on page 3.

You can also control DB2 instances by:

- Maintaining communication protocols
- Setting database manager and database configuration values that affect performance.

Wizards are provided to help you perform complex tasks. For example, a wizard is available to tune the performance of your system. See “Using DB2 Wizards” on page 252 for descriptions of the various wizards and how to start them.

The Control Center provides additional functionality to assist you in managing your servers:



Control Center

Use the Control Center to start another session of the Control Center to administer a server.

Satellite Center

Use the Satellite Center to manage the Satellites that are served by a particular DB2 Control Server. It provides create, remove, modify, and manage functions for Satellites and Groups. You can also create and manage scripts to administer the Satellites.

Data Warehouse Center

Use the Data Warehouse Center to manage warehouses; create and manage warehouse objects, such as sources and targets; define, extract, transform, manipulate, and load steps and processes; and schedule and automate steps.

Command Center

Use the Command Center to enter DB2 commands and SQL statements in an interactive window and see the execution result in a result window. You can scroll through the results and save the output to a file.

Script Center

Use the Script Center to create scripts, which you can store and invoke at a later time. These scripts can contain DB2 commands, SQL statements, as well as operating system commands. Scripts can be scheduled to run unattended. These jobs can be run once or set up to run on a repeating schedule; a repeating schedule is particularly useful for tasks like backup.

Alert Center

Use the Alert Center to monitor your system for early warnings of potential problems or to automate actions to correct problems discovered.

Journal

Use the Journal to view all available information about jobs that are pending execution, executing, or that have completed execution. You can also view the recovery history log, the alerts log, and the messages log; and review the results of jobs that are run unattended.

License Center

Use the License Center to manage licenses and display license status and usage of any DB2 products installed on your system. You can also use the License Center to configure your system for proper license monitoring.

Stored Procedure Builder

Use Stored Procedure Builder to create stored procedures, build stored procedures on local and remote DB2 servers, modify and rebuild existing stored procedures, and run stored procedures for testing and debugging the execution of installed stored procedures.

Tools Setting

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

Information Center

The Information Center provides quick access to DB2 product information. This product information includes such items as: database tasks, reference material, DB2 documentation, warehouse administration information, troubleshooting aids, sample programs for application development, and DB2 web-related URLs.

You can also analyze performance using the DB2 Performance Monitor and Visual Explain. These tools are available from the Control Center.



Use the **DB2 Performance Monitor** to monitor the performance of your system. You can monitor activity by sampling data over a period of time or using data for a particular event. See “Monitoring Databases using DB2 Performance Monitor” on page 232 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 232 for more information.

You can find additional information in the *Administration Guide* or in the online help.

Managing Communications on the Server

The Control Center allows you to maintain or edit server protocol settings in the database manager configuration file.

- Configure database manager parameters by clicking with the right mouse button on an instance and selecting the **Configure** option from the pop-up menu. By default, the setup program automatically detects and configures most communication protocols that it detects on your system.
- Export database information in a profile that can be used to configure clients by clicking with the right mouse button on a system and selecting the **Export Server Profile** option from the pop-up menu.

For information on how to configure server communications, refer to the *Installation and Configuration Supplement*.

Monitoring Databases using DB2 Performance Monitor

With the DB2 Performance Monitor, you can:

- Identify and analyze performance problems in database applications or the database manager.
- Use the early warning system to detect potential problems.
- Automate actions to correct problems that are discovered.
- Define your own statistics, in addition to the default set that is provided.

You can choose to monitor the current state of database activity or collect information when specific events occur. The Performance Monitor allows you to capture point-in-time information at specified intervals. The Event Analyzer allows you to view information about the occurrence of events such as deadlocks and transaction completions.

For additional information, refer to the *Administration Guide* or the online help. You are also able to use the Windows Performance Monitor (supported on Windows NT and Windows 2000) to monitor both database and system performance. For information on how to register DB2 resources and to use the Windows Performance Monitor, refer to the *Administration Guide*.

Viewing SQL Access Plans using Visual Explain

Visual Explain helps database administrators and application developers to:

- View the access plan chosen by the database manager's optimizer for a given SQL statement.
- Tune SQL statements for better performance.
- Design application programs and databases.

- View all the details of an access plan, including the statistics in the system catalogs.
- Decide whether or not to add an index to a table.
- Identify the source of problems by analyzing the access plan or performance of SQL statements.
- Use the portable snapshot function to view snapshots from any remote DB2 server.
- Display access plans for queries on all supported DB2 configurations.

For additional information, refer to the *Administration Guide* or the online help.

Managing Connections to Databases using the Client Configuration Assistant

The Client Configuration Assistant (CCA) helps you manage your database connections to remote servers. The CCA is available on OS/2 and Windows 32-bit operating systems and this is the preferred method to set up any OS/2, Windows 9x, Windows NT, or Windows 2000 client to communicate with a server.

You can use the command line processor to set up DB2 clients on any platform. For more information, see “Chapter 16. Configuring Client-to-Server Communications Using the Command Line Processor” on page 183.

With the CCA, you can:

- Catalog databases so that they can be used by applications. Three methods are available:
 - Use a profile provided by a database administrator to automatically define your connections. Client access is automatically set up for that database.
 - Search the network for available databases and select one. Client access is automatically set up for that database.
 - Manually configure a connection to a database by entering the required connection parameters.
- Remove cataloged databases, or change the properties of a cataloged database.
- Export and import client profiles that contain database and configuration information for a client.
- Test connections to local or remote databases identified on your system.
- Bind applications to a database by selecting utilities or bind files from a list.
- Tune the client configuration parameters on your system. Parameters are logically grouped and suggested settings are provided on the interface as parameters are selected.

- Export client configuration information to a profile.
- Import configuration information from a profile.
- Update the server password.

Managing Warehouses using the Data Warehouse Center

DB2 Universal Database offers the Data Warehouse Center, a component that automates data warehouse processing. You can use the Data Warehouse Center to define the data to include in the warehouse. Then, you can use the Data Warehouse Center to schedule automatic refreshes of the data in the warehouse.

From the Data Warehouse Center, you can manage specific warehousing objects, including subject areas, warehouse sources, warehouse targets, agents, agent sites, steps, and processes.

You can also perform the following tasks from the Data Warehouse Center:

- Define a subject area. You use a subject area to logically group the processes that are related to a particular topic or function.
- Explore the source data and define warehouse sources.
- Create database tables and define warehouse targets.
- Define a process that specifies how to move and transform the source data into the appropriate format for the warehouse.
- Test and schedule steps.
- Define security and monitor database currency.
- Define a star schema model.

Understanding the Administration Server

The Administration Server responds to requests from the DB2 Administration Tools and the Client Configuration Assistant (CCA). The DB2 Administration Tools allow you to start, stop, and set database manager configuration parameters for servers. The Administration Server is used by the CCA to catalog databases for a client.

The Administration Server (DAS) must reside on every server that you want to administer and detect. By default the DAS is DB2AS, which is the default user ID that is created using the db2setup utility.

Developing Applications using the DB2 Application Development Client

The DB2 Application Development Client is a collection of tools that are designed to meet the needs of database application developers. It includes libraries, header files, documented APIs, and sample programs to build character-based, multimedia, or object-oriented applications.

A platform-specific version of the DB2 Application Development Client is available on each server CD-ROM. In addition the Developer Edition boxes contain the Application Development Clients for multiple supported operating systems. The Personal Developer's Edition box contains the Application Development CD-ROMs for OS/2, Windows and Linux. The Universal Developer's Edition box contains the Application Development CD-ROMs for all supported operating systems.

Through a DB2 client, these applications can access all servers and, by using the DB2 Connect product (or the DB2 Connect functionality supplied with DB2 Enterprise - Extended or DB2 Enterprise Edition), they can also access DB2 Universal Database for AS/400, DB2 Universal Database for OS/390, and DB2 for VSE & VM database servers.

The DB2 Application Development Client allows you to develop applications that use the following interfaces:

- Embedded SQL
- Call Level Interface (CLI) development environment (which is compatible with ODBC from Microsoft)
- Java Database Connectivity (JDBC)
- Embedded SQL for Java (SQLJ)
- DB2 Application Programming Interfaces (APIs) that use administrative functions to manage a DB2 database.

The DB2 Application Development Client includes:

- Precompilers for Java, C, C++, COBOL, and FORTRAN.
- Libraries, include files, and code samples to develop applications that use SQLJ and DB2 CLI.
- A single control point for metadata management through the use of templates and tokens.
- JDBC and SQLJ support to develop Java applications and applets.
- Interactive SQL, through the CLP, to prototype SQL statements and perform ad-hoc database queries.
- An API to enable other application development tools to implement precompiler support for DB2 directly with their products.

- An SQL92 and MVS Conformance Flagger to identify embedded SQL statements in applications not conforming to the ISO/ANSO SQL92 Entry Level standard, or which are not supported by DB2 for OS/390.

For complete information on the functionality of the DB2 Application Development Client, and instructions on how to use them, as well as a full list of supported compilers for your platform, refer to the *Application Building Guide*.

Running Your Own Applications

Various types of applications can access DB2 databases:

- Applications developed using a DB2 Application Development Client that include embedded SQL (including Java SQLJ applications and applets), APIs, stored procedures, user-defined functions, calls to DB2 CLI, or calls to JDBC applications and applets.
- ODBC applications such as Lotus Approach.
- Net.Data macros containing HTML and SQL.

The DB2 CLI/ODBC driver is an optional component during a DB2 client install. It is required to run CLI, ODBC, JDBC, and some SQLJ applications.

For more information on running your own applications, refer to the *Installation and Configuration Supplement*.

Appendix E. Using the DB2 Library

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

To access product information online, you can use the Information Center. For more information, see “Accessing Information with the Information Center” on page 251. You can view task information, DB2 books, troubleshooting information, sample programs, and DB2 information on the Web.

DB2 PDF Files and Printed Books

DB2 Information

The following table divides the DB2 books into four categories:

DB2 Guide and Reference Information

These books contain the common DB2 information for all platforms.

DB2 Installation and Configuration Information

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

Cross-platform sample programs in HTML

These samples are the HTML version of the sample programs that are installed with the Application Development Client. The samples are for informational purposes and do not replace the actual programs.

Release notes

These files contain late-breaking information that could not be included in the DB2 books.

The installation manuals, release notes, and tutorials are viewable in HTML directly from the product CD-ROM. Most books are available in HTML on the product CD-ROM for viewing and in Adobe Acrobat (PDF) format on the DB2 publications CD-ROM for viewing and printing. You can also order a printed copy from IBM; see “Ordering the Printed Books” on page 247. The following table lists books that can be ordered.

On OS/2 and Windows platforms, you can install the HTML files under the `sql1lib\doc\html` directory. DB2 information is translated into different

languages; however, all the information is not translated into every language. Whenever information is not available in a specific language, the English information is provided

On UNIX platforms, you can install multiple language versions of the HTML files under the `doc/%L/html` directories, where `%L` represents the locale. For more information, refer to the appropriate *Quick Beginnings* book.

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

- “Viewing Information Online” on page 250
- “Searching Information Online” on page 254
- “Ordering the Printed Books” on page 247
- “Printing the PDF Books” on page 246

Table 20. DB2 Information

Name	Description	Form Number PDF File Name	HTML Directory
DB2 Guide and Reference Information			
<i>Administration Guide</i>	<i>Administration Guide: Planning</i> provides an overview of database concepts, information about design issues (such as logical and physical database design), and a discussion of high availability.	SC09-2946 db2d1x70	db2d0
	<i>Administration Guide: Implementation</i> provides information on implementation issues such as implementing your design, accessing databases, auditing, backup and recovery.	SC09-2944 db2d2x70	
	<i>Administration Guide: Performance</i> provides information on database environment and application performance evaluation and tuning.	SC09-2945 db2d3x70	
	You can order the three volumes of the <i>Administration Guide</i> in the English language in North America using the form number SBOF-8934.		
<i>Administrative API Reference</i>	Describes the DB2 application programming interfaces (APIs) and data structures that you can use to manage your databases. This book also explains how to call APIs from your applications.	SC09-2947 db2b0x70	db2b0

Table 20. DB2 Information (continued)

Name	Description	Form Number PDF File Name	HTML Directory
<i>Application Building Guide</i>	Provides environment setup information and step-by-step instructions about how to compile, link, and run DB2 applications on Windows, OS/2, and UNIX-based platforms.	SC09-2948 db2axx70	db2ax
<i>APPC, CPI-C, and SNA Sense Codes</i>	Provides general information about APPC, CPI-C, and SNA sense codes that you may encounter when using DB2 Universal Database products.	No form number db2apx70	db2ap
	Available in HTML format only.		
<i>Application Development Guide</i>	Explains how to develop applications that access DB2 databases using embedded SQL or Java (JDBC and SQLJ). Discussion topics include writing stored procedures, writing user-defined functions, creating user-defined types, using triggers, and developing applications in partitioned environments or with federated systems.	SC09-2949 db2a0x70	db2a0
<i>CLI Guide and Reference</i>	Explains how to develop applications that access DB2 databases using the DB2 Call Level Interface, a callable SQL interface that is compatible with the Microsoft ODBC specification.	SC09-2950 db2l0x70	db2l0
<i>Command Reference</i>	Explains how to use the Command Line Processor and describes the DB2 commands that you can use to manage your database.	SC09-2951 db2n0x70	db2n0
<i>Connectivity Supplement</i>	Provides setup and reference information on how to use DB2 for AS/400, DB2 for OS/390, DB2 for MVS, or DB2 for VM as DRDA application requesters with DB2 Universal Database servers. This book also details how to use DRDA application servers with DB2 Connect application requesters.	No form number db2h1x70	db2h1
	Available in HTML and PDF only.		

Table 20. DB2 Information (continued)

Name	Description	Form Number PDF File Name	HTML Directory
<i>Data Movement Utilities Guide and Reference</i>	Explains how to use DB2 utilities, such as import, export, load, AutoLoader, and DPROP, that facilitate the movement of data.	SC09-2955 db2dmx70	db2dm
<i>Data Warehouse Center Administration Guide</i>	Provides information on how to build and maintain a data warehouse using the Data Warehouse Center.	SC26-9993 db2ddx70	db2dd
<i>Data Warehouse Center Application Integration Guide</i>	Provides information to help programmers integrate applications with the Data Warehouse Center and with the Information Catalog Manager.	SC26-9994 db2adx70	db2ad
<i>DB2 Connect User's Guide</i>	Provides concepts, programming, and general usage information for the DB2 Connect products.	SC09-2954 db2c0x70	db2c0
<i>DB2 Query Patroller Administration Guide</i>	Provides an operational overview of the DB2 Query Patroller system, specific operational and administrative information, and task information for the administrative graphical user interface utilities.	SC09-2958 db2dwx70	db2dw
<i>DB2 Query Patroller User's Guide</i>	Describes how to use the tools and functions of the DB2 Query Patroller.	SC09-2960 db2wwx70	db2ww
<i>Glossary</i>	Provides definitions for terms used in DB2 and its components. Available in HTML format and in the <i>SQL Reference</i> .	No form number db2t0x70	db2t0
<i>Image, Audio, and Video Extenders Administration and Programming</i>	Provides general information about DB2 extenders, and information on the administration and configuration of the image, audio, and video (IAV) extenders and on programming using the IAV extenders. It includes reference information, diagnostic information (with messages), and samples.	SC26-9929 dmbu7x70	dmbu7
<i>Information Catalog Manager Administration Guide</i>	Provides guidance on managing information catalogs.	SC26-9995 db2dix70	db2di

Table 20. DB2 Information (continued)

Name	Description	Form Number PDF File Name	HTML Directory
<i>Information Catalog Manager Programming Guide and Reference</i>	Provides definitions for the architected interfaces for the Information Catalog Manager.	SC26-9997 db2bix70	db2bi
<i>Information Catalog Manager User's Guide</i>	Provides information on using the Information Catalog Manager user interface.	SC26-9996 db2aix70	db2ai
<i>Installation and Configuration Supplement</i>	Guides you through the planning, installation, and setup of platform-specific DB2 clients. This supplement also contains information on binding, setting up client and server communications, DB2 GUI tools, DRDA AS, distributed installation, the configuration of distributed requests, and accessing heterogeneous data sources.	GC09-2957 db2iyx70	db2iy
<i>Message Reference</i>	Lists messages and codes issued by DB2, the Information Catalog Manager, and the Data Warehouse Center, and describes the actions you should take. You can order both volumes of the Message Reference in the English language in North America with the form number SBOF-8932.	Volume 1 GC09-2978 db2m1x70 Volume 2 GC09-2979 db2m2x70	db2m0
<i>OLAP Integration Server Administration Guide</i>	Explains how to use the Administration Manager component of the OLAP Integration Server.	SC27-0787 db2dpx70	n/a
<i>OLAP Integration Server Metaoutline User's Guide</i>	Explains how to create and populate OLAP metaoutlines using the standard OLAP Metaoutline interface (not by using the Metaoutline Assistant).	SC27-0784 db2upx70	n/a
<i>OLAP Integration Server Model User's Guide</i>	Explains how to create OLAP models using the standard OLAP Model Interface (not by using the Model Assistant).	SC27-0783 db2lpx70	n/a
<i>OLAP Setup and User's Guide</i>	Provides configuration and setup information for the OLAP Starter Kit.	SC27-0702 db2ipx70	db2ip
<i>OLAP Spreadsheet Add-in User's Guide for Excel</i>	Describes how to use the Excel spreadsheet program to analyze OLAP data.	SC27-0786 db2epx70	db2ep

Table 20. DB2 Information (continued)

Name	Description	Form Number PDF File Name	HTML Directory
<i>OLAP Spreadsheet Add-in User's Guide for Lotus 1-2-3</i>	Describes how to use the Lotus 1-2-3 spreadsheet program to analyze OLAP data.	SC27-0785 db2tpx70	db2tp
<i>Replication Guide and Reference</i>	Provides planning, configuration, administration, and usage information for the IBM Replication tools supplied with DB2.	SC26-9920 db2e0x70	db2e0
<i>Spatial Extender User's Guide and Reference</i>	Provides information about installing, configuring, administering, programming, and troubleshooting the Spatial Extender. Also provides significant descriptions of spatial data concepts and provides reference information (messages and SQL) specific to the Spatial Extender.	SC27-0701 db2sbx70	db2sb
<i>SQL Getting Started</i>	Introduces SQL concepts and provides examples for many constructs and tasks.	SC09-2973 db2y0x70	db2y0
<i>SQL Reference, Volume 1 and Volume 2</i>	Describes SQL syntax, semantics, and the rules of the language. This book also includes information about release-to-release incompatibilities, product limits, and catalog views. You can order both volumes of the <i>SQL Reference</i> in the English language in North America with the form number SBOF-8933.	Volume 1 SC09-2974 db2s1x70 Volume 2 SC09-2975 db2s2x70	db2s0
<i>System Monitor Guide and Reference</i>	Describes how to collect different kinds of information about databases and the database manager. This book explains how to use the information to understand database activity, improve performance, and determine the cause of problems.	SC09-2956 db2f0x70	db2f0
<i>Text Extender Administration and Programming</i>	Provides general information about DB2 extenders and information on the administration and configuring of the text extender and on programming using the text extenders. It includes reference information, diagnostic information (with messages) and samples.	SC26-9930 desu9x70	desu9

Table 20. DB2 Information (continued)

Name	Description	Form Number PDF File Name	HTML Directory
<i>Troubleshooting Guide</i>	Helps you determine the source of errors, recover from problems, and use diagnostic tools in consultation with DB2 Customer Service.	GC09-2850 db2p0x70	db2p0
<i>What's New</i>	Describes the new features, functions, and enhancements in DB2 Universal Database, Version 7.	SC09-2976 db2q0x70	db2q0
DB2 Installation and Configuration Information			
<i>DB2 Connect Enterprise Edition for OS/2 and Windows Quick Beginnings</i>	Provides planning, migration, installation, and configuration information for DB2 Connect Enterprise Edition on the OS/2 and Windows 32-bit operating systems. This book also contains installation and setup information for many supported clients.	GC09-2953 db2c6x70	db2c6
<i>DB2 Connect Enterprise Edition for UNIX Quick Beginnings</i>	Provides planning, migration, installation, configuration, and task information for DB2 Connect Enterprise Edition on UNIX-based platforms. This book also contains installation and setup information for many supported clients.	GC09-2952 db2cyx70	db2cy
<i>DB2 Connect Personal Edition Quick Beginnings</i>	Provides planning, migration, installation, configuration, and task information for DB2 Connect Personal Edition on the OS/2 and Windows 32-bit operating systems. This book also contains installation and setup information for all supported clients.	GC09-2967 db2c1x70	db2c1
<i>DB2 Connect Personal Edition Quick Beginnings for Linux</i>	Provides planning, installation, migration, and configuration information for DB2 Connect Personal Edition on all supported Linux distributions.	GC09-2962 db2c4x70	db2c4
<i>DB2 Data Links Manager Quick Beginnings</i>	Provides planning, installation, configuration, and task information for DB2 Data Links Manager for AIX and Windows 32-bit operating systems.	GC09-2966 db2z6x70	db2z6

Table 20. DB2 Information (continued)

Name	Description	Form Number PDF File Name	HTML Directory
<i>DB2 Enterprise - Extended Edition for UNIX Quick Beginnings</i>	Provides planning, installation, and configuration information for DB2 Enterprise - Extended Edition on UNIX-based platforms. This book also contains installation and setup information for many supported clients.	GC09-2964 db2v3x70	db2v3
<i>DB2 Enterprise - Extended Edition for Windows Quick Beginnings</i>	Provides planning, installation, and configuration information for DB2 Enterprise - Extended Edition for Windows 32-bit operating systems. This book also contains installation and setup information for many supported clients.	GC09-2963 db2v6x70	db2v6
<i>DB2 for OS/2 Quick Beginnings</i>	Provides planning, installation, migration, and configuration information for DB2 Universal Database on the OS/2 operating system. This book also contains installation and setup information for many supported clients.	GC09-2968 db2i2x70	db2i2
<i>DB2 for UNIX Quick Beginnings</i>	Provides planning, installation, migration, and configuration information for DB2 Universal Database on UNIX-based platforms. This book also contains installation and setup information for many supported clients.	GC09-2970 db2ixx70	db2ix
<i>DB2 for Windows Quick Beginnings</i>	Provides planning, installation, migration, and configuration information for DB2 Universal Database on Windows 32-bit operating systems. This book also contains installation and setup information for many supported clients.	GC09-2971 db2i6x70	db2i6
<i>DB2 Personal Edition Quick Beginnings</i>	Provides planning, installation, migration, and configuration information for DB2 Universal Database Personal Edition on the OS/2 and Windows 32-bit operating systems.	GC09-2969 db2i1x70	db2i1
<i>DB2 Personal Edition Quick Beginnings for Linux</i>	Provides planning, installation, migration, and configuration information for DB2 Universal Database Personal Edition on all supported Linux distributions.	GC09-2972 db2i4x70	db2i4

Table 20. DB2 Information (continued)

Name	Description	Form Number PDF File Name	HTML Directory
<i>DB2 Query Patroller Installation Guide</i>	Provides installation information about DB2 Query Patroller.	GC09-2959 db2iwx70	db2iw
<i>DB2 Warehouse Manager Installation Guide</i>	Provides installation information for warehouse agents, warehouse transformers, and the Information Catalog Manager.	GC26-9998 db2idx70	db2id
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. The sample programs are provided for informational purposes only. Not all samples are available in all programming languages. The HTML samples are only available when the DB2 Application Development Client is installed. For more information on the programs, refer to the <i>Application Building Guide</i> .	No form number	db2hs
Release Notes			
<i>DB2 Connect Release Notes</i>	Provides late-breaking information that could not be included in the DB2 Connect books.	See note #2.	db2cr
<i>DB2 Installation Notes</i>	Provides late-breaking installation-specific information that could not be included in the DB2 books.	Available on product CD-ROM only.	
<i>DB2 Release Notes</i>	Provides late-breaking information about all DB2 products and features that could not be included in the DB2 books.	See note #2.	db2ir

Notes:

1. The character *x* in the sixth position of the file name indicates the language version of a book. For example, the file name *db2d0e70* identifies the English version of the *Administration Guide* and the file name *db2d0f70* identifies the French version of the same book. The following letters are used in the sixth position of the file name to indicate the language version:

Language	Identifier
Brazilian Portuguese	b

Bulgarian	u
Czech	x
Danish	d
Dutch	q
English	e
Finnish	y
French	f
German	g
Greek	a
Hungarian	h
Italian	i
Japanese	j
Korean	k
Norwegian	n
Polish	p
Portuguese	v
Russian	r
Simp. Chinese	c
Slovenian	l
Spanish	z
Swedish	s
Trad. Chinese	t
Turkish	m

2. Late breaking information that could not be included in the DB2 books is available in the Release Notes in HTML format and as an ASCII file. The HTML version is available from the Information Center and on the product CD-ROMs. To view the ASCII file:
 - On UNIX-based platforms, see the `Release.Notes` file. This file is located in the `DB2DIR/Readme/%L` directory, where `%L` represents the locale name and `DB2DIR` represents:
 - `/usr/lpp/db2_07_01` on AIX
 - `/opt/IBMDB2/V7.1` on HP-UX, PTX, Solaris, and Silicon Graphics IRIX
 - `/usr/IBMDB2/V7.1` on Linux.
 - On other platforms, see the `RELEASE.TXT` file. This file is located in the directory where the product is installed. On OS/2 platforms, you can also double-click the **IBM DB2** folder and then double-click the **Release Notes** icon.

Printing the PDF Books

If you prefer to have printed copies of the books, you can print the PDF files found on the DB2 publications CD-ROM. Using the Adobe Acrobat Reader, you can print either the entire book or a specific range of pages. For the file name of each book in the library, see Table 20 on page 238.

You can obtain the latest version of the Adobe Acrobat Reader from the Adobe Web site at <http://www.adobe.com>.

The PDF files are included on the DB2 publications CD-ROM with a file extension of PDF. To access the PDF files:

1. Insert the DB2 publications CD-ROM. On UNIX-based platforms, mount the DB2 publications CD-ROM. Refer to your *Quick Beginnings* book for the mounting procedures.
2. Start the Acrobat Reader.
3. Open the desired PDF file from one of the following locations:
 - On OS/2 and Windows platforms:
x:\doc\language directory, where *x* represents the CD-ROM drive and *language* represent the two-character country code that represents your language (for example, EN for English).
 - On UNIX-based platforms:
/cdrom/doc/%L directory on the CD-ROM, where */cdrom* represents the mount point of the CD-ROM and *%L* represents the name of the desired locale.

You can also copy the PDF files from the CD-ROM to a local or network drive and read them from there.

Ordering the Printed Books

You can order the printed DB2 books either individually or as a set (in North America only) by using a sold bill of forms (SBOF) number. To order books, contact your IBM authorized dealer or marketing representative, or phone 1-800-879-2755 in the United States or 1-800-IBM-4Y0U in Canada. You can also order the books from the Publications Web page at <http://www.elink.ibm.com/pbl/pbl>.

Two sets of books are available. SBOF-8935 provides reference and usage information for the DB2 Warehouse Manager. SBOF-8931 provides reference and usage information for all other DB2 Universal Database products and features. The contents of each SBOF are listed in the following table:

Table 21. Ordering the printed books

SBOF Number	Books Included
SBOF-8931	<ul style="list-style-type: none"> • Administration Guide: Planning • Administration Guide: Implementation • Administration Guide: Performance • Administrative API Reference • Application Building Guide • Application Development Guide • CLI Guide and Reference • Command Reference • Data Movement Utilities Guide and Reference • Data Warehouse Center Administration Guide • Data Warehouse Center Application Integration Guide • DB2 Connect User's Guide • Installation and Configuration Supplement • Image, Audio, and Video Extenders Administration and Programming • Message Reference, Volumes 1 and 2 • OLAP Integration Server Administration Guide • OLAP Integration Server Metaoutline User's Guide • OLAP Integration Server Model User's Guide • OLAP Integration Server User's Guide • OLAP Setup and User's Guide • OLAP Spreadsheet Add-in User's Guide for Excel • OLAP Spreadsheet Add-in User's Guide for Lotus 1-2-3 • Replication Guide and Reference • Spatial Extender Administration and Programming Guide • SQL Getting Started • SQL Reference, Volumes 1 and 2 • System Monitor Guide and Reference • Text Extender Administration and Programming • Troubleshooting Guide • What's New
SBOF-8935	<ul style="list-style-type: none"> • Information Catalog Manager Administration Guide • Information Catalog Manager User's Guide • Information Catalog Manager Programming Guide and Reference • Query Patroller Administration Guide • Query Patroller User's Guide

DB2 Online Documentation

Accessing Online Help

Online help is available with all DB2 components. The following table describes the various types of help.

Type of Help	Contents	How to Access...
<i>Command Help</i>	Explains the syntax of commands in the command line processor.	<p>From the command line processor in interactive mode, enter:</p> <p style="padding-left: 40px;"><i>? command</i></p> <p>where <i>command</i> represents a keyword or the entire command.</p> <p>For example, <i>? catalog</i> displays help for all the CATALOG commands, while <i>? catalog database</i> displays help for the CATALOG DATABASE command.</p>
<i>Client Configuration Assistant Help</i>	Explains the tasks you can perform in a window or notebook. The help includes overview and prerequisite information you need to know, and it describes how to use the window or notebook controls.	From a window or notebook, click the Help push button or press the F1 key.
<i>Command Center Help</i>		
<i>Control Center Help</i>		
<i>Data Warehouse Center Help</i>		
<i>Event Analyzer Help</i>		
<i>Information Catalog Manager Help</i>		
<i>Satellite Administration Center Help</i>		
<i>Script Center Help</i>		
<i>Message Help</i>	Describes the cause of a message and any action you should take.	<p>From the command line processor in interactive mode, enter:</p> <p style="padding-left: 40px;"><i>? XXXnnnnn</i></p> <p>where <i>XXXnnnnn</i> represents a valid message identifier.</p> <p>For example, <i>? SQL30081</i> displays help about the SQL30081 message.</p> <p>To view message help one screen at a time, enter:</p> <p style="padding-left: 40px;"><i>? XXXnnnnn more</i></p> <p>To save message help in a file, enter:</p> <p style="padding-left: 40px;"><i>? XXXnnnnn > filename.ext</i></p> <p>where <i>filename.ext</i> represents the file where you want to save the message help.</p>

Type of Help	Contents	How to Access...
SQL Help	Explains the syntax of SQL statements.	<p>From the command line processor in interactive mode, enter:</p> <pre>help <i>statement</i></pre> <p>where <i>statement</i> represents an SQL statement.</p> <p>For example, help SELECT displays help about the SELECT statement.</p> <p>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> <pre>? <i>sqlstate</i> or ? <i>class code</i></pre> <p>where <i>sqlstate</i> represents a valid five-digit SQL state and <i>class code</i> represents the first two digits of the SQL state.</p> <p>For example, ? 08003 displays help for the 08003 SQL state, while ? 08 displays help for the 08 class code.</p>

Viewing Information Online

The books included with this product are in Hypertext Markup Language (HTML) softcopy format. Softcopy format enables you to search or browse the information and provides hypertext links to related information. It also makes it easier to share the library across your site.

You can view the online books or sample programs with any browser that conforms to HTML Version 3.2 specifications.

To view online books or sample programs:

- If you are running DB2 administration tools, use the Information Center.
- From a browser, click **File** —> **Open Page**. The page you open contains descriptions of and links to DB2 information:
 - On UNIX-based platforms, open the following page:

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

where %L represents the locale name.

- On other platforms, open the following page:

```
sql1lib\doc\html\index.htm
```

The path is located on the drive where DB2 is installed.

If you have not installed the Information Center, you can open the page by double-clicking the **DB2 Information** icon. Depending on the system you are using, the icon is in the main product folder or the Windows Start menu.

Installing the Netscape Browser

If you do not already have a Web browser installed, you can install Netscape from the Netscape CD-ROM found in the product boxes. For detailed instructions on how to install it, perform the following:

1. Insert the Netscape CD-ROM.
2. On UNIX-based platforms only, mount the CD-ROM. Refer to your *Quick Beginnings* book for the mounting procedures.
3. For installation instructions, refer to the `CDNAVnn.txt` file, where *nn* represents your two character language identifier. The file is located at the root directory of the CD-ROM.

Accessing Information with the Information Center

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

You can open the Information Center by double-clicking the Information Center icon. Depending on the system you are using, the icon is in the Information folder in the main product folder or the Windows **Start** menu.

You can also access the Information Center by using the toolbar and the **Help** menu on the DB2 Windows platform.

The Information Center provides six types of information. Click the appropriate tab to look at the topics provided for that type.

Tasks	Key tasks you can perform using DB2.
Reference	DB2 reference information, such as keywords, commands, and APIs.
Books	DB2 books.
Troubleshooting	Categories of error messages and their recovery actions.
Sample Programs	Sample programs that come with the DB2 Application Development Client. If you did not install the DB2 Application Development Client, this tab is not displayed.
Web	DB2 information on the World Wide Web. To access this information, you must have a connection to the Web from your system.

When you select an item in one of the lists, the Information Center launches a viewer to display the information. The viewer might be the system help viewer, an editor, or a Web browser, depending on the kind of information you select.

The Information Center provides a find feature, so you can look for a specific topic without browsing the lists.

For a full text search, follow the hypertext link in the Information Center to the **Search DB2 Online Information** search form.

The HTML search server is usually started automatically. If a search in the HTML information does not work, you may have to start the search server using one of the following methods:

On Windows

Click **Start** and select **Programs** → **IBM DB2** → **Information** → **Start HTML Search Server**.

On OS/2

Double-click the **DB2 for OS/2** folder, and then double-click the **Start HTML Search Server** icon.

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

Note: The Search function is not available in the Linux, PTX, and Silicon Graphics IRIX environments.

Using DB2 Wizards

Wizards help you complete specific administration tasks by taking you through each task one step at a time. Wizards are available through the Control Center and the Client Configuration Assistant. The following table lists the wizards and describes their purpose.

Note: The Create Database, Create Index, Configure Multisite Update, and Performance Configuration wizards are available for the partitioned database environment.

Wizard	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>Backup Database</i>	Determine, create, and schedule a backup plan.	From the Control Center, right-click the database you want to back up and select Backup → Database Using Wizard .

Wizard	Helps You to...	How to Access...
<i>Configure Multisite Update</i>	Configure a multisite update, a distributed transaction, or a two-phase commit.	From the Control Center, right-click the Databases folder and select Multisite Update .
<i>Create Database</i>	Create a database, and perform some basic configuration tasks.	From the Control Center, right-click the Databases folder and select Create → Database Using Wizard .
<i>Create Table</i>	Select basic data types, and create a primary key for the table.	From the Control Center, right-click the Tables icon and select Create → Table Using Wizard .
<i>Create Table Space</i>	Create a new table space.	From the Control Center, right-click the Table Spaces icon and select Create → Table Space Using Wizard .
<i>Create Index</i>	Advise which indexes to create and drop for all your queries.	From the Control Center, right-click the Index icon and select Create → Index Using Wizard .
<i>Performance Configuration</i>	Tune the performance of a database by updating configuration parameters to match your business requirements.	From the Control Center, right-click the database you want to tune and select Configure Performance Using Wizard . For the partitioned database environment, from the Database Partitions view, right-click the first database partition you want to tune and select Configure Performance Using Wizard .
<i>Restore Database</i>	Recover a database after a failure. It helps you understand which backup to use, and which logs to replay.	From the Control Center, right-click the database you want to restore and select Restore → Database Using Wizard .

Setting Up a Document Server

By default, the DB2 information is installed on your local system. This means that each person who needs access to the DB2 information must install the same files. To have the DB2 information stored in a single location, perform the following steps:

1. Copy all files and subdirectories from `\sql11ib\doc\html` on your local system to a Web server. Each book has its own subdirectory that contains all the necessary HTML and GIF files that make up the book. Ensure that the directory structure remains the same.

2. Configure the Web server to look for the files in the new location. For information, refer to the NetQuestion Appendix in the *Installation and Configuration Supplement*.
3. If you are using the Java version of the Information Center, you can specify a base URL for all HTML files. You should use the URL for the list of books.
4. When you are able to view the book files, you can bookmark commonly viewed topics. You will probably want to bookmark the following pages:
 - List of books
 - Tables of contents of frequently used books
 - Frequently referenced articles, such as the ALTER TABLE topic
 - The Search form

For information about how you can serve the DB2 Universal Database online documentation files from a central machine, refer to the NetQuestion Appendix in the *Installation and Configuration Supplement*.

Searching Information Online

To find information in the HTML files, use one of the following methods:

- Click **Search** in the top frame. Use the search form to find a specific topic. This function is not available in the Linux, PTX, or Silicon Graphics IRIX environments.
- Click **Index** in the top frame. Use the index to find a specific topic in the book.
- Display the table of contents or index of the help or the HTML book, and then use the find function of the Web browser to find a specific topic in the book.
- Use the bookmark function of the Web browser to quickly return to a specific topic.
- Use the search function of the Information Center to find specific topics. See “Accessing Information with the Information Center” on page 251 for details.

Appendix F. Removing DB2 Products

This chapter describes how to remove DB2 products on UNIX-based systems.

For information on removing DB2 products on OS/2 and Windows systems, refer to the *Quick Beginnings* manuals for these platforms.

Stop the Administration Server

You must stop the Administration Server before you remove DB2 products.

To stop the Administration Server:

1. Log in as the Administration Server owner.

2. Run the start up script:

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

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

3. Stop the Administration Server by entering the **db2admin stop** command.

Stop all DB2 instances

You must stop all DB2 instances before you remove DB2.

To stop a DB2 instance:

1. Log in as a user with root authority.

2. To obtain a list of the names of all DB2 instances on your system, enter the **DB2DIR/bin/db2ilist** command

```
where DB2DIR      = /usr/lpp/db2_07_01  on AIX
                  = /opt/IBMd2/V7.1  on HP-UX, PTX, or Solaris
```

3. Log out.

4. Stop the instance.

To stop a database instance:

a. Log in as the instance owner.

b. Run the start up script:

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

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

- c. Back up files in the `INSTHOME/sqllib` directory, if needed, where `INSTHOME` is the home directory of the instance owner.
You might want to save the database manager configuration file, `db2system`, the `db2nodes.cfg` file, or user defined function or fenced stored procedure applications in `INSTHOME/sqllib/function`.
 - d. Stop all database applications by entering the **db2 force application all** command.
 - e. Stop the DB2 database manager by entering the **db2stop** command.
 - f. Confirm that the instance is stopped by entering the **db2 terminate** command.
5. Repeat these steps for each instance.

Remove the Administration Server

You must remove the Administration Server before you remove DB2.

To remove the Administration Server:

1. Log in as the Administration Server owner.
2. Run the start up script:

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

where *INSTHOME* is the home directory of the user who owns the Administration Server.

3. Back up the files in the `ASHOME/sqllib` directory, if needed, where `ASHOME` is the home directory of the user who owns the Administration Server.
4. Log off.
5. Log in as *root* and remove the Administration Server by entering the following command:

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

where

```
where DB2DIR    = /usr/lpp/db2_07_01  on AIX
          = /opt/IBMd2/V7.1      on HP-UX, PTX, or Solaris
```

and *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.

Remove DB2 instances (Optional)

You can remove some or all of the DB2 Version 7 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 7 products, or if you do not want to migrate existing instances to a later version of DB2. **On UNIX Systems:**

To remove an instance:

1. Log in as a user with root authority.
2. Remove the instance by entering the following command:

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

where

```
where DB2DIR    = /usr/lpp/db2_07_01  on AIX
          = /opt/IBMcdb2/V7.1  on HP-UX, PTX, or Solaris
```

The **db2idrop** command removes the instance entry from the list of instances and removes the `INSTHOME/sql11ib` directory, where `INSTHOME` is the home directory of the instance and where `InstName` is the login name of the instance.

3. 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.

Remove DB2 Products

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

To remove DB2 Products on AIX Systems

You can remove DB2 on AIX using the System Management Interface Tool (SMIT) or the **installp** command.

To remove DB2 products on AIX systems using SMIT:

1. Log in as user with root authority.
2. Type **smit install_remove** to proceed directly to the Remove Software Products screen.
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_07_01`.
4. Press Enter to start removing DB2 products.

To remove *all* DB2 Version 7 products, enter the **installp -u db2_07_01** command.

To remove DB2 Products on HP-UX Systems

1. Log in as user with root authority.
2. Use **swremove** to remove some or all DB2 Version 7 products.

To remove DB2 Products on Linux, PTX, SGI IRIX, and Solaris Systems

1. Log in as user with root authority.
2. Enter the **db2_deinstall -n** command. This command is found in the root directory of your DB2 Version 7 product CD.

Note: The **db2_deinstall -n** command will remove *all* DB2 products from your system.

Appendix G. Notices

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

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

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

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

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

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

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

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

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

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

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

IBM Canada Limited
Office of the Lab Director
1150 Eglinton Ave. East
North York, Ontario
M3C 1H7
CANADA

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

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

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

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

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

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

COPYRIGHT LICENSE:

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

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

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

Trademarks

The following terms, which may be denoted by an asterisk(*), are trademarks of International Business Machines Corporation in the United States, other countries, or both.

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

The following terms are trademarks or registered trademarks of other companies:

Microsoft, Windows, and Windows NT are trademarks or registered trademarks of Microsoft Corporation.

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

Tivoli and NetView are trademarks of Tivoli Systems Inc. in the United States, other countries, or both.

UNIX is a registered trademark in the United States, other countries or both and is licensed exclusively through X/Open Company Limited.

Other company, product, or service names, which may be denoted by a double asterisk(**) may be trademarks or service marks of others.

Index

A

- access profiles
 - adding databases 173
 - client 178
 - creating 178
 - server 178
 - using 178
- accessing data
 - using DB2 Connect 222
 - using Net.Data or JDBC 225
- accessing DB2 servers
 - TCP/IP 183
- accessing multiple servers 183, 220
- accessing servers
 - overview 183
- add database wizard 252, 253
- adding databases
 - manually 176
 - using access profiles 173
 - using Discovery 174
- Administration Server
 - overview 234
- AIX
 - mounting CD-ROM 200
- Alert Center 231
- ALTER TABLESPACE 128
- APPC
 - Communications Manager for OS/2 23
 - Communications Server for OS/2 23
 - SNA Server 21
 - software requirements 21, 23, 25
 - SunLink SNA 21, 24
 - supported platforms 21
- application development
 - using Net.Data or JDBC 225
- application programs
 - coordinator node 4
 - migrating to DB2 Enterprise - Extended Edition 4

B

- backup database wizard 252
- books 237, 247

C

- capacity
 - disk storage 7

- catalog node name
 - naming rules 207
- catalog nodes 9
- cataloging
 - databases 188
 - TCP/IP node 187, 188
- client profiles
 - creating 179
 - definition 179
 - importing 180
 - using 179
- clients
 - configuring 183
 - installing 137
 - operating systems supported 219
- columns
 - long-field, partitioning key considerations 10
- Command Center
 - entering DB2 commands 196
 - entering SQL statements 196
 - overview 228
- commands
 - db2 list applications 28
 - db2 list tablespaces 128
 - db2 terminate 28
 - db2cc 160
 - db2imigr 125, 126
 - db2jstrt 158
 - db2sampl 160
 - db2setup 149, 153
 - db2start 129
 - db2stop 28
 - db2uiddl 129
 - rlogin 153
 - sniffle 164
- communication protocols
 - APPC 21, 22, 23, 24, 25
 - configuring 183
 - IPX/SPX 21, 22, 23, 24, 25
 - Named Pipes 23, 25
 - NetBIOS 21, 23, 25
 - TCP/IP 21, 22, 23, 24, 25, 183
- communications
 - configuring the client 183
 - Control Center 232
 - fast communication manager 12
 - managing 183

- communications (*continued*)
 - TCP/IP 183
- concurrency
 - improving with logical nodes 10
- configuration parameters
 - setting DB2 199
 - SYSADM_GROUP 199
- configure multisite update wizard 252
- configuring
 - DB2 clients
 - using the Client Configuration Assistant (CCA) 171
 - TCP/IP 183
 - configuring client communications
 - setting configuration parameters 130, 183
 - using the Command Line Processor 183
 - configuring communications
 - overview 183
- Control Center
 - administering DB2 Connect Enterprise Edition 165
 - administering DB2 for OS/390 165
 - as a Java applet 155
 - as a Java application 155
 - components 228
 - configuring to work with a web server 162
 - customizing db2cc.htm 161
 - functional considerations 162
 - JDBC Applet Server 158
 - listener daemon
 - requirements 228
 - machine configurations 156
 - overview 228
 - running as an applet 160
 - running as an application 160
 - setting up to run as an applet 158
 - supported browsers 157
 - supported Java Runtime Environments (JRE) 157
 - troubleshooting information 164
 - UNIX installation tips 162
- coordinator nodes
 - description 4

- cost-based query optimizers
 - description 4
- create database wizard 253
- create table space wizard 253
- create table wizard 253
- creating profiles
 - clients 179
 - server 179
- creating the sample database
 - cataloging a database 128, 188
 - cataloging a node 187
 - connecting to a database 190

D

- Data Links Manager
 - overview 219
- database access
 - two-phase commit 221
- database administration tools
 - Control Center 228
 - overview 228
- database alias
 - definition 189
 - naming rules 207
- database manager
 - definition 3
- database objects
 - naming rules 208
- database partition server
 - definition 3
- databases
 - definition 3
 - migrating 127
 - naming rules 207
 - partition 3
- DB2 Application Development Client
 - overview 235
- DB2 clients
 - accessing databases 220
 - changing privileges 199
 - installing 137
 - installing on UNIX workstations 153
 - licensing 137
 - OS/2 145
 - overview 220, 233
 - platforms supported 219
 - pre-version 7 137
 - software requirements 21
 - supported platforms 137
 - updating kernel parameters on
 - HP-UX, NUMA-Q/PTX and Solaris 150
 - WIN-OS/2 support 145

- DB2 clients (*continued*)
 - Windows 32-bit operating systems 141
- DB2 Connect
 - overview 216, 222
- DB2 Enterprise - Extended Edition
 - configuration planning 18
 - hardware configuration 6
- DB2 Enterprise Edition
 - memory requirements 18
- DB2 Everywhere
 - overview 213
- DB2 library
 - books 237
 - Information Center 251
 - language identifier for books 245
 - late-breaking information 246
 - online help 248
 - ordering printed books 247
 - printing PDF books 246
 - searching online information 254
 - setting up document server 253
 - structure of 237
 - viewing online information 250
 - wizards 252
- DB2 Security Server
 - starting on Windows NT or Windows 2000 159
- DB2 Universal Database
 - Control Center 228
 - DB2 Performance Monitor 232
 - DB2 Snapshot Monitor 225
 - overview 228
 - platforms supported 228
 - software requirements 19
 - Visual Explain 232
- db2cclst listener daemon 228
- db2classes.exe 160
- db2classes.tar.Z 160
- db2cshrc 38, 57, 76, 93, 112
- db2fcmdm daemon
 - description 12
- db2profile 38, 57, 76, 93, 112
- db2setup
 - using to install DB2 clients 149
- DB2SYSTEM
 - naming rules 211
- db2uidl command 129
- Developer Editions
 - overviews 215
- developing applications
 - using Net.Data or JDBC 225

- Discovery
 - adding databases 174
- disk requirements
 - client 18
 - server 18
- Distributed Computing Environment
 - software requirements 21, 22, 23
- DMS tablespaces
 - migrating Version 5 databases 128

E

- Enterprise - Extended Edition
 - capacity 7
 - overview 215
 - security for an instance 11
 - shared nothing definition 3
 - storage 7
- Enterprise Edition
 - overview 215
- environment variables 12
- Export function 178, 179

F

- files
 - listing for an instance 12
- filesets
 - db2fcmdm daemon 12
 - description 12
- fixed disks
 - hardware requirements 18

G

- gigabyte (GB) 7

H

- hardware requirements
 - fixed disk 18
- host databases
 - overview 216
- HP-UX
 - mounting CD-ROM 201
- HP-UX client
 - updating kernel components 150
- HTML
 - sample programs 245

I

- IBMCATGROUP nodegroup 9
- IBMDEFAULTGROUP nodegroup 9
- IBMTEMPGROUP nodegroup 8
- Import function 178
- importing profiles
 - client 180
- index wizard 253
- Information Center 251

- installing
 - AIX 31, 53, 54, 105, 107
 - CID using SystemView
 - LAN 143, 146
 - client 17, 18
 - DB2 Application Development Clients 137
 - OS/2 145
 - Windows 32-bit operating systems 141
 - DB2 clients 137
 - OS/2 145
 - DB2 clients on UNIX workstations 153
 - errors 143, 146
 - Linux 71, 73
 - log 143, 146
 - Netscape browser 251
 - OS/2 client 145
 - remote clients 153
 - server 17, 18
 - Solaris 106
 - updating kernel components 57, 110, 151
- instance owner 12
- instances
 - definition 11
 - listing files in 12
 - multiple on one processor 12
 - naming restrictions 209
- Intelligent Miner
 - overview 218
- inter-operator parallelism 5
- IPX/SPX
 - software requirements 23, 24
- J**
 - Java Runtime Environment (JRE)
 - defined 155
 - Java Support 225
 - Java Virtual Machine (JVM) 155
 - JDBC Applet Server 158
 - Journal 231
 - JRE
 - supported levels for Control Center 157
- K**
 - kernel configuration parameters
 - updating on UNIX clients 150
- L**
 - LANG environment variable 205
 - language identifier
 - books 245
 - late-breaking information 246

- Linux
 - mounting CD-ROM 201
- listing
 - files in an instance 12
- logical nodes
 - improving concurrency 10
 - purpose 10
- long-field columns
 - partitioning key
 - considerations 10
- M**
 - managing connections 183
 - overview 183, 233
 - using the Client Configuration Assistant 233
 - using the Command Line Processor 183
 - managing databases using the Control Center 228
 - managing server communications
 - overview 232
 - manually adding databases 176
 - memory requirements
 - client 17
 - estimating 17
 - recommended 17
 - server 17
 - Microsoft SNA Server
 - version required 25
 - migrating
 - applications 4
 - databases 28, 127
 - instances 125
 - pre-installation tasks 27
 - rebinding packages 130
 - updating database and database manager configuration 130
 - updating statistics 129
 - Version 5 databases with DMS tablespaces 128
 - migration
 - optional post-migration tasks 128
 - post-installation tasks 125
 - mounting CD-ROM
 - AIX 200
 - HP-UX 201
 - Linux 201
 - PTX 201
 - Solaris 202
- N**
 - naming rules
 - database alias 207

- naming rules (*continued*)
 - database objects 208
 - databases 207
 - general 207
 - groups 209
 - instance names 209
 - password 211
 - userid 209
 - username 209
- Net.Data
 - connecting to the Internet 225
 - overview 225
- Net Search Extender
 - overview 218
- NetBIOS
 - code set 205
 - codepage support 205
 - determining codepage 207
 - on the client 23
- Netfinity Server 23
- Netscape browser
 - installing 251
- nodegroups
 - IBMCATGROUP 9
 - IBMDEFAULTGROUP 9
 - multipartition nodegroup 8
- nodes
 - definition 3
 - FCM daemon 12
 - in the same instance 11
- NUMA-Q/PTX client
 - updating kernel components 151
- O**
 - ODBC
 - running applications on OS/2 147
 - OLAP Server
 - overview 218
 - online help 248
 - online information
 - searching 254
 - viewing 250
- P**
 - parameters
 - SYSADM_GROUP 199
 - partitioning key
 - definition 10
 - partitioning map
 - overview 10
 - partitions
 - definition 3

- passwords
 - naming rules 211
 - PDF 246
 - performance configuration
 - wizard 253
 - Performance Monitor
 - using 232
 - Personal Edition
 - overview 214
 - planning
 - DB2 configuration 17
 - DB2 Connect configuration 17
 - pools
 - agent 10
 - printing PDF books 246
 - privileges
 - required 199
 - product
 - descriptions 213
 - overview 213
 - profiles
 - client 178, 179
 - export 178
 - server 178
 - protocols
 - TCP/IP 183
 - PTX
 - mounting CD-ROM 201
- Q**
- query
 - cost-based optimizer 4
- R**
- Relational Connect
 - overview 217
 - release notes 246
 - restore wizard 253
 - restrictions
 - instance name 209
- S**
- sample programs
 - cross-platform 245
 - HTML 245
 - Satellite Edition
 - overview 214
 - scaling the database system
 - definition 10
 - Script Center 230
 - searching
 - online information 252, 254
 - security
 - for an instance 11
 - server profiles
 - creating 179
 - server profiles (*continued*)
 - definition 178
 - setting configuration
 - parameters 183
 - setting up client communications
 - using the Command Line Processor 183
 - setting up document server 253
 - shared-nothing configuration
 - definition 3
 - storage 7
 - SmartGuides
 - wizards 252
 - software requirements
 - communication protocols 19
 - DB2 Application Development Clients 19, 21, 22, 23, 24, 25
 - DB2 clients 19, 21, 22, 23, 24, 25
 - DB2 Connect 19
 - DB2 Universal Database 19
 - Net.Data 21, 22, 23, 24, 25
 - Solaris
 - mounting CD-ROM 202
 - Solaris client
 - updating kernel components 152
 - Spatial Extender
 - overview 218
 - SQL
 - viewing using Visual Explain 232
 - Stored Procedure Builder 231
 - symmetric multiprocessor architecture 10
 - SYSADM
 - controlling 199
 - SYSADM_GROUP parameter 199
 - system configuration
 - with DB2 Connect 223
 - with DB2 Universal Database 220
- T**
- tables
 - maximum size 7
 - multinode nodegroups 10
 - single-node nodegroups 10
 - TCP/IP
 - client 183
 - configuring 183
 - configuring on OS/2 163
 - enabling localhost on OS/2 163
 - enabling loopback on OS/2 163
 - preventing socket collisions 184
 - setting up client-to-server 183
 - TCP/IP (*continued*)
 - software requirements 22, 23, 24
 - troubleshooting 184
 - verifying on OS/2 164
 - Tivoli Enterprise
 - overview 219
 - Tools Setting 231
 - transparent parallelism 4
- U**
- username
 - naming rules 209
- V**
- verifying the connection
 - TCP/IP 183
 - viewing
 - online information 250
 - Visual Explain
 - overview 232
- W**
- Warehouse Manager
 - overview 217
 - well-known address
 - description 12
 - Windows 2000
 - starting the security server 159
 - Windows NT
 - starting the security server 159
 - wizards
 - add database 252, 253
 - backup database 252
 - completing tasks 252
 - configure multisite update 252
 - create database 253
 - create table 253
 - create table space 253
 - index 253
 - performance configuration 253
 - restore database 253
 - Workgroup Edition
 - overview 215
 - working with DB2 data 219
 - workstation name (nname)
 - naming rules 209

Contacting IBM

If you have a technical problem, please review and carry out the actions suggested by the *Troubleshooting Guide* before contacting DB2 Customer Support. This guide suggests information that you can gather to help DB2 Customer Support to serve you better.

For information or to order any of the DB2 Universal Database products contact an IBM representative at a local branch office or contact any authorized IBM software remarketer.

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

- 1-800-237-5511 for customer support
- 1-888-426-4343 to learn about available service options

Product Information

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

- 1-800-IBM-CALL (1-800-426-2255) or 1-800-3IBM-OS2 (1-800-342-6672) to order products or get general information.
- 1-800-879-2755 to order publications.

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

The DB2 World Wide Web pages provide current DB2 information about news, product descriptions, education schedules, and more.

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

The DB2 Product and Service Technical Library provides access to frequently asked questions, fixes, books, and up-to-date DB2 technical information.

Note: This information may be in English only.

<http://www.elink.ibm.com/pbl/pbl/>

The International Publications ordering Web site provides information on how to order books.

<http://www.ibm.com/education/certify/>

The Professional Certification Program from the IBM Web site provides certification test information for a variety of IBM products, including DB2.

ftp.software.ibm.com

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

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

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

On CompuServe: GO IBMDB2

Enter this command to access the IBM DB2 Family forums. All DB2 products are supported through these forums.

For information on how to contact IBM outside of the United States, refer to Appendix A of the *IBM Software Support Handbook*. To access this document, go to the following Web page: <http://www.ibm.com/support/>, and then select the IBM Software Support Handbook link near the bottom of the page.

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



Part Number: CT79LNA



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

GC09-2964-01



(1P) P/N: CT79LNA

