IBM DB2 Connect

IBM

User's Guide

Version 6

IBM DB2 Connect

User's Guide

Version 6

Before using this information and the product it supports, be sure to read the general information under "Appendix I. Notices" on page 171.

Contents

About This Book

This book contains general information about Version 6 of the following IBM DB2 Connect products, formerly known as Distributed Database Connection Services (DDCS):

- DB2 Connect Personal Edition, formerly known as DDCS Single-User.
- DB2 Connect Enterprise Edition (EE), formerly known as DDCS Multi-user Gateway, for AIX, HP-UX, SCO, Solaris, OS/2, and Windows NT.

This book is intended for programmers and administrators who are responsible for setting up and maintaining DB2 Connect connections between DB2 Universal Database clients and any of the following DRDA application server database management systems:

- DATABASE 2 for OS/390 (DB2 Universal Database for OS/390) Version 5 or higher
- DATABASE 2 for MVS (DB2 for MVS) Version 3 or higher
- DATABASE 2 for VSE and VM (DB2 for VSE & VM)
- DATABASE 2 for AS/400 (DB2 Universal Database for AS/400)
- Any other relational database management system that implements Distributed Relational Database Architecture (DRDA) application server function.

Notes:

1. DB2 Universal Database Version 5 does not require DB2 Connect in order to allow host applications to access DB2 UDB data.
2. DB2 Universal Database for OS/390 Version 5.1 or higher is required in order to use DRDA Level 3 functions, including TCP/IP database connections, and stored procedures with multi-row answer sets.
3. DB2 Universal Database for OS/390 Version 6.1 or higher is required in order to use the DRDA Level 4 functions for which support is provided in DB2 Connect Version 6. These include: support for big integer, large object, row ID, and user-defined distinct data types.

This book explains concepts that are applicable to both DB2 Connect products. For information about a specific platform refer to either:

- *DB2 Connect Personal Edition Quick Beginnings*, for single-user DB2 Connect setup on OS/2, Windows 3.1, Windows 3.11 for Workgroups, Windows 95, or Windows NT.
- *DB2 Connect Enterprise Edition for OS/2 and Windows NT Quick Beginnings*, for multi-user DB2 Connect gateway setup on AIX, HP-UX, OS/2, SCO, Solaris, or Windows NT.

This book also describes features of IBM DB2 Connect products and provides general information about setting up connections to host and AS/400 database servers. In order to fully understand the configuration and communication protocols, you should be familiar with Systems Network Architecture (SNA) protocols, TCP/IP, DRDA, and your host or AS/400 operating system.

DB2 Connect is part of the DB2 Universal Database family of products.

How This Book is Structured

This book contains the following major sections:

- "Appendix I. Notices" on page 171.

Other Information Sources

This section lists other information sources which may be useful.

Using the World Wide Web

You can find the most recent information about DB2 Connect, DB2 Universal Database, and other IBM software products on the World Wide Web. This includes the latest publications, as well as technical hints and tips in the form of Technotes:

1. Set your Web browser to the following URL:

   ```
   http://www.software.ibm.com/data/db2/library/
   ```

2. Select "DB2 Universal Database".

3. For example, search for "Technotes" using the keywords "DDCS", "DRDA", or "Connect".

Related DRDA Publications

The following books contain related information and may be referenced in this manual.

Form number	Book title
SC26-4783	*Distributed Relational Database Architecture Connectivity Guide*
SC26-4773	*Distributed Relational Database Architecture Application Programming Guide*
SC26-4782	*Distributed Relational Database Architecture Problem Determination Guide*
SC26-4650	*Planning for Distributed Relational Database Architecture*
GC26-3195	*Distributed Relational Database Architecture Every Manager's Guide*
G321-5482	*IBM Distributed Data Management Architecture Level 3: Reference*

Related DRDA Server Publications

Related DRDA server publications include the following books from the DB2 Universal Database for AS/400, DB2 for OS/390, and DB2 for VSE & VM libraries.

Form number	Book title
SC41-5702	*AS/400 Distributed Database Programming*
SC41-9609	*AS/400 SAA Structured Query Language/400 Programmer's Guide*
SC41-9608	*AS/400 SAA Structured Query Language/400 Reference*
GC21-8180	*AS/400 Communications Configuration Reference*
SC26-8958	*DB2 Universal Database for OS/390 Application Programming and SQL Reference*
SC26-8960	*DB2 Universal Database for OS/390 Command Reference*
GC26-8970	*DB2 Universal Database for OS/390 Installation Reference*
SC26-8964	*DB2 Universal Database for OS/390 Reference for Remote DRDA Requesters and Servers*
SC26-8966	*DB2 Universal Database for OS/390 SQL Reference*
SC26-8957	*DB2 Universal Database for OS/390 Administration Guide*
SC26-8967	*DB2 Universal Database for OS/390 Utility Guide and Reference*
SH09-8087	*DB2 for VSE & VM SQL Reference*
SC26-3255	*IBM SQL Reference*

Other Related Publications

Form number	Book title
SG24-2006	*Migrating to DB2 Universal Database Version 5*
SG24-2213	*DB2 for OS/390 Version 5 Performance Topics*
SG24-4893	*DB2 Meets NT*
SG24-4894	*The Universal Connectivity Guide to DB2*
SG24-4693	*Getting Started with DB2 Stored Procedures*
SG24-2212	*DRDA Support for TCP/IP in DB2 Universal Database for OS/390 V5.1 and DB2 Universal Database V5.0*
SC33-0814	*CICS for AIX Application Programming Guide*

Form number	Book title
SC33-0931	*CICS for AIX Customization and Operation Guide*
GC09-2829-00	*DB2 Connect Enterprise Edition for UNIX Quick Beginnings*
GC09-2828-00	*DB2 Connect Enterprise Edition for OS/2 and Windows Quick Beginnings*
GC09-2830-00	*DB2 Connect Personal Edition Quick Beginnings*
GG24-4155	*Distributed Relational Database Architecture: Using DDCS for AIX DRDA support with DB2 for MVS/ESA and DB2 Universal Database for AS/400*
GG24-4311	*Distributed Relational Database Architecture Cross Platform Connectivity and Application*
SC23-2443	*Encina for AIX Product Family Overview*

Chapter 1. DB2 Connect Overview

Many organizations now use personal computers and UNIX workstations extensively, while at the same time still keeping much of their most important data on mainframes or minicomputers. Providing access to enterprise data for applications running on PC and UNIX applications is a significant challenge.

IBM DB2 Connect products provide an elegant solution to this situation. They enable database application programs running on DOS, Macintosh, OS/2, UNIX and Windows workstations to access data stored in relational databases on MVS, OS/400, OS/390, VSE and VM hosts. Data stored in certain non-relational databases such as IMS can also be accessed when DB2 Connect is used in conjunction with data replication products such as IBM Data Propagator Non-Relational.

You can also use the DB2 Universal Database client which is included in DB2 Connect to access databases stored on OS/2, Windows, and UNIX systems.

DB2 Connect provides applications with transparent online access to data by implementing a standard architecture for managing distributed data, Distributed Relational Database Architecture (DRDA). For example, you can use DB2 Connect with:

Spreadsheets, such as Lotus 1-2-3** and Microsoft Excel**, to analyze real-time data without the cost and complexity of data extract and import procedures

Decision support tools such as Intersolv Q+E Database Editor, and Crystal Reports, to provide real-time information

Database products such as MS Access** and Lotus Approach**

Development tools such as IBM VisualAge, PowerBuilder, Microsoft VisualBasic**, WATCOM VX-REXX, and VisPro/REXX, to create client/server solutions.

Internet servers such as Lotus Domino, Netscape Enterprise Server, and Microsoft Internet Information Server.

Java tools such as IBM Visual Age for Java, and Symantec Visual Cafe Pro.

IBM Net.Data.

The workstation on which DB2 Connect is installed is called the *DB2 Connect server*. DRDA application servers can be accessed by local clients on the DB2 Connect workstation, and by remote clients. For remote clients, A DB2 Connect server acts as an intermediary between the client workstation and the host or AS/400 database management system.

Figure 1 illustrates how DB2 Connect connects to other systems.
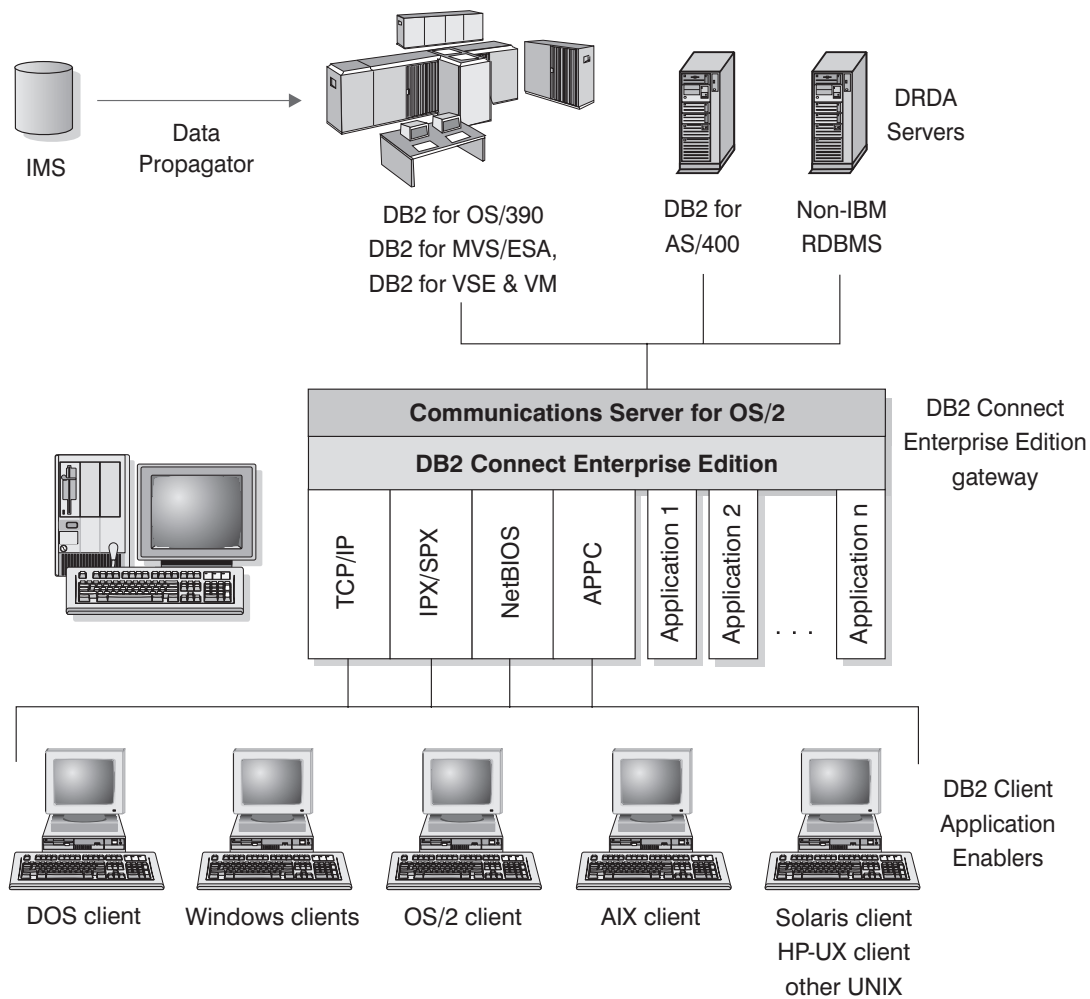
Figure 1. DB2 Connect Connections (OS/2 Example)

Notes:

1. This figure illustrates DB2 Connect Enterprise Edition running on an OS/2 system. DB2 Connect Enterprise Edition is implemented similarly on all supported platforms.

2. This figure includes remote clients. If you do not want remote clients, you do not need to install remote client support. In OS/2 and Windows environments, IBM also offers DB2 Connect Personal Edition, which supports a single user at a time, and does not provide gateway support for downstream remote clients.

3. This figure illustrates remote connections using TCP/IP, APPC, IPX/SPX, or NetBIOS. Depending on the operating system of the DB2 Connect workstation and the operating system of the remote client, different communications protocols may be available. For detailed information about communications protocol support, see *DB2 Connect Personal Edition Quick Beginnings*, or *DB2 Connect Enterprise Edition for OS/2 and Windows NT Quick Beginnings*.

4. You do not need to have DB2 Universal Database installed on the DB2 Connect workstation. If you want a complete relational database management system on the DB2 Connect workstation, order DB2 Universal Database.

5. DB2 Connect does not provide application development tools; however, IBM does provide DB2 software developer kits (SDKs) for DOS, Macintosh, OS/2, UNIX and Windows platforms. You can use these kits to develop applications that work with DB2 Connect. SDKs provide programming tools, utilities, documentation, and code samples. See *Application Building Guide*

6. C programmers developing Windows applications that use the Microsoft** ODBC interface should use the *Microsoft Open Database Connectivity Software Development Kit.* Programmers who want to develop applications using the COBOL programming language can use Micro Focus** COBOL and other development tools offered by Micro Focus, Inc., and IBM. In addition, there are a variety of 4GL development environments, application generators, CASE and other tools available from a number of vendors that can be used to develop applications that use DB2 Connect.

The Database Concept

The term *database* is used throughout this book to describe a relational database management system (RDBMS). Other systems with which DB2 Connect communicates may use the term *database* to describe a slightly different concept. The DB2 Connect term *database* is equivalent to:

MVS   A DB2 for MVS/ESA subsystem identified by its LOCATION NAME.

This name is defined in the Boot Strap Data Set (BSDS) and is also provided in the DSNL004I message (LOCATION=location) which is written when the Distributed Data Facility (DDF) is started.

OS/390
A DB2 Universal Database for OS/390 subsystem identified by its LOCATION NAME.

This name is defined in the Boot Strap Data Set (BSDS) and is also provided in the DSNL004I message (LOCATION=location) which is written when the Distributed Data Facility (DDF) is started.

VSE DB2 for VSE running in a partition identified by its DBNAME

VM DB2 for VM running in a CMS virtual machine identified by its DBNAME

OS/400
DB2 Universal Database for AS/400, an integral part of the OS/400 operating system. Only one database can exist on an AS/400 machine. If the database will be used by applications outside the AS/400 system, the database must be given a name in the relational database directory.

Setting up DB2 Connect

Before you can use DB2 Connect, you must do the following:

1. Install DB2 Connect, and configure both the host or AS/400 server, and workstation communications, as described in either *DB2 Connect Personal Edition Quick Beginnings*, or *DB2 Connect Enterprise Edition for OS/2 and Windows NT Quick Beginnings*.

2. Update the database directories as described in "Chapter 3. Updating Database Directories" on page 19.

   Note that on OS/2, Windows 95, or Windows NT the recommended method is to use the DB2 Universal Database Client Configuration Assistant, and on Windows 3.1 or Windows for Workgroups 3.11 this can be done using the Data Sources Setup dialog.

   On all other platforms the database directories must be updated using the DB2 Command Line Processor (CLP). Both approaches are described in *DB2 Connect Personal Edition Quick Beginnings*, and *DB2 Connect Enterprise Edition for OS/2 and Windows NT Quick Beginnings*.

3. Bind the DB2 Connect utilities to each host or AS/400 database management system, as described in "Chapter 4. Binding Applications and Utilities" on page 33.

   This task can also be performed using the DB2 Universal Database Client Configuration Assistant or the Data Sources Setup dialog where provided.

DB2 Connect and SQL

DB2 Connect forwards SQL statements submitted by application programs to host or AS/400 database servers. DB2 Connect can forward almost any valid SQL statement. The exceptions are documented in "Host or AS/400 Server SQL Statements Rejected by DB2 Connect" on page 84.

There are two types of SQL processing: static SQL and dynamic SQL. Static SQL minimizes the time required to execute an SQL statement by processing it ahead of time. Dynamic SQL is processed when the SQL statement is submitted to the host or AS/400 database server. This makes dynamic SQL more flexible, but potentially slower. The decision to use static or dynamic SQL is made by the application programmer. Both are supported by DB2 Connect.

Different host or AS/400 database servers implement SQL differently. For information about common SQL statements that are supported by all IBM systems, see the *SQL Reference.*

DB2 Connect fully supports the common IBM SQL, as well as the DB2 Universal Database for OS/390, DB2 for MVS/ESA, DB2 for VSE & VM (formerly SQL/DS), and DB2 Universal Database for AS/400 implementations of SQL. IBM SQL is strongly recommended for maintaining database independence. For more information, see "Chapter 8. Programming in a DB2 Connect Environment" on page 71.

Administration Utilities

The following utilities are available to help the DB2 Connect administrator:
- The command line processor lets you issue SQL statements against a host or AS/400 database server database. It flows the SQL statements to the database that you specify.
- Import and export utilities let you copy data to and from a file on a workstation and a host or AS/400 database server database. These files can then be used for importing data into databases, spreadsheets, and other applications running on your workstation.
- The DB2 Connect trace utility lets application developers analyze the flow of the DRDA data stream between the DB2 Connect workstation and the host or AS/400 database server database management system.
- The DB2 Control Center lets administrators create and schedule data replication jobs as well as perform other database administration tasks.
- The DB2 Command Center provides a more graphical version of the command line processor.

For more information about these utilities, see "Chapter 5. DB2 Connect Administration Utilities" on page 41.

In addition, the database system monitor utility lets the system administrator monitor system connections. It also helps the system administrator determine the source of an error. The system administrator can correlate client applications with the corresponding jobs running on the host or AS/400 database server. For more information, see "Chapter 6. Database System Monitor for DB2 Connect" on page 55.

Functions Delivered in Previous Releases

Beginning with the most recent version and release, a summary of the enhancements introduced at each version and release is presented.

DB2 Connect Version 5 Release 2

DB2 Connect Version 5.2 included the following enhancements for the host and AS/400 DRDA functions:
- DCE Cell Directory support
- Enhanced security failure notification
- Enhanced System/390 SYSPLEX exploitation
- Optimized Catalog Access for ODBC and JDBC Applications
- Microsoft Transaction Server support
- New BIND options (DYNAMICRULES)
- Set Client Information API enhancements
- SQLDescribeParam support for DB2 Connect
- Support for Bidirectional Languages
- System Monitor enhancements
- Two-phase commit support enhancements
- Simplified DB2 Syncpoint Manager Configuration
- Support for the SCO** operating system
- Support for Big Integer, Large Object, Row ID, and User Defined Distinct data types.

DB2 Connect Version 5.0
- New easier to purchase packaging:
  - A single DB2 Connect Personal Edition package that contains OS/2, Windows 3.1, Windows 95, and Windows NT versions of the product. This package contains everything that is needed to get started, including a complimentary copy of Lotus Approach.

- A single DB2 Connect Enterprise Edition package that contains OS/2, Windows NT, and all UNIX versions.
- Capability:
  - New Level 3 ODBC driver with many improvements
  - Updated JDBC driver for better Java support
  - Support for stored procedures that return multi-row result sets and multiple result sets (requires DB2 Universal Database for OS/390 Version 5.1 or higher)
  - Built-in replication support
  - Generic bind option: you can specify any bind option supported by the host database.
  - SYSPLEX exploitation (DB2 Connect Personal Edition only; requires DB2 Universal Database for OS/390 Version 5.1 or higher)
- Usability:
  - New installation method
  - TCP/IP database connections are much easier to configure (requires DB2 Universal Database for OS/390 Version 5.1 or higher, or DB2 Universal Database for AS/400 Version 4.2)
  - Integrated SNA support with point-and click configuration (DB2 Connect Personal Edition only)
  - New point and click configuration utility for configuring host connections.
  - Much easier process for connecting desktop client systems to DB2 Connect Enterprise Edition servers. Clients can discover DB2 Connect servers and all of the databases that are defined on each server
  - Improved ODBC traces with detailed information for performance analysis
  - Control Center and other GUI tools that simplify several DBA tasks
- Security:
  - DCE security (requires DB2 Universal Database for OS/390 Version 5.1 or higher)
  - Ability to run ODBC applications without having to authorize each user to base tables. Users can now bind their ODBC driver in such a way as to allow applications to run under the authority of the person that bound the ODBC driver.
- Performance:
  - Faster access to the DB2 catalog for ODBC applications
  - Reduced network traffic:
    - Early close for cursors
    - Deferred prepare

- Reduced byte count on Compound SQL
- Several other network flow enhancements
- Support for ASCII storage on the host (requires DB2 Universal Database for OS/390 Version 5.1 or higher)

- Connectivity:
  - Support for DRDA over TCP/IP connections to other IBM DRDA Application Servers, as they introduce support for TCP/IP.
  - SNA over TCP/IP via integrated MPTN support (requires AnyNet on the host).
  - Support for additional SNA connectivity options:
    - IBM Communication Server for Windows NT
    - IBM Personal Communications
- Other:
  - Ability to initiate 2-phase commit transactions over TCP/IP (requires DB2 Universal Database for OS/390 Version 5.1 or higher)
  - Ability for desktop applications to participate in a 2-phase commit transactions without the need for a gateway (TCP/IP only, requires DB2 Universal Database for OS/390 V5.1 or higher)
  - Ability to use DB2 Universal Database for OS/390 for added reliability of transaction coordination (requires DB2 Universal Database for OS/390 Version 5.1 or higher, and TCP/IP)
  - Numerous other enhancements and fixes affecting all aspects of system performance, reliability, and usability.

DDCS Version 2 Release 4

Distributed Database Connection Services (DDCS) for Windows Single-User Version 2.4 introduced:

- A Data Sources Setup tool to help you define connections to host and AS/400 servers quickly and easily.
- Wall Data Rumba, to provide you with the communications support required to make these connections.
- A DB2 Password Expiration Maintenance utility (DB2PEM), which enables you to change your DB2 for MVS/ESA password without logging on to TSO.
- Enhancements to improve the performance and flexibility of DB2 Connect:
  - Deferred Prepare, which improves the performance of ODBC and other dynamic SQL applications by attaching the PREPARE request to a subsequent request instead of sending it separately.
  - Asynchronous ODBC, which improves the availability of ODBC applications. Previously, these might have appeared to be delayed while processing long queries in some network situations.

– On AIX and OS/2, support for multi-threaded applications, which gives non-ODBC applications the ability to maintain multiple database connections with their own contexts.

DDCS Version 2 Release 3

New features in DDCS Version 2 Release 3.1 included:
- Two-phase commit for DRDA connections using the LU6.2 Syncpoint Manager (SPM) on OS/2 and AIX.

New features in DDCS Version 2 Release 3.0 included:
- Client application performance could be improved by running stored procedures on DB2 for MVS/ESA Version 4.1 and DB2 Universal Database for AS/400 Version 3.1 servers. See "Stored Procedures" on page 80.
- Able to work with multiple databases in a single transaction. See "Multisite Updates (Two-Phase Commit)" on page 13.
- Able to improve performance by concatenating SQL statements. See "NOT ATOMIC Compound SQL" on page 81 and "Using Import and Export Utilities" on page 41.
- Able to implement chargeback accounting by using accounting strings. See "Implementing Chargeback Accounting" on page 84.
- Able to use many new bind options when binding applications to a DRDA application server. See "The BIND Command" on page 38.
- When using a DCE directory, the ability to consolidate the directory information needed by all your clients in a central repository. See "Appendix E. Using DCE Directory Services" on page 143.
- Greater flexibility in SQLCODE processing. See "Chapter 9. SQLCODE Mapping" on page 89.
- Diagnostic information stored in a readable format and consolidated in a single location (the first failure service log). For more information, see the *Troubleshooting Guide*.
- The DDCSSETP environment variable was replaced by BIND and PREPARE options such as `SQLERROR CONTINUE`, simplifying operations.
- Various other performance improvements were also implemented.

Chapter 2. Distributed Relational Database Architecture Concepts

Distributed Relational Database Architecture (DRDA) is a set of protocols that permits multiple database systems, both IBM and non-IBM, as well as application programs to work together. Any combination of relational database management products that use DRDA can be connected to form a distributed relational database management system. DRDA coordinates communication between systems by defining what must be exchanged and how it must be exchanged.

DRDA and DB2 Connect

DB2 Connect implements the DRDA architecture to reduce the cost and complexity of accessing data stored in DB2 Universal Database for AS/400, DB2 Universal Database for OS/390, DB2 for MVS/ESA, DB2 for VSE & VM, and other DRDA-compliant database servers. By fully exploiting the DRDA architecture, DB2 Connect offers a well-performing, low-cost solution with the system management characteristics that customers demand.

In DRDA terminology, an *application requester* is the code that handles the application end of a distributed connection; it is the application that is requesting data. An *application server* is the code that handles the database end of the connection. In the DB2 Connect environment, the DB2 Connect workstation can only function as an application requester on behalf of application programs.

Figure 2 on page 12 shows the flow of data between the DB2 Connect workstation and the DRDA server in the case where there are local clients only. In addition, a private protocol exists between the DB2 Connect workstation and any remote clients.

11
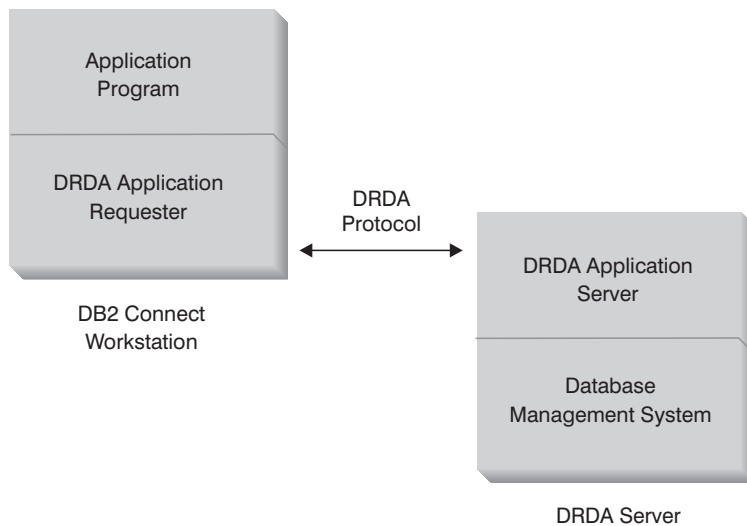
Figure 2. DRDA Flow in DB2 Connect

To implement the connections between DRDA server database management systems and database clients, DRDA uses the following architectures:

- Character Data Representation Architecture (CDRA)
- Distributed Data Management Architecture (DDM)
- Formatted Data Object Content Architecture (FD:OCA)
- Systems Network Architecture (SNA)
- SNA Management Services Architecture (MSA)
- Transmission Control Protocol/Internet Protocol (TCP/IP).

These are used as architectural building blocks. The data streams which flow over the network are specified by DRDA architecture, which documents a data stream protocol supporting distributed relational database access.

A request is routed to the correct destination by means of directories that contain various types of communication information and the name of the DRDA server database being accessed.

Remote Unit of Work

A unit of work is a single logical transaction. It consists of a sequence of SQL statements in which either all of the operations are successfully performed or the sequence as a whole is considered unsuccessful.

Remote unit of work lets a user or application program read or update data at one location per unit of work. It supports access to one database within a unit of work. While an application program can update several remote databases, it can only access one database within a unit of work.

Remote unit of work has the following characteristics:
- Multiple requests (SQL statements) per unit of work are supported.
- Multiple cursors per unit of work are supported.
- Each unit of work can access only one database.
- The application program either commits or rolls back the unit of work. In certain error circumstances, the database server or DB2 Connect may roll back the unit of work.

Multisite Updates (Two-Phase Commit)

Multisite update, also known as Distributed Unit of Work (DUOW) and Two-Phase commit, is a function that enables your applications to update data in multiple remote database servers with guaranteed integrity. A good example of a multisite update is a banking transaction that involves transfer of money from one account to another in a different database server. In such a transaction it is critical that updates that implement debit operation on one account do not get committed unless updates required to process credit to the other account are committed as well. The multisite update considerations apply when data representing these accounts is managed by two different database servers.

DB2 products provide comprehensive support for multisite update. This support is available for applications developed using regular SQL as well as applications that utilize Transaction Monitor products that implement X/Open XA interface specification. Examples of such Transaction Monitor products include IBM TxSeries (CICS and Encina), Message and Queuing Series, Component Broker Series, San Francisco Project as well as Microsoft Transaction Server (MTS), BEA Tuxedo, NCR TopEnd and several others. There are different setup requirements depending on whether native SQL multisite update or TP Monitor multisite upate is used.

Both the native SQL and TP Monitor multi-site update programs must be precompiled with the `CONNECT 2 SYCNCPOINT TWOPHASE` options. Both can use the SQL Connect statement to indicate which database they want to be used for the SQL statements that follow. If there is no TP Monitor to tell DB2 it is going to coordinate the transaction (as indicated by DB2 receiving the `xa_open` calls from the TP monitor to establish a database connection), then the DB2 software will be used to coordinate the transaction.

When using TP monitor multi-site update, the application must request commit or rollback by using the TP monitor's API, e.g. CICS SYNCPOINT, Encina Abort(), MTS SetAbort(). When using native SQL muilti-update, the normal SQL COMMIT and ROLLBACK must be used.

TP Monitor multi-site update can coordinate a transaction that accesses both DB2 and non-DB2 resource managers such as Oracle, Informix, SQLServer, etc. Native SQL multisite update is used with DB2 servers only.

For a multisite update transaction to work each of the databases participating in a distributed transaction must be capable of supporting Distributed Unit of Work. At the time of this writing the following DB2 servers provided DUOW support that enabled them to participate in distributed transactions:
• DB2 Common Server V2
• DB2 Universal Database V5 and V6
• DB2 for MVS/ESA V3.1 and 4.1
• DB2 for OS/390 V5.1
• DB2 Universal Database for OS/390 V6.1
• DB2/400 V3.1 or later
• DB2 Server for VM and VSE V5.1

A distributed transaction can update any mix of supported database servers. For example, your application can update several tables in DB2 UDB on Windows NT, a DB2 for OS/390 database and a DB2/400 database all within a single transaction.

Host and AS/400 database servers require DB2 Connect to participate in a distributed transaction originating from PC, UNIX, and web applications. In addition many of the multisite update scenarios that involve host and AS/400 database servers require that the Syncpoint Manager (SPM) component be configured. The need for SPM is dictated by the choice of protocol (SNA vs. TCP/IP) and use of a TP monitor. See Table 1 on page 15 for a summary of scenarios that require use of the SPM. The table shows that DB2 Connect is required for any access to the host or AS/400 from Intel or UNIX machines. In addition, for multisite updates, the SPM component of DB2 Connect is required if the access is via SNA or uses a TP monitor.

Table 1. Host and AS/400 Multisite Update Scenarios that Require SPM

Host and AS/400 multisite update scenarios that require SPM.				
TP Monitor Used?	Protocol	SPM Needed?	Product Required (choose One)	Host and AS/400 Database Supported
Yes	TCP/IP	Yes	• DB2 Connect Enterprise Edition • DB2 UDB Enteprise Edition • DB2 UDB Extended Enterprise Edition	• DB2 for OS/390 V5.1 • DB2 UDB for OS/390 V6.1
Yes	SNA	Yes	• DB2 Connect Enterprise Edition* • DB2 UDB Enterprise Edition* • DB2 UDB Extended Enterprise Edition* **Note:** *AIX, OS/2 and Windows NT platforms only.	• DB2 for MVS/ESA V3.1 and 4.1 • DB2 for OS/390 V5.1 • DB2 UDB for OS/390 V6.1 • DB2/400 V3.1 or later • DB2 Server for VM or VSE V5.1
No	TCP/IP	No	• DB2 Connect Personal Edition • DB2 Connect Enterprise Edition • DB2 UDB Enterprise Edition • DB2 UDB Extended Enterprise Edition	• DB2 for OS/390 V5.1 • DB2 UDB for OS/390 V6.1

Table 1. Host and AS/400 Multisite Update Scenarios that Require SPM (continued)

No	SNA	Yes	• DB2 Connect Enterprise Edition* • DB2 UDB Enterprise Edition* • DB2 UDB Extended Enterprise Edition* **Note:** *AIX, OS/2 and Windows NT platforms only	• DB2 for MVS/ESA V3.1 and 4.1 • DB2 for OS/390 V5.1 • DB2 UDB for OS/390 V6.1 • DB2/400 V3.1 or later • DB2 Server for VM and VSE V5.1

For more information about two-phase commit, see:

- *Administration Guide*
- *DB2 Connect Enterprise Edition for OS/2 and Windows NT Quick Beginnings*
- *DB2 Connect Personal Edition Quick Beginnings* (this edition does not include the DB2 Syncpoint Manager).

For more information about Transaction Monitors, and setting up a TP monitor environment, see:

- *CICS for AIX Customization and Operation Guide*
- *DB2 Connect Enterprise Edition for OS/2 and Windows NT Quick Beginnings.*

DRDA and Data Access

Although DRDA defines database communication protocols, it does not define the programming interfaces, or APIs, that should be used by application programmers. In general, DRDA can be used by an application program to pass any request that a target DRDA server can execute. All of the DRDA servers available today can execute SQL requests forwarded by an application program via DB2 Connect.

IBM provides application programmers with tools to generate SQL requests for DOS, Windows, OS/2, and several UNIX platforms. These tools are part of the DB2 Software Developer Kits (DB2 SDKs). DB2 SDKs support two different API types: embedded SQL and the DB2 Call Level Interface (DB2 CLI). These can be used by programmers building applications in a variety of programming languages. For more information about these APIs, see *Application Building Guide*.

Application developers can also use APIs provided by other companies. For example, the Microsoft ODBC API is used by many Windows application programmers to develop database applications. DB2 Connect provides an ODBC driver that supports applications developed using the ODBC API. IBM does not provide tools for developing ODBC applications; these tools are provided by Microsoft Corporation.

Chapter 3. Updating Database Directories

DB2 Connect uses the following directories to manage information about databases that it connects to:

- The *node directory*, which contains network address and communication protocol information for every host or AS/400 database server that DB2 Connect accesses.
- The *database connection services* (DCS) directory, which contains information specific to host or AS/400 database server databases.
- The *system database directory*, which contains name, node, and authentication information for every database that DB2 Connect accesses.

Notes:

1. Before updating these directories, you should configure communications on the host or AS/400 database server and workstations. This is described in *DB2 Connect Personal Edition Quick Beginnings*, and *DB2 Connect Enterprise Edition for OS/2 and Windows NT Quick Beginnings*.

2. Note that on OS/2, Windows 95, and Windows NT, database directories can be updated using the DB2 Universal Database Client Configuration Assistant (CCA). On Windows 3.1 and Windows for Workgroups 3.11, the database directories can be updated using the Data Sources Setup tool.

   On all other platforms the database directories must be updated using the DB2 Command Line Processor (CLP).

3. "Updating the Directories" on page 30 provides example command syntax. For further information refer to the *Command Reference*.

4. If you are using DCE, for each host or AS/400 database server database that you connect to you must update these directories or store equivalent information in a global DCE directory. For more information about DCE, see "Appendix E. Using DCE Directory Services" on page 143 and the *Administration Guide*. This chapter assumes that you are *not* using DCE Directory Services.

Collecting Information

"Appendix C. Directory Customization Worksheet" on page 137 shows the information that you need to collect. You may find it convenient to make a copy of the worksheet and write in the values for your system.

Node Directory

You can specify the following information in the node directory:

Node name
> A nickname for the host or AS/400 database server system on which the remote database resides. This name is user-defined. Write the same node name in both the Node Directory Parameters table and the System Database Directory Parameters table.

> Format: 1–8 single-byte alphanumeric characters, including the number sign (#), at sign (@), dollar sign ($), and underscore (_). It cannot begin with an underscore or a number.

Protocol
> Can be APPC or TCPIP.

Symbolic destination name
> When defining an APPC node, use the symbolic destination name that was specified in the CPI Communications Side Information Table (for example, the name of the CPI-C Symbolic Destination Properties when using Microsoft SNA Server). You should get this value from the person who either installed and/or configured SNA. The symbolic destination name is case sensitive (you may encounter an SQL1338 return code if there is a mismatch between upper and lower case names).

Security type
> The type of security checking that will be done. For APPC nodes, the valid options are SAME, PROGRAM, and NONE. For TCP/IP nodes, SECURITY SOCKS is an option which specifies that the node will be SOCKS-enabled, in which case the SOCKS_NS and SOCKS_SERVER environment variables are mandatory and must be set to enable SOCKS. For more information, see "Chapter 7. Security" on page 63, and refer to the *Command Reference*.

TCP/IP remote hostname
> When defining a TCP/IP node, either the remote TCP/IP hostname, or the remote TCP/IP address. If a hostname is specified, then it must be resolved at the DB2 Connect workstation, either through Domain Name Server (DNS) lookup, or by an entry in the local TCP/IP hosts file.

> For DB2 for OS/390 remote hosts, the hostname appears in the DSNL004I message (DOMAIN=hostname) when the Distributed Data Facility (DDF) is started.

TCP/IP service name
> When defining a TCP/IP node, either the remote TCP/IP service

name or port number. This must be defined to TCP/IP at the remote host. Port number 446 has been registered as the default port number for DRDA.

For DB2 for OS/390 remote hosts, the port number is defined in the Boot Strap Data Set (BSDS) as PORT and is also provided in the DSNL004I message (TCPPORT=portnumber) when the Distributed Data Facility (DDF) is started.

Note: A second port used for two-phase commit resynchronization operations over TCP/IP connections is assigned by the server. For example, the DB2 Universal Database for OS/390 bootstrap dataset assigns a port number (RESPORT) to be used for resynchronization for inbound connections to DB2 Universal Database for OS/390 only. No service name need be defined for this.

DCS Directory

You can specify the following information in the DCS directory:

Database name
A user-defined nickname for the host or AS/400 database database server. Use the same database name in both the DCS Directory Parameters table and the System Database Directory Parameters table.

Format: 1–8 single-byte alphanumeric characters, including the number sign (#), at sign (@), dollar sign ($), and underscore (_). It cannot begin with an underscore or a number.

Target database name
The database on the host or AS/400 database server system, as follows:

MVS/ESA
The LOCATION value as defined in the DB2 for MVS/ESA Bootstrap Dataset (BSDS) ″LOCATION=locname″ field and which is also provided in the DSNL004I message (LOCATION=location) when the Distributed Data Facility (DDF) is started.

OS/390
The LOCATION value as defined in the DB2 Universal Database for OS/390 Bootstrap Dataset (BSDS) ″LOCATION=locname″ field and which is also provided in the DSNL004I message (LOCATION=location) when the Distributed Data Facility (DDF) is started.

VSE or VM
The database name (DBNAME)

OS/400

>The relational database name (RDBNAME)

Other For OS/2, Windows NT, and UNIX-based systems, the database alias.

Application requester name

>The name of the application requester that forwards SQL requests to DRDA application servers. The application requester handles requests on behalf of an application program.

>Format: AR <application_requester_name>

>The default is the DB2 Connect application requester.

Parameter string

>If you want to change the defaults, specify any or all the following parameters in the following order. Note that the parameter string cannot be set using the Client Configuration Assistant, and that when using the CLP the parameter string must be enclosed in either single quotes (for example on OS/2 or Windows NT), or in double quotes (for example on AIX):

>*map-file*

>>The name of an SQLCODE mapping file that overrides the default SQLCODE mapping. To turn off SQLCODE mapping, specify **NOMAP**. For more information, see "Chapter 9. SQLCODE Mapping" on page 89 .

>**,D** This is the second positional parameter. If it is specified the application will disconnect from the host or AS/400 database server database when one of the following SQLCODES is returned:

```
SQL30000N
SQL30040N
SQL30050N
SQL30051N
SQL30053N
SQL30060N
SQL30070N
SQL30071N
SQL30072N
SQL30073N
SQL30074N
SQL30090N
```

>>When the disconnect parameter **,D** is not specified, a disconnect will be performed only when the following SQLCODEs are returned:

```
SQL30020N
SQL30021N
SQL30041N
SQL30061N
SQL30081N
```

For explanations of these codes, see the *Message Reference.*

Note: If DB2 Connect disconnects due to an error, a rollback will be done automatically.

,,INTERRUPT_ENABLED

This is the third positional parameter. If INTERRUPT_ENABLED is configured in the DCS directory at the DB2 Connect workstation, and a client application issues an interrupt while connected to the host or AS/400 database server, DB2 Connect will perform the interrupt by dropping the connection and rolling back the unit of work. This interrupt behavior is supported on AIX, OS/2, and Windows NT.

The application will receive sqlcode (-30081) indicating that the connection to the server has been terminated. The application must then establish a new connection with the host or AS/400 database server, in order to process additional database requests. Note that on platforms other than AIX V4.1 and later, SNA Server V3.1 and later, OS/2, and Windows NT, DB2 Connect does not support the option of automatically disconnecting when an application using it receives an interrupt request.

Note: This support works for TCP/IP connections on any platforms. The client may kill the socket, but - depending on the server implementation - there may or may not be an outstanding receive. DB2 Universal Database for OS/390 utilizes asynchronous socket calls and therefore is able to detect the loss of the connection and roll back any long-running SQL statements that are in progress.

,,,,,SYSPLEX

This parameter, the 6th positional parameter, can be used to explicitly enable DB2 Connect SYSPLEX support for a particular database.

Note that a new profile (environment or registry) variable has also been introduced, called DB2SYSPLEX_SERVER, and it can be used to disable the SYSPLEX support at the workstation level.

,,,,,,LOCALDATE=″<value>″

This parameter, the seventh positional parameter, is used to enable DB2 Connect date formatting support. This is implemented using a date mask for the *<value>* as follows:

Suppose you issue the following CLP (command line processor) statements:

```
catalog appc node nynode remote nycpic security program
catalog dcs database nydb1 as new_york
catalog database nydb1 as newyork1 at node nynode
    authentication dcs
```

The database alias ″newyork1″ is to be used for accessing a host database without date transformation because no date mask has been specified.

However, with the new date formatting support, you can now use the following CLP commands. Note that in this case, because the CLP is being used, and the parameter string is itself being specified using double quotes, the LOCALDATE value has to be specified inside two pairs of double quotes.Note the use of the operating system escape character ″\″ (backslash) in order to ensure that the double quotes are not stripped from the LOCALDATE specification. See also "Specifying the Parameter String" on page 28.

```
catalog dcs database nydb2 as new_york
    parms \",,,,,,LOCALDATE=\"\"YYYYMMDD\"\"\"
catalog database nydb2 as newyork2 at node nynode
    authentication dcs
```

The database alias ″newyork2″ gives you access to the same host database but, in addition, it has a date format mask specified. This example illustrates that the date format mask is specified using the keyword LOCALDATE and is the seventh positional parameter in the PARMS field of a DCS directory entry.

In order for the date mask to be valid, ALL of the following must be true:

1. There can only be at most one sequence each of Y's, M's, and D's where Y is a year digit, M is a month digit, and D is a day digit.
2. The maximum number of Y's in a sequence is 4.
3. The maximum number of M's in a sequence is 2.
4. The maximum number of D's in a sequence is 2.

For instance, the following are all valid date masks:

```
"YYyyMmDd"   - Y, M, and D digits are case-insensitive
"MM+DD+YYYY" - OK to have a mask longer than 10 bytes
               and to have characters other than Y, M,
               and D in the mask
"abcYY+MM"   - OK not to have a sequence of D's
```

The following are all invalid date masks:

```
"YYYYyMMDD"  - invalid because there are 5 Y's
               in a sequence
"YYYYMDDM"   - invalid because there are 2 sequences
               of M's
```

If a date format mask is invalid, no error will be issued. It will just be ignored. Just because a date mask is valid does not mean it will be used. Date format transformation based on a valid date mask will only be performed if ALL of the following are true:

1. There is no SQL error.
2. The output is a date value in ISO-like (ISO and JIS) format.
3. The output data area is at least 10 bytes long. This is the minimum size of an output data area in order for a data value to be stored there even if NO date format transformation is to be performed. This requirement applies even if the date format mask ends up being shorter than 10 bytes.
4. There is a valid date format mask specified in the DCS directory entry and this mask fits in the output data area.

,,,,,,,CHGPWD_SDN=<name>

This parameter, the eighth positional parameter, is used to specify the symbolic destination name to be used for Password Expiration Management (PEM). The value specified for <name> is case sensitive.

The appendix ″Changing Your MVS Password″ in the DB2 Connect Quick Beginnings manuals, shows an

example of cataloging a dcs database directory using CHGPWD_SDN, as follows:

```
catalog dcs database db1 as dsn_db_1
    parms ",,,,,,,CHGPWD_SDN=pempgm"
```

,,,,,,,,BIDI=<ccsid>
This parameter, the ninth positional parameter, is used to specify the Bidirectional (BiDi) CCSID to be used to override the default server database BiDi CCSID. For example:

```
",,,,,,,,BIDI=xyz"
```

where **xyz** is the CCSID override (see 1 on page 27).

For a list of what BiDi CCSIDs are supported along with their string types, please refer to the *Administration Guide.*

The following BiDi attributes are required for correct handling of BiDi data on different platforms:

- Numeral shape (ARABIC vs HINDI)
- Orientation (RIGHT-TO-LEFT vs LEFT-TO-RIGHT)
- Shaping (SHAPED vs UNSHAPED)
- Symmetric swapping (YES or NO)
- Text type (LOGICAL vs VISUAL)

Since defaults on different platforms are not the same, problems appear when DB2 data is sent from one platform to another. For example, Windows platforms use LOGICAL UNSHAPED data, while data on MVS and OS/390 is usually in SHAPED VISUAL format. Therefore, without any support for BiDi attributes, data sent from DB2 for MVS or OS/390 to DB2 Connect on Windows displays incorrectly.

When data is exchanged between DB2 Connect and a database on a server, it is usually the receiver that performs conversion on the incoming data. The same convention would normally apply to BiDi layout transformation also, which is in addition to the usual code page conversion. However, currently no host DB2 product supports BiDi-specific CCSIDs or BiDi layout transformation. Therefore, DB2 Connect has been enhanced with the optional ability to perform BiDi layout transformation on data it is about to send to the server database in addition to data received from the server database.

In order for DB2 Connect to perform BiDi layout transformation on outgoing data to a server database, the BiDi CCSID of the server database will have to

be overridden (see 2). This is accomplished through the use of the BIDI parameter in the PARMS field of the DCS database directory entry for the server database.

The use of this feature is best illustrated with an example.

Suppose you have a Hebrew DB2 client running CCSID 62213 (BiDi string type 5) and you would like to access a DB2 host database running CCSID 424 (BiDi string type 4). However, you know that the data contained in the DB2 host database is instead based on CCSID 8616 (BiDi string type 6).

There are two problems in this situation. The first is that the DB2 host database does not know the difference between the BiDi string types with CCSIDs 424 and 8616. The second problem is that the DB2 host database does not recognize the DB2 client CCSID of 62213. It only supports CCSID 862, which is based on the same code page as CCSID 62213.

You will need to make sure that data sent to the DB2 host database is in BiDi string type 6 format to begin with and also let DB2 Connect know that it has to perform BiDi layout transformation on data it receives from the DB2 host database. You will use the following cataloging for the DB2 host database:

```
catalog dcs database nydb1 as TELAVIV parms ",,,,,,,,BIDI=8616"
```

This tells DB2 Connect to override the DB2 host database CCSID of 424 with 8616. This override includes the following processing:

1. DB2 Connect will connect to the DB2 host database using CCSID 862.
2. DB2 Connect will perform BiDi layout transformation on data it is about to send to the DB2 host database from CCSID 62213 (BiDi string type 5) to CCSID 62221 (BiDi string type 6).
3. DB2 Connect will perform BiDi layout transformation on data it receives from the DB2 host database from CCSID 8616 (BiDi string type 6) to CCSID 62213 (BiDi string type 5).

Notes:
1. The environment variable or registry value DB2BIDI has to be set to YES in order for the BIDI parameter to take effect.
2. Please note that if you would like DB2 Connect to perform layout transformation on data it is about to send to the DB2 host database even though you do not have to override its CCSID, you still have to add the BIDI parameter in the DCS database directory PARMS field. In this case, the CCSID that you should provide would be the default DB2 host database CCSID.
3. In some cases, use of a bidirectional CCSID may cause the SQL query itself to be modified such that it is not recognized by the DB2 server. Specifically, you should try to avoid using IMPLICIT CONTEXTUAL and

IMPLICIT RIGHT-TO-LEFT CCSIDs when a different string type can be used. CONTEXTUAL CCSIDs can produce unpredictable results if the SQL query contains quoted strings. Avoid using quoted strings in SQL statements, and use host variables instead when possible.

If a specific bidirectional CCSID is causing problems which cannot be rectified by following these recommendations, then you should set the environment variable or registry value DB2BIDI to NO.

Specifying the Parameter String

Here are examples of some parameter strings you could specify.

For example, you could specify any of the following where ″\″ (backslash) is the operating system escape character:

On AIX:

```
NOMAP
/u/username/sqllib/map/dcs1new.map,D
,D
,,INTERRUPT_ENABLED
NOMAP,D,INTERRUPT_ENABLED,,,SYSPLEX,LOCALDATE=\"\"YYMMDD\"\",,
```

On OS/2 or Windows NT:

```
NOMAP
d:\sqllib\map\dcs1new.map,D
,,INTERRUPT_ENABLED
NOMAP,D,INTERRUPT_ENABLED,,,SYSPLEX,LOCALDATE=\"\"YYMMDD\"\",,
```

Alternatively you can accept the defaults by not specifying a parameter string.

Note: Because of the need to specify two pairs of double quotes when specifying the LOCALDATE mask in the parameter string, you must use the operating system escape character ″\″ (backslash), for example:

```
db2 catalog dcs db x as y parms \",,,,,,LOCALDATE=\"\"YYMMDD\"\"\"
```

This results in the following DCS directory entry:

```
DCS 1 entry:

 Local database name           = X
 Target database name          = Y
 Application requestor name     =
 DCS parameters                = ,,,,,,LOCALDATE="YYMMDD"
 Comment                       =
 DCS directory release level   = 0x0100
```

System Database Directory

You can specify the following information in the system database directory:

Database name
>The same value that you wrote in the DCS Directory Parameters table.

Database alias
>An alias for the host or AS/400 database server. This name will be used by any application program that accesses the database. By default, the value that you specify for Database name is used.
>
>Format: 1–8 single-byte alphanumeric characters, including the number sign (#), at sign (@), dollar sign ($), and underscore (_). It cannot begin with an underscore or a number.

Node name
>The same value that you wrote in the Node Directory Parameters table.

Authentication
>Specifies where the validation of the user's name and password will be done. The valid options are: SERVER, SERVER_ENCRYPT, CLIENT, DCE, DCS, and DCS_ENCRYPT. For more information, see "Chapter 7. Security" on page 63.

Defining Multiple Entries for the Same Database

For each database, you must define at least one entry in each of the three directories (node directory, DCS directory, and system database directory). In some cases, you might want to define more than one entry for the database.

For example, you might want to turn off SQLCODE mapping for applications that were ported from the host or AS/400 database server but accept the default mapping for applications that were developed for the client/server environment. You would do this as follows:

- Define one entry in the node directory.
- Define two entries in the DCS directory, with different database names. For one entry, specify NOMAP in the parameter string.
- Define two entries in the system database directory, with different database aliases and the two database names that you specified in the DCS directory.

Both aliases access the same database, one with SQLCODE mapping and the other without SQLCODE mapping.

Updating the Directories

You can use the command line processor CATALOG command on any DB2
Connect system; the "Add Database SmartGuide" of the Client Configuration
Assistant on OS/2, Windows 95, or Windows NT; or the Data Sources Setup
tool on Windows 3.1 or Windows 3.11 for Workgroups. If you have the DB2
Software Developer Kit (SDK), you can also create an application program to
catalog entries. For information about APIs, see the *Administrative API
Reference* and the *Command Reference*. Note that to catalog a database, you
must have *sysadm* authority.

To update the directories using the command line processor, do the following:

1. Use one of the following commands to update the node directory:
   - For a node having an APPC connection:

     ```
     db2 CATALOG APPC NODE nodename
     REMOTE symbolic_destination_name
     SECURITY security_type
     ```

     For example:

     ```
     db2 CATALOG APPC NODE DB2NODE REMOTE DB2CPIC SECURITY PROGRAM
     ```

   - For a DB2 Universal Database for OS/390 Version 5.1 or a DB2
     Universal Database for AS/400 Version 4.2 database having a TCP/IP
     connection:

     ```
     db2 CATALOG TCPIP NODE nodename
     REMOTE hostname or IP address
     SERVER service_name or port_number
     SECURITY security_type
     ```

     For example:

     ```
     db2 CATALOG TCPIP NODE MVSIPNOD REMOTE MVSHOST SERVER DB2INSTC
     ```

     The default DRDA port number for TCP/IP connections is 446.

2. Use the following command to update the DCS directory:

   ```
   db2 CATALOG DCS DATABASE database_name
   AS target_database_name
   [AR application_requester]
   [PARMS "parameter string"]
   ```

   For example:

   ```
   db2 CATALOG DCS DATABASE DB2DB AS NEW_YORK3
   ```

   Or, for OS/2 or Windows NT:

   ```
   db2 CATALOG DCS DATABASE DB2DB AS NEW_YORK3 PARMS "NOMAP,D"
   ```

   Or, for AIX:

```
db2 CATALOG DCS DATABASE DB2DB AS NEW_YORK3 PARMS '"NOMAP,D"'
```

> **Note:** See "Specifying the Parameter String" on page 28 for information about using the operating system escape character when specifying the LOCALDATE mask in the parameter string.

3. Use the following command to update the system database directory:
   ```
   db2 CATALOG DATABASE database_name
   AS alias
   AT NODE nodename
   AUTHENTICATION authentication_type
   ```

   For example:
   ```
   db2 CATALOG DATABASE DB2DB AS NYC3 AT NODE DB2NODE AUTHENTICATION DCS
   ```

If you have remote clients, you must also update directories on each remote client. For more information, see *DB2 Connect Enterprise Edition for OS/2 and Windows NT Quick Beginnings*.

Chapter 4. Binding Applications and Utilities

Application programs developed using embedded SQL must be bound to each database with which they will operate. On platforms where these functions are available, you can do this using the Command Center and the Client Configuration Assistant.

Binding should be performed once per application, for each database. During the bind process, database access plans are stored for each SQL statement that will be executed. These access plans are supplied by application developers and are contained in *bind files*, which are created during precompilation. Binding is simply a process of processing these bind files by a host or AS/400 database server. For more information about binding, refer to the *Application Development Guide*.

Because several of the utilities supplied with DB2 Connect are developed using embedded SQL, they must be bound to a host or AS/400 database server before they can be used with that system. If you do not use the DB2 Connect utilities and interfaces listed in Table 2 on page 36, you do not have to bind them to each of your host or AS/400 database servers. The lists of bind files required by these utilities are contained in the following files:

ddcsmvs.lst
For MVS or OS/390

ddcsvse.lst
For VSE

ddcsvm.lst
For VM

ddcs400.lst
For OS/400

Binding one of these lists of files to a database will bind each individual utility to that database.

If DB2 Connect Enterprise Edition (formerly DDCS Multi-user Gateway) is installed, the DB2 Connect utilities must be bound to each host or AS/400 database server; once from each type of client platform, before they can be used with that system.

For example, if you have 10 OS/2 clients, 10 Windows clients, and 10 AIX clients connecting to DB2 Universal Database for OS/390 via a DB2 Connect Enterprise Edition for Window NT gateway, do the following:

1. Bind ddcsmvs.lst from one of the Windows clients.
2. Bind ddcsmvs.lst from one of the OS/2 clients.
3. Bind ddcsmvs.lst from one of the AIX clients.
4. Bind ddcsmvs.lst from the DB2 Connect gateway workstation.

Note: This assumes all the clients are at the same service level. If they are not then, in addition, you may need to bind from each client of a particular service level. Refer to "Appendix F. Binding Utilities for Back-Level Clients" on page 155 if you have clients prior to DB2 Version 2.1.

In addition to DB2 Connect utilities, any other applications that use embedded SQL must also be bound to each database that you want them to work with. An application that is not bound will usually produce an SQL0805N error message when executed. You might want to create an additional bind list file for all of your applications that need to be bound.

For each host or AS/400 database server that you are binding to, do the following:

1. Make sure that you have sufficient authority for your host or AS/400 database server management system:

   MVS or OS/390

   The authorizations required are:

   SYSADM or

   SYSCTRL or

   BINDADD *and* CREATE IN COLLECTION NULLID

   Note: The BINDADD and the CREATE IN COLLECTION NULLID privileges provide sufficient authority **only** when the packages do not already exist. For example, if you are creating them for the first time.

   If the packages already exist, and you are binding them again, then the authority required to complete the task(s) depends on who did the original bind.

   A If you did the original bind and you are doing the bind again, then having any of the above listed authorities will allow you to complete the bind.

   B If your original bind was done by someone else and you are doing the second bind, then you will require either the SYSADM **or** the SYSCTRL authorities to complete the bind. Having just the BINDADD and the CREATE IN

COLLECTION NULLID authorities will not allow you to complete the bind. It is still possible to create a package if you do not have either SYSADM or SYSCTRL privileges. In this situation you would need the BIND privilege on each of the existing packages that you intend to replace.

VSE or VM

The authorization required is DBA authority. If you want to use the GRANT option on the bind command (to avoid granting access to each DB2 Connect package individually), the NULLID user ID must have the authority to grant authority to other users on the following tables:

system.syscatalog

system.syscolumns

system.sysindexes

system.systabauth

system.syskeycols

system.syssynonyms

system.syskeys

system.syscolauth

On the VSE or VM system, you can issue:

```
grant select on table to nullid with grant option
```

OS/400

*CHANGE authority or higher on the NULLID collection.

2. Issue commands similar to the following:

```
db2 connect to DBALIAS user USERID using PASSWORD
db2 bind path@ddcsmvs.lst blocking all
    sqlerror continue messages ddcsmvs.msg grant public
db2 connect reset
```

Where *DBALIAS*, *USERID*, and *PASSWORD* apply to the host or AS/400 database server, ddcsmvs.lst is the bind list file for MVS, and *path* is the location of the bind list file.

For example *drive*:\sqllib\bnd\ applies to all Intel operating systems, and *INSTHOME*/sqllib/bnd/ applies to all UNIX operating systems, where *drive* is the logical drive where DB2 Connect was installed and *INSTHOME* is the home directory of the DB2 Connect instance.

You can use the `grant` option of the `bind` command to grant EXECUTE privilege to PUBLIC or to a specified user name or group ID. If you do not use the `grant` option of the `bind` command, you must GRANT EXECUTE (RUN) individually.

To find out the package names for the bind files, enter the following command:

```
ddcspkgn @bindfile.lst
```

For example:

```
ddcspkgn @ddcsmvs.lst
```

might yield the following output:

```
Bind File                       Package Name
-----------------------------   ------------------------------
f:\sqllib\bnd\db2ajgrt.bnd      SQLAB6D3
```

For your reference, Table 2 shows the bind files and package names that are used by different components of DB2 Connect. In some cases, different bind files and packages are used on different operating systems.

Table 2. Bind Files and Packages

Component	Bind file	Package	MVS or OS/390	VSE	VM	OS/400
Binder (used by the GRANT bind option)	db2ajgrt.bnd	sqlab*xxx*	yes	yes	yes	yes
DB2 Call Level Interface						
Isolation level CS	db2clics.bnd	sqll1*xxx*	yes	yes	yes	yes
Isolation level RR	db2clirr.bnd	sqll2*xxx*	yes	yes	yes	yes
Isolation level UR	db2cliur.bnd	sqll3*xxx*	yes	yes	yes	yes
Isolation level RS	db2clirs.bnd	sqll4*xxx*	yes	yes	yes	yes
Isolation level NC	db2clinc.bnd	sqll5*xxx*	no	no	no	yes
Using MVS table names	db2clims.bnd	sqll7*xxx*	yes	no	no	no
Using OS/400 table names (OS/400 3.1 or later)	db2clias.bnd	sqlla*xxx*	no	no	no	yes
Using VSE/VM table names	db2clivm.bnd	sqll8*xxx*	no	yes	yes	no
Command Line Processor						
Isolation level CS	db2clpcs.bnd	sqlc2*xxx*	yes	yes	yes	yes
Isolation level RR	db2clprr.bnd	sqlc3*xxx*	yes	yes	yes	yes

Table 2. Bind Files and Packages (continued)

Component	Bind file	Package	MVS or OS/390	VSE	VM	OS/400
Isolation level UR	db2clpur.bnd	sqlc4*xxx*	yes	yes	yes	yes
Isolation level RS	db2clprs.bnd	sqlc5*xxx*	yes	yes	yes	yes
Isolation level NC	db2clpnc.bnd	sqlc6*xxx*	no	no	no	yes
REXX						
Isolation level CS	db2arxcs.bnd	sqla1*xxx*	yes	yes	yes	yes
Isolation level RR	db2arxrr.bnd	sqla2*xxx*	yes	yes	yes	yes
Isolation level UR	db2arxur.bnd	sqla3*xxx*	yes	yes	yes	yes
Isolation level RS	db2arxrs.bnd	sqla4*xxx*	yes	yes	yes	yes
Isolation level NC	db2arxnc.bnd	sqla5*xxx*	no	no	no	yes
Utilities						
Export	db2uexpm.bnd	sqlub*xxx*	yes	yes	yes	yes
Import	db2uimpm.bnd	sqluf*xxx*	yes	yes	yes	yes

where *xxx* depends on the type of client platform and what service you
have applied.

The last 3 characters of a package name represent:

```
- The last digit of the year the product was released
- The platform of the product
    where:  A = AIX                D = OS/2
            W = Windows 3.1        N = Windows NT/95/98
            M = Macintosh          G = Silicon Graphics
            H = HP-UX              U = Solaris
            X = Sinix Nixdorf      S = SCO
            P = Power PC           X = Parallel edition
```

To determine these values for DB2 Connect Version 6 execute the *ddcspkgn*
utility, for example:

```
ddcspkgn @ddcsmvs.lst
```

Optionally, this utility can be used to determine the package name of
individual bind files, for example:

```
ddcspkgn bindfile.bnd
```

If your DB2 for MVS/ESA system has the fix for APAR PN60988 installed
(or if it is a later release than Version 3 Release 1), you can also add the
bind files for isolation level NC to the *ddcsmvs.lst file*.

Bind options are discussed in detail in the *Command Reference.*

Notes:

a. Using the bind option `sqlerror continue` is required; however, this option is automatically specified for you when you bind applications using the DB2 tools or the command line processor. Specifying this option turns bind errors into warnings, so that binding a file containing errors can still result in the creation of a package. In turn, this allows one bind file to be used against multiple servers even when a particular server implementation may flag the SQL syntax of another to be invalid. For this reason, binding any of the list files `ddcsxxx.lst` against any particular host or AS/400 database server should be expected to produce some warnings. For example, when binding against DB2 for VM, numerous warning messages may result since DB2 for VM does not permit cursors to be declared as `"WITH HOLD"`.

b. If you are connecting to a DB2 Universal Database database through DB2 Connect, use the bind list `db2ubind.lst` and do not specify `sqlerror continue`, which is only valid when connecting to a host or AS/400 database server. Also, to connect to a DB2 Universal Database database, we recommend that you use the DB2 clients provided with DB2 and not DB2 Connect.

3. Use similar statements to bind each application or list of applications.

4. If you have remote clients from a previous release of DB2, you may need to bind the utilities on these clients to DB2 Connect. See "Appendix F. Binding Utilities for Back-Level Clients" on page 155 for more information.

The BIND Command

The DB2 **bind** command binds an application to a specific database. If you precompile and bind in separate operations, the options that you specify in the bind will override the options that you specified in the precompile step.

The *Command Reference* describes the syntax of the BIND command that you must use when binding an application to a host or AS/400 database server through DB2 Connect. Be sure to check that you are referring to the DRDA-specific description.

Note: Some parameters of the BIND command may not be supported by your host or AS/400 database server. For more information, see the documentation supplied with your host or AS/400 database server RDBMS.

Rebinding

After you bind your application (and create the package on the host or AS/400 database server), you may find that you need to recreate that package. It is possible to do this without having the original bind file, by using the command line processor **REBIND PACKAGE** command or corresponding API.

The benefits of using this command are:
- You can take advantage of changes in the system by re-optimizing and building new package sections without having the original bind file.
- You can recreate packages that have been made inoperative or invalid.
- You can recreate packages that have been invalidated by migration.
- You can improve performance by using an explicit rebind rather than using an implicit rebind or a bind.
- You can change characteristics. For example, with DB2 Universal Database for OS/390, you can change the qualifier of unqualified tables for testing or migration purposes.

If you want to modify a program, the bind options, or any owner information, you should use the **BIND** command. Also, if the package does not exist in the database, or if you want to see all bind errors (not only the first detected error), you should use the **BIND** command.

To run this command, you need the level of authority required by your host or AS/400 database server. If you are not connected to a database, the command will cause an implicit connect to be done to the default database (if you have connect privileges).

The syntax of the command line processor command is described in the *Command Reference.*

To find out the package name for a bind file, enter the following command:

```
ddcspkgn bindfile.bnd
```

Chapter 5. DB2 Connect Administration Utilities

This chapter describes utilities that help you perform administration tasks. It contains the following sections:

- "Command Line Processor"
- "Using Import and Export Utilities"
- "Trace Utility (ddcstrc)" on page 43.

Command Line Processor

The command line processor lets you issue SQL statements against a host or AS/400 database server, preceded by **db2**. For differences between host and AS/400 SQL and DB2 Connect SQL, see "Host or AS/400 Server SQL Statements Supported by DB2 Connect" on page 83 and "Host or AS/400 Server SQL Statements Rejected by DB2 Connect" on page 84.

You can also enter **db2** with no SQL statement. This puts you into interactive input mode, where you can type SQL commands without prefixing them by **db2**. Enter **quit** to leave CLP interactive input mode,or enter **terminate** to end the command line processor and the database connection.

Before you can use the command line processor, it must be bound, as described in "Chapter 4. Binding Applications and Utilities" on page 33.

Note: When using the command line processor in UNIX-based systems without being in interactive input mode, you must put double quotes around special characters (such as * and ?) when issuing SQL statements.

For more information on SQL commands, see the *SQL Reference*.

Using Import and Export Utilities

The import and export utilities let you move data from a host or AS/400 database to a file on the DB2 Connect workstation or vice versa. You can then use this data with any other application or RDBMS that supports this import/export format. For example, you can export data from DB2 for MVS/ESA into a delimited ASCII file and later import it into a DB2 for OS/2 database.

You can perform export and import functions from a database client or from the DB2 Connect workstation.

Notes:

1. The data to be imported or exported must comply with the size and data type restrictions of both databases.

2. To improve import performance, you can use compound SQL. Specify COMPOUND=*number* in the import API or the CLP *filetype-mod* string parameter to group the specified number of SQL statements into a block. This may reduce network overhead and improve response time.

3. For information on the syntax of the import and export utilities from the command line processor, see the *Command Reference* manual.

Moving Data from a Workstation to a Host or AS/400 Database Server

To export to a host or AS/400 database server:

1. Export the rows of information from the DB2 table into a PC/IXF file.

2. If the host or AS/400 database server database does not contain a table having attributes compatible with the information to be imported into it, create a compatible table.

3. Using the IMPORT command with the INSERT option, import the PC/IXF file to a table in the DRDA server database.

Moving Data from a host or AS/400 database server to a Workstation

To import data from a host or AS/400 database server:

1. Export the rows of information from the server database table to a PC/IXF file.

2. Use the PC/IXF file for importing to a DB2 table.

Restrictions

With the DB2 Connect program, import or export operations must meet the following conditions:

- The file type must be PC/IXF.
- Index definitions are not stored on export or used on import.
- A table with attributes that are compatible with those of the data must exist before you can import to it. Importing through the DB2 Connect program cannot create a table because INSERT is the only supported option.
- A commit count interval must not be specified with import.

If these conditions are violated, the operation will fail and an error message will be generated.

Mixed Single-Byte and Double-Byte Data

If you import and export mixed data (columns containing both single-byte and double-byte data), consider the following:

- On systems that store data in EBCDIC (MVS, OS/390, OS/400, VM, and VSE), shift-out and shift-in characters mark the start and end of double-byte data. When you define column lengths for your database tables, be sure to allow enough room for these characters.
- Variable-length character columns are recommended unless the data in a column has a consistent pattern. If it does, fixed length is acceptable.

Replacement for SQLQMF Utility

The function of the SQLQMF utility with DDCS for OS/2 has been replaced by the DB2 Connect Import/Export functions. The advantages are:

- No need for QMF on the host
- No need to logon to the host (a TSO id is still required on DB2 for MVS/ESA or DB2 Universal Database for OS/390)
- Supports DB2 for MVS, DB2 for OS/390, DB2 for OS/400, and DB2 for VM and VSE
- Good performance achieved by using compound SQL
- Supports several file formats, in addition to ASCII
- Can be run from a client machine with no SNA connectivity.

Refer to the *Command Reference* for further information on using these commands.

Trace Utility (ddcstrc)

The **ddcstrc** utility provides a record of the data interchanged between the DB2 Connect workstation (on behalf of the database client) and the host or AS/400 database server management system.

As a database administrator (or application developer), you may find it useful to understand how this flow of data works, because this knowledge can help you determine the origin of a particular problem. For example, say you issue a CONNECT TO database statement for a host or AS/400 database server, but the command fails and you receive an unsuccessful return code. If you understand exactly what information was conveyed to the host or AS/400 database server management system, you may be able to determine the cause of the failure even if the return code information is general. Many failures are caused by simple user errors.

Output from `ddcstrc` lists the data streams exchanged between the DB2 Connect workstation and the host or AS/400 database server management system. Data sent to the host or AS/400 database server is labeled SEND BUFFER and data received from the host or AS/400 database server is labeled RECEIVE BUFFER. If a receive buffer contains SQLCA information, it will be followed by a formatted interpretation of this data and labeled SQLCA. The SQLCODE field of an SQLCA is the *unmapped* value as returned by the host or AS/400 database server. (For more on mapping, see "Chapter 9. SQLCODE Mapping" on page 89.) The send and receive buffers are arranged from the oldest to the most recent within the file. Each buffer has:

- The process ID
- A SEND BUFFER, RECEIVE BUFFER, or SQLCA label. The first DDM command or object in a buffer is labeled DSS TYPE.

The remaining data in send and receive buffers is divided into five columns, consisting of:

- A byte count.
- Columns 2 and 3 represent the DRDA data stream exchanged between the two systems, in ASCII or EBCDIC.
- An ASCII representation of columns 2 and 3.
- An EBCDIC representation of columns 2 and 3.

For more information about DDM, see:
- *DB2 for OS/390 Reference for Remote DRDA Requesters and Servers*
- *Distributed Relational Database Reference*
- *Distributed Data Management Architecture Level 3: Reference*

Trace Syntax

This command is invoked from the operating system command prompt with the following syntax

```
>>──┬───────┬──ddcstrc──on──┬──┬─-i─┬──┬──┬──────────────┬───────────>
    └─path─┘               ▲  ├─-r─┤     └─-l──=──length─┘
                           │  ├─-s─┤
                           └──┴─-c─┘

>──off──┬──────────────────┬──┬──────────────┬──────────────────────><
        └─-t──=──tracefile─┘  └─-p──=──pid────┘
```

Figure 3. Syntax of the ddcstrc Command

Note: The syntax of this command may vary slightly depending on the operating system that you are using. For example, / may be used in place of - for the OS/2 operating system.

Trace Parameters

on   Turns on DB2 Connect tracing of DRDA flows with the host or AS/400 database server.

off   Turns off DB2 Connect tracing of DRDA flows with the host or AS/400 database server.

-**i**   Timestamps will be included in the trace information.

-**r**   Traces DRDA data streams received from the host or AS/400 server system.

-**s**   Traces DRDA data streams sent to the host or AS/400 database server.

-**c**   Traces the SQLCA received from the host or AS/400 database server.

        The default is -r, -s, and -c.

-**l**=*length*
        Specifies the size of the buffer used to store the trace information. The default is 1M, and the minimum is 64K.

-**t**=*tracefile*
        Specifies the destination for the trace; *tracefile* may be the name of a file or a standard device. If a file name is specified without a complete path, the current path is used for the missing parts. The default file name is `ddcstrc.dmp`.

-**p**=*pid* Traces events only for this process. If -p is not specified, all processes for the user's instance are written to the output file.

        Note: For a remote client, the *pid* can be found in the Agent ID field returned by the database system monitor.
        For more information, see "Chapter 6. Database System Monitor for DB2 Connect" on page 55.

Trace Output

ddcstrc writes the following information to *tracefile*:

-r
- Type of DRDA reply/object
- Receive buffer

-s
- Type of DRDA request
- Send buffer

-c
- SQLCA

CPI-C error information
- Receive function return code
- Severity
- Protocol used
- API used
- Function
- CPI-C return code
- Error number
- Internal return code.

SNA error information
- Receive function return code
- Severity
- Protocol used
- Function
- Partner LU name
- Error number.

TCP/IP error information
- Receive function return code
- Severity
- Protocol used
- API used
- Function
- Error number.

Note:

1. A value of zero for the exit code indicates that the command completed successfully, and a non-zero value indicates that it did not.
2. The fields returned vary based on the API used. The SNA API is used only for 2PC SPM connections.
3. The fields returned vary based on the platform on which DB2 Connect is running, even for the same API.
4. If ddcstrc sends the output to a file that already exists, the old file will be erased unless the permissions on the file do not allow it to be erased.

Analyzing the Trace Output File

The following pages show example output illustrating some DRDA data streams exchanged between DB2 Connect workstations and a host or AS/400 database server. From the user's viewpoint, a `CONNECT TO` database command has been issued using the command line processor.

Figure 4 on page 48 uses DB2 Connect Enterprise Edition Version 6 and DB2 Universal Database for OS/390 Version 5.1 over an APPC connection.

Figure 5 on page 50 uses DB2 Connect Enterprise Edition Version 6 and DB2 Universal Database for OS/390 Version 5.1 over a TCP/IP connection.

```
1       DB2 fnc_data       gateway_drda_ar      sqljcsend (1.35.10.80)
        pid 95212; tid 537115484; node 0; cpid 0; sec 0; nsec 0; tpoint 177

        SEND BUFFER:  EXCSAT RQSDSS         (ASCII)          (EBCDIC)
          0 1 2 3 4 5 6 7  8 9 A B C D E F  0123456789ABCDEF  0123456789ABCDEF
  0000  006AD04100010064 10410020115E8482  .j.A...d.A. .^..  .|}..........;db
  0010  F282974040404040 4040404040404040  ...@@@@@@@@@@@@  2bp
  0020  4040F0F0F0F1F7F3 C5C3000C116DA685  @@...........m..    000173EC..._we
  0030  81A2859340400013 115AC4C2F240C396  ....@@...Z...@..  asel   ...]DB2 Co
  0040  95958583A340F54B F200141404140300  .....@.K........  nnect 5.2.......
  0050  0414440003240700 05240F0003000D11  ..D..$...$......  ................
  0060  47D8C4C2F261F6F0 F0F00085D0010002  G....a..........  .QDB2/6000.e}...
  0070  007F200100162110 E2C1D56DC6D9C1D5  .. ...!....m....  ."......SAN_FRAN
  0080  C3C9E2C3D6404040 40400006210F2407  .....@@@@@..!.$.  CISCO     ......
  0090  000D002FD8E3C4E2 D8D3C1E2C3000C11  .../............  ....QTDSQLASC...
  00A0  2EE2D8D3F0F5F0F2 F0003C210437E2D8  ..........<!.7..  .SQL05020.....SQ
  00B0  D3F0F5F0F1F0C1C9 E740404040404040  .........@@@@@@@  L05010AIX
  00C0  4040404040404040 8482F28297404040  @@@@@@@@.....@@@        db2bp
  00D0  4040404040404040 40404040848483A2  @@@@@@@@@@@@....           ddcs
  00E0  A4A2F14000000A00 350006119C0352    ...@....5.....R  us1 ...........

2       DB2 fnc_data       gateway_drda_ar      sqljcrecv (1.35.10.81)
        pid 95212; tid 537115484; node 0; cpid 0; sec 0; nsec 0; tpoint 178

        RECEIVE BUFFER:  EXCSATRD OBJDSS      (ASCII)          (EBCDIC)
          0 1 2 3 4 5 6 7  8 9 A B C D E F  0123456789ABCDEF  0123456789ABCDEF
  0000  0056D04300010050 14430010115EC4C2  .V.C...P.C...^..  ..}....&.....;DB
  0010  C1C1F0F7F6C3F4C5 F0F8001414041403  ................  AA076C4E08......
  0020  0004144400032407 0005240F00030008  ...D..$...$.....  ................
  0030  1147D8C4C2F20014 116DE2C1D56DC6D9  .G.......m...m..  ..QDB2..._SAN_FR
  0040  C1D5C3C9E2C3D640 4040000C115AC4E2  .......@@@...Z..  ANCISCO   ...]DS
  0050  D5F0F5F0F1F1007B D002000200752201  .......{.....u".  N05011.#}.......
  0060  000611490000000D 002FD8E3C4E2D8D3  ...I...../......  ..........QTDSQL
  0070  F3F7F0000C112EC4 E2D5F0F5F0F1F100  ...............  370....DSN05011.
  0080  0A000350006119C01 F4001B2135C3C1C9  ..5........!5...  ........4....CAI
  0090  C2D4D6D4D34BD6D4 E7E3F4C8F0C1B043  .....K.........C  BMOML.OMXT4H0A..
  00A0  CBC735AB002D244E 0006244C00010023  ..5..-$N..$L...#  .G.....+...<....
  00B0  244D0006244FFFFF 001911E9C3C1C9C2  $M..$O..........  .(...!.....ZCAIB
  00C0  D4D6D4D34BE2C6D3 E44040404007F6C4  ....K....@@@...  MOML.SFLU   .6D
  00D0  C2                                 .               B
```

Figure 4. Example of Trace Output (APPC connection) (Part 1 of 2)

```
3        DB2 fnc_data     gateway_drda_ar     sqljcsend (1.35.10.80)
         pid 95212; tid 537115484; node 0; cpid 0; sec 0; nsec 0; tpoint 177

         SEND BUFFER:   RDBCMM  RQSDSS       (ASCII)           (EBCDIC)
          0 1 2 3 4 5 6 7  8 9 A B C D E F   0123456789ABCDEF  0123456789ABCDEF
   0000  000AD00100010004 200E              ........ .        ..}.......

4        DB2 fnc_data     gateway_drda_ar     sqljcrecv (1.35.10.81)
         pid 95212; tid 537115484; node 0; cpid 0; sec 0; nsec 0; tpoint 178

         RECEIVE BUFFER:  ENDUOWRM RPYDSS    (ASCII)           (EBCDIC)
          0 1 2 3 4 5 6 7  8 9 A B C D E F   0123456789ABCDEF  0123456789ABCDEF
   0000  002BD05200010025 220C000611490004  .+.R...%"....I..  ..}.............
   0010  00162110E2C1D56D C6D9C1D5C3C9E2C3  ..!....m........  ....SAN_FRANCISC
   0020  D640404040400005 211501000BD00300  .@@@@@..!.......  O      .......}..
   0030  0100052408FF                       ...$..            ......

5        DB2 fnc_data     gateway_drda_ar     sqljmsca (1.35.10.108)
         pid 95212; tid 537115484; node 0; cpid 0; sec 0; nsec 0; tpoint 179
         SQLCA

         SQLCAID:  SQLCA
         SQLCABC:  136
         SQLCODE:  0
         SQLERRML: 0
         SQLERRMC:
         SQLERRP:  DSN
         SQLERRD[0->5]: 00000000, 00000000, 00000000, 00000000, 00000000, 00000000
         SQLWARN(0->A):  ,  ,  ,  ,  ,  ,  ,  ,  ,  ,
         SQLSTATE: 00000
```

Figure 4. Example of Trace Output (APPC connection) (Part 2 of 2)

```
1      DB2 fnc_data      gateway_drda_ar      sqljcsend (1.35.10.80)
       pid 80286; tid 537125164; node 0; cpid 0; sec 0; nsec 0; tpoint 177

       SEND BUFFER:  EXCSAT RQSDSS           (ASCII)          (EBCDIC)
         0 1 2 3 4 5 6 7  8 9 A B C D E F    0123456789ABCDEF  0123456789ABCDEF
  0000  006ED04100010068 10410020115E8482    .n.A...h.A. .ˆ..   .>}..........;db
  0010  F282974040404040 4040404040404040    ...@@@@@@@@@@@@    2bp
  0020  4040F0F0F0F1F3F9 F9C5000C116DA685    @@..........m..      0001399E..._we
  0030  81A2859340400013 115AC4C2F240C396    ....@@...Z...@..   asel   ...]DB2 Co
  0040  95958583A340F54B F200181404140300    .....@.K........   nnect 5.2.......
  0050  0514740005240700 05240F0003144000    ..t..$...$....@.   .............. .
  0060  05000D1147D8C4C2 F261F6F0F0F00010    ....G....a......   .....QDB2/6000..
  0070  D0410002000A106D 000611A20003003C    .A.....m.......<   }......._...s....
  0080  D04100030036106E 000611A200030016    .A...6.n........   }......>...s....
  0090  2110E2C1D56DC6D9 C1D5C3C9E2C3D640    !....m.........@   ..SAN_FRANCISCO
  00A0  40404040000C11A1 9781A2A2A6969984    @@@@...........   ....password
  00B0  000A11A0A4A28599 8984009CD0010004    ..............   ....userid..}...
  00C0  0096200100162110 E2C1D56DC6D9C1D5    .. ...!....m...   .o......SAN_FRAN
  00D0  C3C9E2C3D6404040 40400006210F2407    .....@@@@@..!.$.   CISCO     ......
  00E0  000D002FD8E3C4E2 D8D3C1E2C3000C11    .../...........   ....QTDSQLASC...
  00F0  2EE2D8D3F0F5F0F2 F0003C210437E2D8    ..........<!.7..   .SQL05020.....SQ
  0100  D3F0F5F0F1F0C1C9 E740404040404040    .........@@@@@@@   L05010AIX
  0110  4040404040404040 8482F28297404040    @@@@@@@@.....@@@         db2bp
  0120  4040404040404040 40404040A4A28599    @@@@@@@@@@@@....            user
  0130  8984404000000172 135C7F9F1F5F0C6C4    ..@@...!5.......   id  .....G9150FD
  0140  C24BC7C1F1C51ACA 06195122000A0035    .K........Q"...5   B.GA1E..........
  0150  0006119C0352                         .....R            ......

2      DB2 fnc_data      gateway_drda_ar      sqljcrecv (1.35.10.81)
       pid 80286; tid 537125164; node 0; cpid 0; sec 0; nsec 0; tpoint 178

       RECEIVE BUFFER:  EXCSATRD OBJDSS       (ASCII)          (EBCDIC)
         0 1 2 3 4 5 6 7  8 9 A B C D E F    0123456789ABCDEF  0123456789ABCDEF
  0000  005AD04300010054 14430010115EC4C2    .Z.C...T.C...ˆ..   .]}..........;DB
  0010  C1C1F0F7F6C3F4C5 F0F80018140414,03   ..............   AA076C4E08......
  0020  00051474000052407 0005240F00031440    ...t..$...$....@   ..............
  0030  000500081147D8C4 C2F200014116DE2C1    .....G.......m..   ......QDB2..._SA
  0040  D56DC6D9C1D5C3C9 E2C3D6404040000C    .m.........@@@..   N_FRANCISCO   ..
  0050  115AC4E2D5F0F5F0 F1F10010D0430002    .Z..........C..   .]DSN05011..}...
  0060  000A14AC000611A2 00030015D0420003    ..............B..   .......s....}...
  0070  000F121900061149 0000000511A40000    .......I........   .............u..
  0080  51D0020004004B22 0100061149000000    Q.....K"....I...   .}..............
  0090  0D002FD8E3C4E2D8 D3F3F7F0000C112E    ../...........   ...QTDSQL370....
  00A0  C4E2D5F0F5F0F1F1 000A00350006119C    ...........5....   DSN05011........
  00B0  01F4001E244E0006 244C00010014244D    ....$N..$L....$M   .4...+...<.....(
  00C0  0006244FFFFFF000A 11E80915155C01BE    ..$O.........\..   ...!.....Y...*..
```

Figure 5. Example of Trace Output (TCP/IP connection) (Part 1 of 2)

```
3       DB2 fnc_data      gateway_drda_ar      sqljcsend (1.35.10.80)
        pid 80286; tid 537125164; node 0; cpid 0; sec 0; nsec 0; tpoint 177

        SEND BUFFER: RDBCMM  RQSDSS         (ASCII)          (EBCDIC)
         0 1 2 3 4 5 6 7  8 9 A B C D E F   0123456789ABCDEF  0123456789ABCDEF
  0000  000AD00010010004 200E              ........ .        ..}.......

4       DB2 fnc_data      gateway_drda_ar      sqljcrecv (1.35.10.81)
        pid 80286; tid 537125164; node 0; cpid 0; sec 0; nsec 0; tpoint 178

        RECEIVE BUFFER:  ENDUOWRM RPYDSS    (ASCII)          (EBCDIC)
         0 1 2 3 4 5 6 7  8 9 A B C D E F   0123456789ABCDEF  0123456789ABCDEF
  0000  002BD05200010025 220C000611490004  .+.R...%"....I..  ..}.............
  0010  00162110E2C1D56D C6D9C1D5C3C9E2C3  ..!....m........  ....SAN_FRANCISC
  0020  D640404040400005 211501000BD00300  .@@@@@..!.......  O      .......}..
  0030  0100052408FF                       ...$..            ......

5       DB2 fnc_data      gateway_drda_ar      sqljmsca (1.35.10.108)
        pid 80286; tid 537125164; node 0; cpid 0; sec 0; nsec 0; tpoint 179
        SQLCA

        SQLCAID:  SQLCA
        SQLCABC:  136
        SQLCODE:  0
        SQLERRML: 0
        SQLERRMC:
        SQLERRP:  DSN
        SQLERRD[0->5]: 00000000, 00000000, 00000000, 00000000, 00000000, 00000000
        SQLWARN(0->A):  ,  ,  ,  ,  ,  ,  ,  ,  ,  ,
        SQLSTATE: 00000
```

Figure 5. Example of Trace Output (TCP/IP connection) (Part 2 of 2)

The following information is captured in the traces:

- The process ID (PID) of the client application
- The RDB_NAME cataloged in the database connection services (DCS) directory
- The DB2 Connect CCSID(s)
- The host or AS/400 database server CCSID(s)
- The host or AS/400 database server management system with which the DB2 Connect system is communicating.

The first buffer contains the Exchange Server Attributes (EXCSAT) and Access RDB (ACCRDB) commands sent to the host or AS/400 database server management system. It sends them as a result of a CONNECT TO database command.

The next buffer contains the reply that DB2 Connect received from the host or AS/400 database server management system. It contains an Exchange Server Attributes Reply Data (EXCSATRD) and an Access RDB Reply Message (ACCRDBRM).

Analyzing EXCSAT and ACCRDB

The EXCSAT command contains the workstation name of the client specified by the Server Name (SRVNAM) object, which is code point X'116D', according to DDM specification. The EXCSAT command is found in the first buffer. Within the EXCSAT command, the values X'116DA68581A28593' (coded in CCSID 500) are translated to *weasel* once the X'116D' is removed.

The EXCSAT command also contains the EXTNAM (External Name) object, which is often placed in diagnostic information on the host or AS/400 database management system. It consists of a 20-byte application ID followed by an 8-byte process ID (or 4-byte process ID and 4-byte thread ID). It is represented by code point X'115E', and in this example its value is db2bp_32 padded with blanks followed by 0000BE5C. On a UNIX-based database client, this value can be correlated with the **ps** command, which returns process status information about active processes to standard output.

The ACCRDB command contains the RDB_NAME in the RDBNAM object, which is code point X'2110'. The ACCRDB command follows the EXCSAT command in the first buffer. Within the ACCRDB command, the values X'2110E2C1D56DC6D9C1D5C3C9E2C3D6' are translated to SAN_FRANCISCO once the X'2110' is removed. This corresponds to the target database name field in the DCS directory.

The accounting string (described in "Implementing Chargeback Accounting" on page 84) has code point X'2104'.

The code set configured for the DB2 Connect workstation is shown by locating the CCSID object CCSIDSBC (CCSID for single-byte characters) with code point X'119C' in the ACCRDB command. In this example, the CCSIDSBC is X'0352', which is 850.

If the additional objects CCSIDDBC (CCSID for double-byte characters) and CCSIDMBC (CCSID for mixed-byte characters), with code points X'119D' and X'119E' respectively, are present, the DB2 Connect workstation is configured for DBCS code page support. Since the example output file does not include the two additional code points, the workstation is not configured for DBCS.

Note: The TCP/IP flows contains two new commands: ACCSEC used to access the security manager and exchange supported security mechanisms, and SECCHK, which contains the authentication tokens

used to authenticate the end user of the connection. ACCSEC and
SECCHK only appear for TCP/IP connections, and they do so between
EXCSAT and ACCRDB.

Analyzing EXCSATRD and ACCRDBRM

CCSID values are also returned from the host or AS/400 database server in
the Access RDB Reply Message (ACCRDBRM) within the second buffer. This
buffer contains the EXCSATRD followed by the ACCRDBRM. The example
output file contains CCSID values for the host or AS/400 database server
system of 500 (X'01F4', SBCS CCSID).

If DB2 Connect does not recognize the code page coming back from the host
or AS/400 database server, SQLCODE -332 will be returned to the user with
the source and target code pages. If the host or AS/400 database server
doesn't recognize the code set sent from DB2 Connect, it will return
VALNSPRM (Parameter Value Not Supported, with DDM code point X'1252'),
which gets translated into SQLCODE -30073 for the user.

The ACCRDBRM also contains the parameter PRDID (Product-specific
Identifier, with code point X'112E'). The value is X'C4E2D5F0F5F0F1F0'. This
hex string corresponds to DSN05010 in EBCDIC. According to standards, DSN
is DB2 for MVS/ESA or DB2 Universal Database for OS/390. The version, 5.1,
is also indicated. ARI is DB2 for VSE & VM, SQL is DB2 Common Server, and
QSQ is DB2 Universal Database for AS/400.

Analyzing Subsequent Buffers

You can analyze subsequent send and receive buffers for additional
information. The third buffer contains a commit. The **commit** command
instructs the host or AS/400 database server management system to commit
the current unit of work. The fourth buffer is received from the host or
AS/400 database server database management system as a result of a commit
or rollback. It contains the End Unit of Work Reply Message (ENDUOWRM),
which indicates that the current unit of work has ended. In this example, it
contains a null SQLCA, indicated by DDM code point X'2408' followed by
X'FF'. A null SQLCA (X'2408FF') indicates success (SQLCODE 0). When a
receive buffer contains an SQLCA (possibly a null SQLCA), ddcstrc will follow
this receive buffer with a formatted interpretation of the SQLCA information.

This SQLCA is the result of
an attempt to delete rows from a nonexistent table.

```
1       DB2 fnc_data       gateway_drda_ar       sqljcrecv (1.35.10.81)
        pid 48732; tid 1; node 0; cpid 0; sec 0; nsec 0; tpoint 178

        RECEIVE BUFFER:  SQLCARD OBJDSS       (ASCII)          (EBCDIC)
         0 1 2 3 4 5 6 7  8 9 A B C D E F  0123456789ABCDEF 0123456789ABCDEF
  0000  0065D0030001005F 240800FFFFFF34F4  .e......_$.....4. ..}....^.......4
  0010  F2F7F0F4C4E2D5E7 D6E3D34000E2C1D5  ...........@.... 2704DSNXOTL .SAN
  0020  6DC6D9C1D5C3C9E2 C3D64040404040FF  m.........@@@@@. _FRANCISCO     .
  0030  FFFE0C0000000000 000000FFFFFFFF00  ................ ................
  0040  0000000000000040 4040404040404040  .......@@@@@@@@@ .......
  0050  40400000000FC4C4 C3E2E4E2F14BD4E8  @@...........K..    ....DDCSUS1.MY
  0060  E3C1C2D3C5                         .....            TABLE

2       DB2 fnc_data       gateway_drda_ar       sqljmsca (1.35.10.108)
        pid 48732; tid 1; node 0; cpid 0; sec 0; nsec 0; tpoint 179
        SQLCA

        SQLCAID:  SQLCA
        SQLCABC:  136
        SQLCODE:  -204
        SQLERRML: 15
        SQLERRMC: DDCSUS1.MYTABLE
        SQLERRP:  DSNXOTL
        SQLERRD[0->5]: FFFFFE0C, 00000000, 00000000, FFFFFFFF, 00000000, 00000000
        SQLWARN(0->A):  ,  ,  ,  ,  ,  ,  ,  ,  ,  ,
        SQLSTATE: 42704
```

Figure 6. Receive Buffer Example

Chapter 6. Database System Monitor for DB2 Connect

This chapter summarizes the capabilities of the DB2 System Monitor for DB2 Connect users. The following System Monitor enhancements are provided in DB2 Connect Version 6:

- A graphical user interface equivalent of the LIST DCS APPLICATIONS CLP command. This capability is provided through the DB2 Command Center. See "LIST DCS APPLICATIONS EXTENDED" on page 59

For more detailed information about the database system monitor and how it works, see the *System Monitor Guide and Reference*.

Monitoring Connections for Remote Clients

You can use the Database System Monitor with DB2 Connect Enterprise Edition to monitor the connections for *remote clients*. Owing to a limitation, at present you *cannot* use the System Monitor to monitor connections for applications running on the DB2 Connect gateway workstation.

For example, when an error occurs at the host or AS/400 system, the system administrator can determine if the problem was on the DB2 Connect workstation. The database system monitor correlates:

- The DRDA correlation token (CRRTKN), for unprotected conversations.
- The logical unit of work identifier (LUWID), for two-phase conversations protected by an SNA Syncpoint Manager (SPM).
- The unit of work id (UOWID), for two-phase connections protected by the DRDA-3 Syncpoint Manager (as used over TCP/IP connections).
- The DB2 Connect connection identifier (the Application ID).

This information shows which DB2 Connect connection caused the problem, which allows the system administrator to force the individual client application from the system without affecting the other clients using the DB2 Connect connection.

Turning on Monitor Switches for DB2 Connect

The system monitor is always active. However, if you want to get more detail in GET SNAPSHOT output, you should turn on the corresponding *monitor switches*. The monitor switches relevant to DB2 Connect are STATEMENT (for statement-level information), and UOW (for transaction-level information).

To change monitor switches, use the `DB2 UPDATE MONITOR SWITCHES` command. Refer to the *Command Reference* for the syntax of this command. Here is an example, which creates DB2 System Monitor statistics for Units of Work (UOW):

```
db2 update monitor switches using uow on
```

Listing the Status of Monitor Switches

To list the status of monitor switches, use the `DB2 GET MONITOR SWITCHES` command.

Listing DCS Application Status

As of DB2 Connect V5.2, the System Monitor provides three forms of the `LIST DCS APPLICATIONS` command, as follows:

- LIST DCS APPLICATIONS
- LIST DCS APPLICATIONS SHOW DETAIL
- LIST DCS APPLICATIONS EXTENDED.

LIST DCS APPLICATIONS

To view the information provided by the monitor at the application level, issue the `DB2 LIST DCS APPLICATIONS` command. It returns the following information for an APPC connection (DB2 Connect Enterprise Edition Version 6 to DB2 Universal Database for OS/390):

```
Auth Id  Application Name      Appl.      Host Application ID
                               Handle
-------- --------------------- ---------- --------------------------------
USERID   db2bp_41              0          CAIBMOML.OMXT4H0A.A79EAA3C6E29
```

It returns the following information for a TCP/IP connection (DB2 Connect Enterprise Edition Version 6 to DB2 Universal Database for OS/390):

```
Auth Id  Application Name      Appl.      Host Application ID
                               Handle
-------- --------------------- ---------- --------------------------------
USERID   db2bp_41              2          0915155C.9704.1517172201BE
```

Auth.Id
    The authorization ID that was used to log on to the host or AS/400 database server. This identifies who is running the application.

Application Name
    The name of the application running at the client as known to DB2

Connect. Only the first 20 bytes after the last path separator are available. The application name is not available for applications running on DB2 for OS/2 Version 1.

Appl. Handle
The agent that is executing on the DB2 Connect workstation. You can use this element to link the database system monitor information to other diagnostic information. (For example, see "Trace Utility (ddcstrc)" on page 43.) The agent ID is also required when using the FORCE USERS command or API.

Host Application ID
One of the following:

- The DRDA correlation token (CRRTKN), for unprotected conversations.
- The logical unit of work identifier (LUWID), for two-phase conversations protected by an SNA Syncpoint Manager (SPM).
- The unit of work id (UOWID), for two-phase connections protected by the DRDA-3 Syncpoint Manager (as used over TCP/IP connections).

This unique identifier is generated when the application connects to the host or AS/400 database server. You can use this element in conjunction with the Application ID to correlate the client and server parts of the application information.

LIST DCS APPLICATIONS SHOW DETAIL

If the `DB2 LIST DCS APPLICATIONS SHOW DETAIL` command format is specified, additional information is shown, including:

```
Auth Id Application Name   Appl.    Client Application Id         Seq# Client
                           Handle                                     DB Alias
------- ------------------ ------- ---------------------------- ---- --------
NEWTON  db2bp              0        09151251.07D3.980925183850        0001 MVSDB

 Client Client   Client   Host Application Id        Seq# Host DB Name
 Node   Release  Codepage
 ------- -------- -------- -------------------------- ---- -------------------
 antman  SQL05020 819      G9151251.G7D4.980925183851 0000 GILROY

Host
Release
--------
DSN05011
```

This report is unformatted, and for this reason you may find the "LIST DCS APPLICATIONS EXTENDED" on page 59 report more useful.

Client Application ID
> Uniquely identifies the application connected to the DB2 Connect workstation. There are different formats for the application ID, which are dependent on the communication protocol between the client and the DB2 Connect workstation. For more information on the formats, see the *Administration Guide*.
>
> This value lets you correlate connections from clients to the DB2 Connect workstation and from the DB2 Connect workstation to the host or AS/400 database server.

Client Sequence no (Seq#)
> The client sequence number is the transaction sequence number. It is used to help correlate a transaction spread over different systems.

Client DB alias
> The alias of the database provided by the application to connect to the database. This element can be used to identify the actual database that the application is accessing. The mapping between this name and the database name could be done by using the database directories at the client node and the database manager server node.

Client NNAME (Node)
> Identifies the node where the client application is executing. The information varies according to the client protocol in use. For example, for a client connected via NetBIOS, this is the value of the NNAME database manager configuration parameter. For a client connected via TCP/IP, this is the host name.

Client Product ID (Client)
> The product and version that is running on the client. The client product IDs will be:

- SQL01010 for Version 1 of DB2 for OS/2
- SQL01011 for Version 1 of UNIX-based DB2 products and Client Application Enablers.
- SQL02010 for Version 2 of DB2 products and Client Application Enablers.
- SQL02020 for Version 2.1.2 of DB2 products and Client Application Enablers.
- SQL05000 for Version 5.0 of DB2 Universal Database and DB2 Connect products and their clients.
- SQL05020 for Version 6 of DB2 Universal Database and DB2 Connect products and their clients.

Code Page ID
> The code page identifier at the node where the monitored application started.

You can use this information to ensure that data conversion is supported between the application code page and the database code page (or for host or AS/400 database server databases, the host or AS/400 database server CCSID).

If the application code page is different from that under which the database system monitor is running, this code page element can help you to manually convert data that was passed from the application and displayed by the database system monitor. For example, you can use it to help translate the Application Name.

Outbound Sequence No
> This represents the outbound sequence number. It is used for correlating transactions on different systems.

Host Database Name
> The real name of the database to which the application is connected. In the DCS directory, this is the *target database name*.

Host Product ID
> The product and version that is running on the server. It is in the form *PPPVVRRM*, where:
>
> **PPP** Identifies the host or AS/400 database server product (for example, DSN for DB2 Universal Database for OS/390, ARI for DB2 for VSE & VM, or QSQ for DB2 Universal Database for AS/400)
>
> **VV** Is a two-digit version number, such as 01
>
> **RR** Is a two-digit release number, such as 01
>
> **M** Is a one-digit modification level.

LIST DCS APPLICATIONS EXTENDED

You can use the LIST DCS APPLICATIONS command with the option EXTENDED in order to generate an Extended Report. The Extended Report lists all the fields that are listed when the SHOW DETAIL option is specifed on the command, plus nine new fields:

- DCS application status
- Status change time
- Client platform
- Client protocol
- Host Coded Character Set Identifier (CCSID).
- Client login ID
- Process ID of client application
- Database alias at the gateway

- DCS database name

While the existing command options list the fields horizontally, one line per application, the new option lists them vertically, one field per line.

Here is the new syntax of the command:

```
LIST DCS APPLICATIONS [SHOW DETAIL | EXTENDED ]
```

And here is sample output from this command, when using the new option EXTENDED:

```
                List of DCS Applications - Extended Report

Client application ID                   = 09151251.0AA7.981015204853
  Sequence number                       = 0001
  Authorization ID                      = NEWTON
  Application name                      = db2bp
  Application handle                    = 1
  Application status                    = waiting for request
  Status change time                    = 10-15-1998 16:50:29.489160
  Client node                           = antman
  Client release level                  = SQL05020
  Client platform                       = AIX
  Client protocol                       = TCP/IP
  Client codepage                       = 819
  Process ID of client application      = 39324
  Client login ID                       = smith
  Host application ID                   = G9151251.GAA8.981015204854
  Sequence number                       = 0000
  Database alias at the gateway         = MVSDB
  DCS database name                     = DCSDB
  Host database name                    = GILROY
  Host release level                    = DSN05011
  Host CCSID                            = 500
```

The application status field contains one of the following three values:

1. `connect pending - outbound`. This means that the request to connect to a host database has been issued, and DB2 Connect is waiting for the connection to be established.

2. `waiting for request`. This means that the connection with the host database has been established, and that DB2 Connect is waiting for an SQL statement from the client application

3. `waiting for reply`. This means that the SQL statement has been sent to the host database.

Also, the status change time is only shown in the report if the System Monitor UOW switch was turned on during processing. Otherwise, ″Not Collected″ will be shown.

Using the DB2 Control Center to List Extended DCS Applications Information

You can use the DB2 Version 6 Control Center to perform DB2 Connect gateway monitoring. This section shows how you can use the Control Center to display the same report as provided by the `LIST DCS APPLICATIONS EXTENDED` command.

To view the extended report for any application:

1. Expand the tree under the `systems` icon of the Control Center to display Systems->Instances->Gateway Connections. If you right-click the mouse on any instance under the ″Gateway Connections″ folder, a pop-up menu will appear. Select the Applications... item from this menu. The Applications window appears. This window has a tabbed notebook appearance, with one tab titled ″Applications″. If there are gateway applications in your instance, there will be a second tab titled ″Gateway Applications″.

2. The main window of each page contains columns of information corresponding to the fields of the `LIST DCS APPLICATIONS EXTENDED` report. The first six columns, visible in the window, provide the following data:

   ```
   Client Node
   Application Name
   Client Application ID
   Host Application ID
   Database Alias at Gateway
   Status
   ```

   The rest of the fields of the report can be viewed by moving the horizontal scroll bar at the bottom of the window.

All of the fields listed by the `LIST DCS APPLICATIONS EXTENDED` command are present in this view.

Using the Windows NT Performance Monitor

More DB2 objects and counters have been added to the Windows NT Performance Monitor in Version 6. The new objects are: DB2 DCS Databases and DB2 DCS Applications.

If you are using your Windows NT system as a DB2 Connect gateway, you can now use the Windows NT Performance Monitor to graph all the DCS database and application counters provided by the DB2 System Monitor, i.e, the ones that can be displayed by issuing the GET SNAPSHOT FOR ALL DCS DATABASES or GET SNAPSHOT FOR ALL DCS APPLICATIONS commands. Refer to the *Administration Guide* for more DCS APPLICATIONS commands,

or for more information regarding using the Windows/NT Performance
Monitor with DB2 UDB and DB2 Connect.

Chapter 7. Security

This chapter describes DB2 Connect security considerations including Authentication types. It also provides some additional hints and tips on security for DB2 Universal Database for OS/390 users.

If you wish to know more about setting up security with DCE please refer to the DB2 Universal Database *Administration Guide*, and the database and DCE manuals for your host or AS/400 database server.

Note: When using DB2 Connect with DCE security, DCE software is required to be installed at the DB2 client workstation and on the host or AS/400 database server, but it is not necessary to install it on the DB2 Connect workstation. For further information about software prerequisites for DCE, please refer to your DB2 Connect *Quick Beginnings* manual.

Authentication

As DB2 Connect administrator, in cooperation with your host or AS/400 database administrator, you can determine where user names and passwords are validated. There are five possibilities:
- Validation at the client
- Validation at the DB2 Connect workstation
- Validation at both the DB2 Connect workstation and the host or AS/400 server
- Validation at the host or AS/400 server
- Validation at a DCE security server.

You determine where validation occurs by setting the Authentication type parameter in the system database directory, and the Security type parameter in the node directory for APPC or APPN nodes. For more information about updating these directories, see "Chapter 3. Updating Database Directories" on page 19.

Notes:

1. DB2 Connect itself performs no user validation. If you want to have the DB2 Connect workstation perform validation, the local security subsystem will be used to verify the userid and password provided with each `CONNECT` request. Therefore, when you set up a DB2 Connect Enterprise Edition gateway, if you will use `AUTHENTICATION=SERVER`, you must set up all the necessary userids and passwords on the gateway system.

2. If you use DCE Directory Services, authentication works differently. For more information, see "Security with DCE Directory Services" on page 150.

The following authentication types are allowed with DB2 Connect Version 5:

CLIENT
> The user name and password are validated at the client.

SERVER
> The user name and password are validated at the DB2 Connect workstation. When no authentication is specified, **SERVER** is assumed.

SERVER_ENCRYPT
> As for SERVER authentication, the user name and password are validated at the DB2 Connect workstation, but the transferred passwords are encrypted at the client and decrypted at the DB2 Connect workstation.

DCS The user name and password are validated at the host or AS/400 database server.

DCS_ENCRYPT
> As for DCS authentication, the user name and password are validated at the host or AS/400 database server, but the transferred passwords are encrypted at the client and, depending on the authentication type specified at the DB2 Connect workstation, decrypted at the DB2 Connect workstation or host or AS/400 database server.

DCE The user name and password are validated at the DCE security server.

SERVER_ENCRYPT and DCS_ENCRYPT authentication have the same semantics as SERVER and DCS authentication in terms of authentication location. They differ in that any transferred passwords will be encrypted at the source (the client or the DB2 Connect server) and decrypted at the target (the DB2 Connect server or the host or AS/400 database server) as specified by the authentication type catalogued at the source. Encrypted and non-encrypted values with matching authentication locations can then be used to choose different encryption combinations between client and Connect server or Connect server and host or AS/400 database server, while not affecting where authentication takes place. Here are some examples of how this might be used in a gateway scenario, where ″gateway″ is used to denote the Connect server:

Authentication at Client	Authentication at Gateway	Authentication Location	Client-Gateway Encryption?	Gateway-Server Encryption?
SERVER_ENCRYPT	SERVER	gateway	yes	no
DCS_ENCRYPT	DCS	server	yes	no
DCS	DCS_ENCRYPT	server	no	yes
DCS_ENCRYPT	DCS_ENCRYPT	server	yes	yes

The only APPC security parameter supported when either SERVER_ENCRYPT or DCS_ENCRYPT are used is SECURITY=NONE.

Notes:

1. For any system database directory entry that DB2 Connect uses for establishing a connection, if the authentication parameter is not specified, then DB2 Connect will use authentication **SERVER**.

2. As with DB2 Universal Database client-server communications, the authentication type is not required at a remote client attached to a DB2 Connect Enterprise Edition gateway, but it may be specified there in order to help optimize performance, since then it does not need to be gotten from the gateway, thus reducing the elapsed time for transactions.

3. In the case of a discrepancy between the value at the client and value at the gateway, the value specified at the DB2 Connect gateway takes precedence.

Security Types

This section lists the various combinations of authentication and security settings that are supported with DB2 Connect Version 6 over both APPC and TCP/IP connections.

The discussion which follows applies to both types of connection.

Security Types for APPC Connections

The following security types are allowed for APPC connections, in order to specify what security information will flow at the communications layer:

SAME Only the user name is passed to the host or AS/400 database server.

PROGRAM
The user name and password are passed to the host or AS/400 database server.

NONE
No security information flows.

Table 3 shows the possible combinations of these values and the authentication type specified on the DB2 Connect workstation, and where validation is performed for each combination. Only the combinations shown in this table are supported by DB2 Connect over APPC connections.

Table 3. Valid Security Scenarios for APPC connections

Case	Authentication setting in the database directory entry at the DB2 Connect workstation	Security	Validation
1	CLIENT	SAME	Client
2	SERVER	SAME	DB2 Connect server
3	SERVER	PROGRAM	DB2 Connect server and host or AS/400 database server
4	SERVER_ENCRYPT OR DCS_ENCRYPT	NONE	host or AS/400 database server
5	DCS	PROGRAM	host or AS/400 database server
6	DCE	NONE	DCE security server

If remote clients are connected to a DB2 Connect Enterprise Edition server, specify the following authentication and security types:

- If a remote client is connected to a DB2 Connect server via APPC, specify a security type of NONE at the remote client.
- If the authentication type in the database manager configuration at the DB2 Connect server is CLIENT, specify CLIENT at each remote client.
- If the authentication type at the DB2 Connect server is SERVER, SERVER_ENCRYPT, DCS, or DCS_ENCRYPT, specify one of these types at each remote client. (Which of these 4 types you specify at the remote client makes no difference.)

Notes:

1. For AIX systems, all login users using APPC security type SAME must belong to the AIX **system** group.
2. For AIX systems with remote clients, the instance of the DB2 Connect product running on the DB2 Connect workstation must belong to the AIX **system** group.
3. Access to a host or AS/400 database server is controlled by its own security mechanisms or subsystems; for example, the Virtual Telecommunications Access Method (VTAM) and Resource Access Control Facility (RACF). Access to protected database objects is controlled by the SQL **GRANT** and **REVOKE** statements.

Security Types for TCP/IP Connections

The TCP/IP communication protocol does not support security options at the network protocol layer. Thus only the authentication type controls where authentication takes place. Only the combinations shown in this table are supported by DB2 Connect over TCP/IP connections.

Table 4. Valid Security Scenarios for TCP/IP connections

Case	Authentication setting in the database directory entry at the DB2 Connect workstation	Validation
1	CLIENT	Client
2	SERVER OR SERVER_ENCRYPT	DB2 Connect workstation
3	Not applicable	None
4	DCS OR DCS_ENCRYPT	host or AS/400 database server
5	DCE	DCE security server

Discussion of Security Types

The following discussion applies to both APPC and TCP/IP connections, as described above and listed in Table 3 on page 66 and Table 4. Each case is described in more detail, as follows:

- In case 1, the user name and password are validated only at the remote client. (For a local client, the user name and password are validated only at the DB2 Connect server.)

  The user is expected to be authenticated at the location he or she first signs on to. The user ID is sent across the network, but not the password. Use this type of security only if all client workstations have adequate security facilities that can be trusted.

- In case 2, the user name and password are validated at the DB2 Connect server only. The password is sent across the network from the remote client to the DB2 Connect server but not to the host or AS/400 database server.

- In case 3, the user name and password are validated at both the DB2 Connect server and the host or AS/400 database server. The password is sent across the network from the remote client to the DB2 Connect workstation and from the DB2 Connect workstation to the host or AS/400 database server.

  Because validation is performed in two places, the same set of user names and passwords must be maintained at both the DB2 Connect server and the host or AS/400 database server.

- In case 4, the user name and password are validated at the host or AS/400 database server only. The user ID and password are sent across the network from the remote client to the DB2 Connect server and from the DB2 Connect server to the host or AS/400 database server.
- In case 5, a DCE encrypted ticket is obtained by the client from the DCE security server. The ticket is passed unaltered through DB2 Connect to the server, where it is validated by the server using DCE Security Services.

Additional Hints and Tips About Security

This section provides some additional hints and tips about security for users of DB2 Connect Version 6.

Extended Security Codes

Until DB2 Universal Database for OS/390 Version 5.1, connect requests that provided user IDs or passwords could fail with SQL30082 reason code 0, but no other indication as to what might be wrong.

DB2 Universal Database for OS/390 Version 5.1 introduces an enhancement which provides support for extended security codes. Specifying extended security will provide additional diagnostics, such as (PASSWORD EXPIRED) in addition to the reason code.

In order to exploit this, the DB2 Universal Database for OS/390 ZPARM installation parameter for extended security should be set to the value YES. Use the DB2 Universal Database for OS/390 installation panel DSN6SYSP to set EXTSEC=YES. You can also use DDF panel 1 (DSNTIPR) to set this. The default value is EXTSEC=NO. In the case of an expired password, PC, UNIX, Apple Macintosh, and Web applications using DB2 Connect will receive error message SQL01404.

TCP/IP Security Already Verified

If you wish to provide support for the DB2 Universal Database security option AUTHENTICATION=CLIENT, then use DB2 Universal Database for OS/390 installation panel DSNTIP4 (DDF panel 2) to set TCP/IP already verified security to YES.

Desktop ODBC and Java Application Security

Workstation ODBC and Java applications use dynamic SQL. This may create security concerns in some installations. DB2 Universal Database for OS/390 introduces a new bind option DYNAMICRULES(BIND) that allows execution of

dynamic SQL under the authorization of either the owner or the binder. Refer to the *Command Reference* to see how DYNAMICRULES can be specified through DB2 Connect.

DB2 Universal Database and DB2 Connect Version 5 provide a new CLI/ODBC configuration parameter CURRENTPACKAGESET in the DB2CLI.INI configuration file. This should be set to a schema name that has the appropriate privileges. An SQL SET CURRENT PACKAGESET schema statement will automatically be issued after every connect for the application.

Use the ODBC Manager to update DB2CLI.INI. See *Installation and Configuration Supplement* for further information.

Password Change Support

If an SQL CONNECT statement returns a message indicating that the user ID's password has expired, with DB2 Connect Version 6 it is now possible to change the password without signing on to TSO. Through DRDA, DB2 Universal Database for OS/390 can change the password for you.

The old password along with the new password and the verify password must be supplied by the user. If the security specified at the DB2 Connect Enterprise Edition server is DCS then a request to change the password is sent to the DB2 Universal Database for OS/390 database server. If the security specified is SERVER then the password on the DB2 Connect server is changed.

An additional benefit is that a separate LU definition is no longer required, as was the case with DB2 Connect Version 5.0. Refer to the DB2 Connect Enterprise Edition *Quick Beginnings* manual for additional information.

Chapter 8. Programming in a DB2 Connect Environment

This section provides some information about creating applications that use DB2 Connect. For more information, refer to the *CLI Guide and Reference*, the *Command Reference*, and the *Application Development Guide*.

Programming in a Distributed Environment

DB2 Connect lets an application program access data in any host or AS/400 server database. For example, an application running on OS/2 can access data in a DB2 Universal Database for OS/390 database. You can create new applications, or modify existing applications to run in a host or AS/400 environment. You can also develop applications in one environment and port them to another.

DB2 Connect enables you to use the following items with host database products such as DB2 Universal Database for OS/390, as long as the item is supported by the host database product:
* Embedded SQL, both static and dynamic
* Compound SQL, as described in "NOT ATOMIC Compound SQL" on page 81
* Stored procedures, as described in "Stored Procedures" on page 80
* The DB2 Call Level Interface
* The Microsoft ODBC API
* JDBC.

Some SQL statements differ among relational database products. You may encounter the following situations:
* SQL statements that are the same for all the database products that you use regardless of standards
* SQL statements that are documented in the *SQL Reference* and are therefore available in all IBM relational database products
* SQL statements that are unique to one database system that you access.

SQL statements in the first two categories are highly portable, but those in the third category will first require changes.

In general, SQL statements in Data Definition Language (DDL) are not as portable as those in Data Manipulation Language (DML).

DB2 Connect accepts some SQL statements that are not supported by DB2 Universal Database. DB2 Connect passes these statements on to the host or AS/400 server.

For information on limits on different platforms, such as the maximum column length, refer to the *SQL Reference.*

If you move a CICS application from OS/390 or VSE to run under another CICS product (for example, CICS for AIX), it can also access the OS/390 or VSE database using DB2 Connect. Refer to the *CICS/6000 Application Programming Guide* and the *CICS Customization and Operation* manual for more details.

When you program in a host or AS/400 environment, you should consider the following specific factors:
- Using Data Definition Language (DDL)
- Using Data Manipulation Language (DML)
- Using Data Control Language (DCL)
- Connecting and disconnecting
- Precompiling
- Defining a sort order
- Managing referential integrity
- Locking
- Differences in SQLCODEs and SQLSTATEs
- Using system catalogs
- Isolation levels
- Stored procedures
- NOT ATOMIC compound SQL
- Distributed unit of work
- SQL statements supported or rejected by DB2 Connect.

Using Data Definition Language (DDL)

DDL statements differ among the IBM database products because storage is handled differently on different systems. On host or AS/400 server systems, there can be several steps between designing a database and issuing a CREATE TABLE statement. For example, a series of statements may translate the design of logical objects into the physical representation of those objects in storage.

The precompiler passes many such DDL statements to the host or AS/400 server when you precompile to a host or AS/400 server database. The same statements would not precompile against a database on the system where the

application is running. For example, in an OS/2 application the CREATE STORGROUP statement will precompile successfully to a DB2 Universal Database for OS/390 database, but not to a DB2 for OS/2 database.

Using Data Manipulation Language (DML)

In general, DML statements are highly portable. SELECT, INSERT, UPDATE, and DELETE statements are similar across the IBM relational database products. Most applications primarily use DML SQL statements, which are supported by the DB2 Connect program.

Numeric Data Types

When numeric data is transferred to DB2 Universal Database, the data type may change. Numeric and zoned decimal SQLTYPEs (supported by DB2 Universal Database for AS/400) are converted to fixed (packed) decimal SQLTYPEs.

Mixed-Byte Data

Mixed-byte data can consist of characters from an extended UNIX code (EUC) character set, a double-byte character set (DBCS) and a single-byte character set (SBCS) in the same column. On systems that store data in EBCDIC (OS/390, OS/400, VSE, and VM), shift-out and shift-in characters mark the start and end of double-byte data. On systems that store data in ASCII (such as OS/2 and UNIX), shift-in and shift-out characters are not required.

If your application transfers mixed-byte data from an ASCII system to an EBCDIC system, be sure to allow enough room for the shift characters. For each switch from SBCS to DBCS data, add 2 bytes to your data length. For better portability, use variable-length strings in applications that use mixed-byte data.

Long Fields

Long fields (strings longer than 254 characters) are handled differently on different systems. A host or AS/400 server may support only a subset of scalar functions for long fields; for example, DB2 Universal Database for OS/390 allows only the **LENGTH** and **SUBSTR** functions for long fields. Also, a host or AS/400 server may require different handling for certain SQL statements; for example, DB2 for VSE & VM requires that with the INSERT statement, only a host variable, the SQLDA, or a NULL value be used.

Large Object (LOB) Data Type

The LOB data type is supported by DB2 Connect.

User Defined Types (UDTs)

Only User Defined Distinct Types are supported by DB2 Connect. Abstract Data Types are not.

ROWID Data Type

The ROWID data type is handled by DB2 Connect Version 6 as VARCHAR for bit data.

64-bit Integer (BIGINT) data type

Eight byte (64-bit) integers are supported by DB2 Connect Version 6. The BIGINT internal datatype is used in order to provide support for the cardinality of very large databases, while retaining data precision.

Using Data Control Language (DCL)

Each IBM relational database management system provides different levels of granularity for the GRANT and REVOKE SQL statements. Check the product-specific publications to verify the appropriate SQL statements to use for each database management system.

Connecting and Disconnecting

DB2 Connect supports the CONNECT TO and CONNECT RESET versions of the CONNECT statement, as well as CONNECT with no parameters. If an application calls an SQL statement without first performing an explicit CONNECT TO statement, an *implicit* connect is performed to the default application server (if one is defined).

When you connect to a database, information identifying the relational database management system is returned in the SQLERRP field of the SQLCA. If the application server is an IBM relational database, the first three bytes of SQLERRP contain one of the following:

DSN   DB2 Universal Database for OS/390

ARI   DB2 for VSE & VM

QSQ   DB2 Universal Database for AS/400

SQL   DB2 Universal Database.

If you issue a CONNECT TO or null CONNECT statement while using DB2 Connect, the country code or territory token in the SQLERRMC field of the SQLCA is returned as blanks; the CCSID of the application server is returned in the code page or code set token.

You can explicitly disconnect by using the CONNECT RESET statement (for type 1 connect), the RELEASE and COMMIT statements (for type 2 connect), or the DISCONNECT statement (either type of connect, but not in a TP monitor environment). Type 2 connect is used for distributed unit of work, as described in "Multisite Updates (Two-Phase Commit)" on page 13.

If a connection is not explicitly disconnected and the application ends normally, DB2 Connect commits the resulting data implicitly.

Note: An application can receive SQLCODEs indicating errors and still end normally; DB2 Connect commits the data in this case. If you do not want the data to be committed, you must issue a ROLLBACK command.

The FORCE command lets you disconnect selected users or all users from the database. This is supported for host or AS/400 server databases; the user can be forced off the DB2 Connect workstation.

Precompiling

There are some differences in the precompilers for different IBM relational database systems. The precompiler for DB2 Universal Database differs from the host or AS/400 server precompilers in the following ways:
- It makes only one pass through an application.
- When binding against DB2 Universal Database databases, objects must exist for a successful bind. VALIDATE RUN is not supported.

Blocking

The DB2 Connect program supports the DB2 database manager blocking bind options:

UNAMBIG
> Only unambiguous cursors are blocked (the default).

ALL   Ambiguous cursors are blocked.

NO   Cursors are not blocked.

The DB2 Connect program uses the block size defined in the DB2 database manager configuration file for the RQRIOBLK field. Current versions of DB2 Connect support block sizes up to 32 767. If larger values are specified in the DB2 database manager configuration file, DB2 Connect uses a value of 32 767 but does not reset the DB2 database manager configuration file. Blocking is handled the same way using the same block size for dynamic and static SQL.

Note: Most host or AS/400 server systems consider dynamic cursors ambiguous, but DB2 Universal Database systems consider some dynamic cursors unambiguous. To avoid confusion, you can specify BLOCKING ALL with DB2 Connect.

Specify the block size in the DB2 database manager configuration file by using the CLP, the Control Center, or an API, as listed in the *Administrative API Reference* and *Command Reference.*

Package Attributes

A package has the following attributes:

Collection ID
> The ID of the package. It can be specified on the PREP command.

Owner
> The authorization ID of the package owner. It can be specified on the PREP or BIND command.

Creator
> The user name that binds the package.

Qualifier
> The implicit qualifier for objects in the package. It can be specified on the PREP or BIND command.

Each host or AS/400 server system has limitations on the use of these attributes:

DB2 Universal Database for OS/390
> All four attributes can be different. The use of a different qualifier requires special administrative privileges. For more information on the conditions concerning the usage of these attributes, refer to the *Command Reference* for DB2 Universal Database for OS/390.

DB2 for VSE & VM
> All of the attributes must be identical. If USER1 creates a bind file (with PREP), and USER2 performs the actual bind, USER2 needs DBA authority to bind for USER1. Only USER1's user name is used for the attributes.

DB2 Universal Database for AS/400
> The qualifier indicates the collection name. The relationship between qualifiers and ownership affects the granting and revoking of privileges on the object. The user name that is logged on is the creator and owner unless it is qualified by a collection ID, in which case the collection ID is the owner. The collection ID must already exist before it is used as a qualifier.

DB2 Universal Database
> All four attributes can be different. The use of a different owner requires administrative authority and the binder must have CREATEIN privilege on the schema (if it already exists).

Note: DB2 Connect provides support for the *SET CURRENT PACKAGESET* command for DB2 Universal Database for OS/390 and DB2 Universal Database.

C Null-terminated Strings

The CNULREQD bind option overrides the handling of null-terminated strings that are specified using the LANGLEVEL option. Refer to *Application Development Guide* for a description of how null-terminated strings are handled when prepared with the LANGLEVEL option set to MIA or SAA1. By default, CNULREQD is set to YES. This causes null-terminated strings to be interpreted according to MIA standards. If connecting to a DB2 Universal Database for OS/390 server it is strongly recommended to set CNULREQD to YES. You need to bind applications coded to SAA1 standards (with respect to null-terminated strings) with the CNULREQD option set to NO. Otherwise, null-terminated strings will be interpreted according to MIA standards, even if they are prepared using LANGLEVEL set to SAA1.

Standalone SQLCODE and SQLSTATE

Standalone SQLCODE and SQLSTATE variables, as defined in ISO/ANS SQL92, are supported through the LANGLEVEL SQL92E precompile option. An SQL0020W warning will be issued at precompile time, indicating that LANGLEVEL is not supported. This warning applies only to the features listed under LANGLEVEL MIA in the *Command Reference*, which is a subset of LANGLEVEL SQL92E.

Defining a Sort Order

The differences between EBCDIC and ASCII cause differences in sort orders in the various database products, and also affect ORDER BY and GROUP BY clauses. One way to minimize these differences is to create a user-defined collating sequence that mimics the EBCDIC sort order. You can specify a collating sequence only when you create a new database. For more information, refer to the *Application Development Guide*, the *Administrative API Reference* and the *Command Reference*.

Note: Database tables can now be stored on DB2 Universal Database for OS/390 in ASCII format. This permits faster exchange of data between DB2 Connect and DB2 Universal Database for OS/390, and removes the need to provide field procedures which must otherwise be used to convert data and resequence it.

Managing Referential Integrity

Different systems handle referential constraints differently:

DB2 Universal Database for OS/390
> An index must be created on a primary key before a foreign key can be created using the primary key. Tables can reference themselves.

DB2 for VSE & VM
> An index is automatically created for a foreign key. Tables cannot reference themselves.

DB2 Universal Database for AS/400
> An index is automatically created for a foreign key. Tables can reference themselves.

DB2 Universal Database
> For DB2 Universal Database databases, an index is automatically created for a unique constraint, including a primary key. Tables can reference themselves.

Other rules vary concerning levels of cascade.

Locking

The way in which the database server performs locking can affect some applications. For example, applications designed around row-level locking and the isolation level of cursor stability are not directly portable to systems that perform page-level locking. Because of these underlying differences, applications may need to be adjusted.

The DB2 Universal Database for OS/390 and DB2 Universal Database products have the ability to time-out a lock and send an error return code to waiting applications.

Differences in SQLCODEs and SQLSTATEs

Different IBM relational database products do not always produce the same SQLCODEs for similar errors. You can handle this problem in either of two ways:

- Use the SQLSTATE instead of the SQLCODE for a particular error.

  SQLSTATEs have approximately the same meaning across the database products, and the products produce SQLSTATEs that correspond to the SQLCODEs.

- Map the SQLCODEs from one system to another system.

  By default, DB2 Connect maps SQLCODEs and tokens from each IBM host or AS/400 server system to your DB2 Universal Database system. You can specify your own SQLCODE mapping file if you want to override the

default mapping or you are using a database server that does not have SQLCODE mapping (a non-IBM database server). You can also turn off SQLCODE mapping. For more information, see "Chapter 9. SQLCODE Mapping" on page 89.

Using System Catalogs

The system catalogs vary across the IBM database products. Many differences can be masked by the use of views. For information, see the documentation for the database server that you are using.

The catalog functions in CLI get around this problem by presenting support of the same API and result sets for catalog queries across the DB2 family.

Numeric Conversion Overflows on Retrieval Assignments

Numeric conversion overflows on retrieval assignments may be handled differently by different IBM relational database products. For example, consider fetching a float column into an integer host variable from DB2 Universal Database for OS/390 and from DB2 Universal Database. When converting the float value to an integer value, a conversion overflow may occur. By default, DB2 Universal Database for OS/390 will return a warning SQLCODE and a null value to the application. In contrast, DB2 Universal Database will return a conversion overflow error. It is recommended that applications avoid numeric conversion overflows on retrieval assignments by fetching into appropriately sized host variables.

Isolation Levels

DB2 Connect accepts the following isolation levels when you prep or bind an application:

RR   Repeatable Read

RS   Read Stability

CS   Cursor Stability

UR   Uncommitted Read

NC   No Commit

The isolation levels are listed in order from most protection to least protection. If the host or AS/400 server does not support the isolation level that you specify, the next higher supported level is used.

Table 5 on page 80 shows the result of each isolation level on each host or AS/400 application server.

Table 5. Isolation Levels

DB2 Connect	DB2 Universal Database for OS/390	DB2 for VSE & VM	DB2 Universal Database for AS/400	DB2 Universal Database
RR	RR	RR	note 1	RR
RS	note 2	RR	COMMIT(*ALL)	RS
CS	CS	CS	COMMIT(*CS)	CS
UR	note 3	CS	COMMIT(*CHG)	UR
NC	note 4	note 5	COMMIT(*NONE)	UR

Note:

1. There is no equivalent COMMIT option on DB2 Universal Database for AS/400 that matches RR. DB2 Universal Database for AS/400 supports RR by locking the whole table.

2. Results in RR for Version 3.1, and results in RS for Version 4.1 with APAR PN75407 or Version 5.1.

3. Results in CS for Version 3.1, and results in UR for Version 4.1 or Version 5.1.

4. Results in CS for Version 3.1, and results in UR for Version 4.1 with APAR PN60988 or Version 5.1.

5. Isolation level NC is not supported with DB2 for VSE & VM.

With DB2 Universal Database for AS/400, you can access an unjournalled table if an application is bound with an isolation level of UR and blocking set to ALL, or if the isolation level is set to NC.

Stored Procedures

- Invocation

  A client program can invoke a server program by issuing an SQL CALL statement. Each server works a little differently to the other servers in this case.

  OS/390

  The schema name must be no more than 8 bytes long, the procedure name must be no more than 18 bytes long, and the stored procedure must be defined in the SYSIBM.SYSPROCEDURES catalog on the server.

  VSE or VM

  The procedure name must not be more than 18 bytes long and must be defined in the SYSTEM.SYSROUTINES catalog on the server.

  OS/400

  The procedure name must be an SQL identifier. You can also use the DECLARE PROCEDURE or CREATE PROCEDURE statements to specify the actual path name (the schema-name or collection-name) to locate the stored procedure.

All CALL statements to DB2 for AS/400 from REXX/SQL must be dynamically prepared and executed by the application as the CALL statement implemented in REXX/SQL maps to CALL USING DESCRIPTOR.

For the syntax of the SQL CALL statement, refer to the *SQL Reference.* For information on how to use stored procedures when writing application programs, refer to the *Application Development Guide.*

You can invoke the server program on DB2 Universal Database with the same parameter convention that server programs use on DB2 Universal Database for OS/390, DB2 Universal Database for AS/400, or DB2 for VSE & VM. For more information on invoking DB2 Universal Database stored procedures, see the *Application Development Guide.* For more information on the parameter convention on other platforms, refer to the DB2 product documentation for that platform.

All the SQL statements in a stored procedure are executed as part of the SQL unit of work started by the client SQL program.
- Do not pass indicator values with special meaning to or from stored procedures.

  Between DB2 Universal Database, the systems pass whatever you put into the indicator variables. However, when using DB2 Connect, you can only pass 0, -1, and -128 in the indicator variables.
- You should define a parameter to return any errors or warning encountered by the server application.

  A server program on DB2 Universal Database can update the SQLCA to return any error or warning, but a stored procedure on DB2 Universal Database for OS/390 or DB2 Universal Database for AS/400 has no such support. If you want to return an error code from your stored procedure, you must pass it as a parameter. The SQLCODE and SQLCA is only set by the server for system detected errors.
- DB2 for VSE & VM Version 6 or higher and DB2 Universal Database for OS/390 Version 5.1 or higher are the only host or AS/400 Application Servers that can return the result sets of stored procedures at this time.

NOT ATOMIC Compound SQL

Compound SQL allows multiple SQL statements to be grouped into a single executable block. This may reduce network overhead and improve response time.

DB2 Connect supports NOT ATOMIC compound SQL. This means that processing of compound SQL continues following an error. (With ATOMIC

compound SQL, which is not supported by DB2 Connect, an error would roll back the entire group of compound SQL.)

Statements will continue execution until terminated by the application server. In general, execution of the compound SQL statement will be stopped only in the case of serious errors.

NOT ATOMIC compound SQL can be used with all of the supported host or AS/400 application servers.

If multiple SQL errors occur, the SQLSTATEs of the first seven failing statements are returned in the SQLERRMC field of the SQLCA with a message that multiple errors occurred. For more information, refer to the *SQL Reference.*

Multisite Update with DB2 Connect

DB2 Connect allows you to perform a multisite update, also known as an update of multiple databases within a single distributed unit of work (DUOW). Whether you can use this capability depends on several factors:

- Your application program must be precompiled with the CONNECT 2 and SYNCPOINT TWOPHASE options.
- If you have SNA network connections, you can use two-phase commit support provided by the DB2 Syncpoint Manager (SPM) function of DB2 Connect Enterprise Edition Version 6 on AIX, OS/2, and Windows NT. This enables the following database servers to participate in a distributed unit of work:
  - DB2 for AS/400 Version 3.1 or later
  - DB2 for MVS/ESA Version 3.1 or later
  - DB2 for OS/390 Version 5.1 or later
  - DB2 for VM & VSE Version V5.1 or later.

  The above is true for native DB2 UDB applications and applications coordinated by an external Transaction Processing Monitor such as IBM TXSeries, CICS for Open Systems, Encina Monitor, and Microsoft Transaction Server.
- If you have TCP/IP network connections then the following database servers can participate in a distributed unit of work:
  - DB2 for OS/390 V5.1 or later.

  If the application is controlled by a Transaction Processing Monitor such as IBM TXSeries, CICS for Open Sytems, Encina Monitor, or Microsoft Transaction Server, then you must use the DB2 Syncpoint Manager. With

DB2 UDB Version 6 (including DB2 UDB EE, DB2 UDB EEE, and DB2 Connect EE), the DB2 Syncpoint Manager has been enhanced to support TCP/IP connections.

If the application is a native DB2 UDB application, then the DB2 Syncpoint Manager is not required for two phase commit.

If a common DB2 Connect Enterprise Edition gateway is used by both native DB2 applications and TP Monitor Applications to access host data over TCP/IP connections then the DB2 Syncpoint Manager must be used.

Note: If a single DB2 Connect Enterprise Edition gateway will be used to access host data using both SNA and TCP/IP network protocols and two phase commit is required, then the DB2 Syncpoint Manager must be used. This is true for both DB2 Applications and TP Monitor applications.

Notes:

1. DB2 Common Server Version 2.1 databases can be updated with two-phase commit in a unit of work *only* when DB2 Universal Database for OS/390 Version 5.1 is *not* the transaction manager database.

2. If the TM_DATABASE for the transaction is DB2 Universal Database for OS/390 Version 5.1, then DB2 CS V2.1 databases participating in that transaction become *READ-ONLY* for the client application.

3. DB2 Universal Database Version 6 client applications can participate in distributed units of work with different levels of database server only if DB2 Universal Database for OS/390 Version 5.1 is the *TM_DATABASE* for the transaction. For further information refer to either *DB2 Connect Enterprise Edition for OS/2 and Windows NT Quick Beginnings*, or *DB2 Connect Personal Edition Quick Beginnings*.

4. The database levels supported for two-phase commit transactions over TCP/IP depend on the DB2 client level, the TM_DATABASE level, and the participant database levels.

   We strongly recommend that *all* clients accessing DB2 Universal Database Version 6 and DB2 Universal Database for OS/390 Version 5 databases be at DB2 Universal Database Version 6. DB2 Version 2.1 clients *cannot* initiate two-phase commit transactions if DB2 Universal Database for OS/390 Version 5 databases participate in the transaction.

Host or AS/400 Server SQL Statements Supported by DB2 Connect

The following statements compile successfully for host or AS/400 server processing but not for processing with DB2 Universal Database systems:

- ACQUIRE
- DECLARE (modifier.(qualifier.)table_name TABLE ...

- LABEL ON

These statements are also supported by the command line processor.

The following statements are supported for host or AS/400 server processing but are not added to the bind file or the package and are not supported by the command line processor:
- DESCRIBE statement_name INTO descriptor_name USING NAMES
- PREPARE statement_name INTO descriptor_name USING NAMES FROM ...

The precompiler makes the following assumptions:
- Host variables are input variables
- The statement is assigned a unique section number.

Host or AS/400 Server SQL Statements Rejected by DB2 Connect

The following SQL statements are not supported by DB2 Connect and not supported by the command line processor:
- COMMIT WORK RELEASE
- DECLARE state_name, statement_name STATEMENT
- DESCRIBE statement_name INTO descriptor_name USING xxxx (where xxxx is ANY, BOTH, or LABELS)
- PREPARE statement_name INTO descriptor_name USING xxxx FROM :host_variable (where xxxx is ANY, BOTH, or LABELS)
- PUT ...
- ROLLBACK WORK RELEASE
- SET :host_variable = CURRENT ...

DB2 for VSE & VM extended dynamic SQL statements are rejected with -104 and syntax error SQLCODEs.

Implementing Chargeback Accounting

Many DB2 Universal Database for OS/390 installations implement resource monitoring practices which allow systems administrators to associate resource usage with individual user access. This can be used to charge individual users or their departments for the resources they consume. This practice is commonly referred to as *charge-back accounting*. DB2 Connect products allow system administrators to monitor mainframe resources consumed by users accessing databases through DB2 Connect.

You can use accounting strings to send accounting data from DB2 Connect to the host or AS/400 application server. An accounting string combines system-generated data with user-supplied data. This data lets system administrators associate resource usage with each user's access, and charge the users accordingly.

The accounting string is sent using the DRDA parameter PRDDTA. Because the content of this parameter is not architected in DRDA, there is no guarantee that your application server will recognize the data as accounting data. Currently, PRDDTA is supported as follows:

MVS and OS/390
The string is stored as an accounting record.

VSE or VM
The string is ignored.

OS/400
The string is ignored.

other For OS/2 and UNIX-based systems, the string is ignored.

The accounting string consists of 56 bytes generated by DB2 Connect (the prefix) followed by up to 199 bytes specified by the user (the suffix), for a maximum length of 255.

Table 6 shows the system-generated fields. Each of these fields is padded to the right with blanks.

Table 6. Accounting String Fields Generated by DB2 Connect

Field Name	Length	Description.
acct_str_len	1	A hexadecimal value representing the length of the accounting string minus 1. For example, X'3C'.
client_prdid	8	The product ID of the client software. For example, the product ID for DB2 Universal Database Version 6 is SQL05020.
client_platform	18	The platform the client is on, for example AIX, OS/2, DOS, or Windows.

Table 6. Accounting String Fields Generated by DB2 Connect (continued)

Field Name	Length	Description.
client_appl_name	20	The first 20 characters of the user's application name, for example, `payroll`.
client_authid	8	The authid of the user's application, for example, `SMITH`.
suffix_len	1	A hexadecimal value representing the length of the user-supplied suffix. X'00' means that there is no user-supplied suffix.

The user-defined suffix is one of the following:

- The value specified by an application with the sqlesact() API
- The value of the DB2ACCOUNT environment variable
- The value of the DFT_ACCOUNT_STR (default accounting string) configuration parameter
- A null string.

If the suffix is longer than 199 characters, it is truncated. To be sure that the accounting string is converted correctly when it is transmitted to the host or AS/400 database server, you should use only the characters A to Z, 0 to 9 and underscore (_).

The API method of setting the accounting string is recommended. Your application should call the API before connecting to a database. If you want to change the accounting string within the application (for example, to send a different string when you connect to a different database), call the API again. Otherwise, the value remains in effect until the end of the application.

If the sqlesact() API is not called before the first database connection request, the DB2ACCOUNT environment variable is read. This value remains in effect until the end of the application or background command line processor process. To specify a new accounting string suffix after the first database connection, either use the sqlesact() API or end the application or background CLP process and restart it with DB2ACCOUNT set to its new value.

If no DB2ACCOUNT value exists, the value of the DFT_ACCOUNT_STR system configuration parameter is used. This default can be useful for database clients that do not have the ability to forward an accounting string to DB2 Connect. If this does not exist, a null string is used.

The following are examples of accounting strings:

```
x'3C'SQL020100S/2          cheque          SMITH  x'05'DEPT1

x'37'SQL020100S/2          cheque          SMITH  x'00'
```

In the first example, the user-defined suffix is DEPT1. In the second example, it is a null string.

Useful Publications

The following publications may help you develop applications that run in a distributed environment:

- The application programming books for the specific database products may contain information that differs from one product to another.
- The SQL reference books for the specific database products will help you ensure that an application contains only supported SQL statements, with the correct syntax.
- The *DB2 Universal Database for OS/390 Version 5 Reference for Remote DRDA Requesters and Servers* provides the latest information about chargeback accounting for DB2 Universal Database for OS/390 users.
- *IBM SQL Reference* provides a high-level discussion of differences among the IBM relational database products, and also discusses how to handle some specific differences.
- The DRDA publications provide information on planning, connectivity, programming, and problem determination in the DRDA environment. For a list of titles and order numbers, see "Related DRDA Publications" on page ix.

Chapter 9. SQLCODE Mapping

Different IBM relational database products do not always produce the same SQLCODEs for similar errors. Even when the SQLCODE is the same, it may be accompanied by *tokens* that are specified differently. (The token list is passed in the SQLERRMC field of the SQLCA.) By default, DB2 Connect maps SQLCODEs and tokens from each IBM host or AS/400 database server to the appropriate DB2 UDB SQLCODEs.

Turning Off SQLCODE Mapping

If you want to turn off SQLCODE mapping, specify **NOMAP** in the parameter string of the DCS Directory or the DCE routing information object. For information about updating the DCS directory, see "Chapter 3. Updating Database Directories" on page 19. For information about using DCE, see "Appendix E. Using DCE Directory Services" on page 143.

If you port an application directly from a host or AS/400 database server (such as DB2 Universal Database for OS/390), you might want to turn off SQLCODE mapping. This would let you use the application without changing the SQLCODEs that it references.

Tailoring the SQLCODE Mapping

By default, DB2 Connect maps SQLCODEs and tokens from each IBM host or AS/400 database server to the appropriate DB2 UDB SQLCODEs. The following files are copies of the default SQLCODE mapping:

dcs1dsn.map
> Maps DB2 for MVS/ESA and DB2 Universal Database for OS/390 SQLCODEs

dcs1ari.map
> Maps DB2 for VSE & VM SQLCODEs

dcs1qsq.map
> Maps DB2 Universal Database for AS/400 SQLCODEs

No mapping is required for OS/2 and UNIX-based DB2 systems.

If you want to override the default SQLCODE mapping or you are using a host or AS/400 database server that does not have SQLCODE mapping (a non-IBM database server), you can copy one of these files and use it as the

basis for your new SQLCODE mapping file. (By copying the file rather than editing it directly, you ensure that you can always refer to the original SQLCODE mapping if necessary.)

Specify the file name of your new SQLCODE mapping file in the parameter string of the DCS Directory or the DCE routing information object. For information about updating the DCS directory, see "Chapter 3. Updating Database Directories" on page 19. For information about using DCE, see "Appendix E. Using DCE Directory Services" on page 143.

Each mapping file is an ASCII file, which is created and edited using an ASCII editor. At initial installation, the file is stored in the map directory in the installation path.

The file can contain the following special types of lines:

&& The logical beginning of the file. All lines before the first occurrence of && are considered free-form comments and ignored. If the file contains nothing after &&, no SQLCODE mapping is performed. (You can also turn off SQLCODE mapping with the NOMAP parameter, as described previously.)

* As the first character on a line, indicates a comment.

W As the only character on a line, indicates that warning flags should be remapped. (By default, the original warning flags are passed.) The W must be uppercase.

All other lines after && must be either blank or mapping statements in the following form:

```
input_code [, output_code [, token_list]]
```

input_code is one of the following:

sqlcode The SQLCODE from the host or AS/400 database server.

U All undefined negative SQLCODEs (those not listed in this file) are mapped to the specified *output_code*. If no *output_code* is specified on this line, the original SQLCODE is used. This character must be uppercase.

P All undefined positive SQLCODEs (those not listed in this file) are mapped to the specified *output_code*. If no *output_code* is specified on this line, the original SQLCODE is used. This character must be uppercase.

cc*nn* The SQLSTATE class code from the host or AS/400 database server. *nn* is one of the following:

00 Unqualified successful completion

01	Warning
02	No data
21	Cardinality violation
22	Data exception
23	Constraint violation
24	Invalid cursor state
26	Invalid SQL statement identifier
40	Transaction Rollback
42	Access violation
51	Invalid application state
55	Object not in prerequisite state
56	Miscellaneous SQL or Product Error
57	Resource not available or operator intervention
58	System error

The specified *output_code* is used for all SQLCODEs with this class code that are not specified explicitly in the mapping file. If no *output_code* is specified on this line, the original SQLCODE is mapped to itself with no tokens copied over.

The characters **cc** must be lowercase.

If the same *input_code* appears more than once in the mapping file, the first occurrence is used.

output_code is the output SQLCODE. If no value is specified, the original SQLCODE is used.

If you specify an output code, you can also specify one of the following:

(s) The input SQLCODE plus the product ID (ARI, DSN or QSQ) will be put into the SQLCA message token field.

The original SQLCODE is returned as the only token. This option is designed to handle undefined SQLCODEs, with the exception of +965 and -969. If +965 or -969 is the *output_code*, the token list returned in the SQLERRMC field of the SQLCA includes the original SQLCODE, followed by the product identifier, followed by the original token list.

The character **s** must be lowercase.

(*token-list***)**

> A list of tokens, separated by commas. Specify only a comma to skip a particular token. For example, the form (*,t2,,t4*) means that the first and third output tokens are null.
>
> Each token has the form of a number (*n*), optionally preceded by **c**, optionally followed by **c** or **i**. It is interpreted as follows:

c   The datatype of the token in this position is CHAR (the default). If **c** comes before *n*, it refers to the input token; if it comes after *n*, it refers to the output token. The character **c** must be lowercase.

i   The datatype of the token in this position is INTEGER. If **i** comes after *n*, it refers to the output token. **i** should not come before *n*, because IBM host or AS/400 database server products support only CHAR tokens. The character **i** must be lowercase.

n   A number or numbers indicating which host or AS/400 database server tokens are used. They are arranged in the order desired for placement in the output SQLCA. The number indicates the host or AS/400 database server token; the arrangement indicates the order in which the tokens will be placed in the SQLCA.

> For example, the host or AS/400 database server might return two tokens, 1 and 2. If you want token 2 to appear before token 1 in the output SQLCA, specify (2,1).
>
> Multiple token numbers can be combined to form one CHAR output token by connecting them with periods.
>
> Commas are used to separate output tokens. If no token is specified before a comma, no output token is included in the SQLCA for that position. Any tokens occurring in the output SQLCA following the last specified token are mapped to a null token.

Figure 7 on page 93 shows a sample SQLCODE mapping file.

```
&&
 -007    ,   -007   ,   (1)
 -010
 -060    ,   -171   ,   (2)
...
  -204    ,   -204   ,   (c1.2c)
...
  -633    ,   -206   ,   (,c1i)

 -30021  ,  -30021 ,   (c1c,c2c)

 cc00    ,  +000
...
  U      ,   -969   ,   (s)
  P      ,   +965   ,   (s)
```

Figure 7. An SQLCODE Mapping File

Each mapping statement in the file is described below:

1. The SQLCODE is mapped from -007 to -007. The first input token received from the host or AS/400 database server is used as the first output token, and it defaults to CHAR. No other tokens are transferred.

2. The SQLCODE is mapped from -010 to -010 (no output SQLCODE is specified). No tokens are put into the output SQLCA.

3. The SQLCODE is mapped from -060 to -171. The first input token received from the host or AS/400 database server is discarded. The second is used as the first token in the output SQLCA, and it is CHAR. There is no second token in the output SQLCA.

4. The SQLCODE is mapped from -204 to -204. The first and second tokens received from the host or AS/400 database server are CHAR. These two input tokens are combined to form one CHAR output token, which will be the first output token in the SQLCA.

5. The SQLCODE is mapped from -633 to -206. The first input token received from the host or AS/400 database server is CHAR. It is converted to INTEGER and is used as the second token in the output SQLCA. The first token in the output SQLCA is null, as indicated by a comma.

6. The SQLCODE is mapped from -30021 to -30021. The first and second input tokens received from the host or AS/400 database server are CHAR, and they are used as the first and second tokens in the output SQLCA.

7. All SQLCODEs in SQLCAs with SQLSTATEs in the 00 class will be mapped to SQLCODE +000.

8. All undefined SQLCODEs are mapped to -969. This option should be used only if all mappable codes are listed, including all those that are identical and require no mapping. The **(s)** option indicates that the token list to be returned in the SQLERRMC field of the SQLCA includes the original

SQLCODE, followed by the product the error occurred in, followed by the original token list. If the **U** entry is not included, all unlisted codes are passed without any mapping.

9. All undefined positive SQLCODEs are mapped to +965. This option should be used only if all mappable codes are listed, including all those that are identical and require no mapping. The **(s)** option indicates that the token list to be returned in the SQLERRMC field of the SQLCA includes the original SQLCODE, followed by the product the warning occurred in, followed by the original token list. If the **P** entry is not included, all unlisted positive codes are passed without any mapping.

Chapter 10. Problem Determination

The DB2 Connect environment involves multiple software, hardware and communications products. Problem determination is best approached by a process of elimination and refinement of the available data to arrive at a conclusion (the location of the error).

The following topics are provided to assist in the problem determination process:

- "Other Information Sources"
- "Gathering Relevant Information" on page 96
- "Initial Connection is Not Successful" on page 97
- "Problems Encountered after an Initial Connection" on page 98
- "Diagnostic Tools" on page 99
- "Most Common DB2 Connect Problems" on page 100.

After gathering the relevant information and based on your selection of the applicable topic proceed to the referenced section.

Other Information Sources

This section lists additional sources of information.

Using the Troubleshooting Guide

For more information on DB2 Connect and DB2 Universal Database problem determination topics, please refer to the *Troubleshooting Guide.*

Using the World Wide Web

You can find the most recent information about DB2 Connect problem determination hints and tips in the DB2 Product and Service Technical Library on the World Wide Web:

1. Set your Web browser to the following URL:

   ```
   http://www.software.ibm.com/data/db2/library/
   ```
2. Select "DB2 Universal Database".
3. Search for "Technotes" using the keywords "DDCS", or "Connect".

APPC, CPI-C, and SNA Sense Codes Documentation

Documentation on APPC, CPI-C, and SNA Sense Codes that was provided has been repackaged and it is provided as a PostScript file, and an HTML book (English only).

The PostScript version of this book can be found on your DB2 Version 6 CD in the ″\doc\en\ps″ directory, The file name is as follows:

- ″db2ape60.ps.Z″ on UNIX based operating systems (use uncompress to create the printable file)
- ″db2ape60.psz″ on OS/2 and Windows operating systems (use PKUNZIP).

The HTML copy of this book is always installed when DB2 Connect is installed, and this book is integrated with the DB2 Information Center.

Gathering Relevant Information

Problem determination includes narrowing the scope of the problem and investigating the possible causes. The proper starting point is to gather the relevant information and determine what you know, what data has not been gathered and what paths you can eliminate. At a minimum answer the following questions.

- Has the initial connection been successful?
- Is the hardware functioning properly?
- Are the communication paths operational?
- Have there been any communication network changes that would make previous directory entries invalid?
- Has the database been started?
- Is the communication breakdown between client and DB2 Connect workstation, DB2 Connect workstation and host or AS/400 database server, all clients or one client?
- What can you determine by the content of the message and the tokens returned in the message?
- Will using diagnostic tools provide any assistance at this time?
- Are other machines performing similar tasks working correctly?
- If this is a remote task, is it successful if performed locally?

Initial Connection is Not Successful

Review the following questions and ensure that the installation steps were followed.

1. *Did the installation processing complete successfully?*
   - Were all the prerequisite software products available?
   - Were the memory and disk space adequate?
   - Was remote client support installed?
   - Was the installation of the communications software completed without any error conditions?

2. *For UNIX-based systems was an instance of the product created?*
   - As root did you create a user and a group to become the instance owner and sysadm group?

3. *If applicable, was the license information processed successfully?*
   - For UNIX-based systems, did you edit the nodelock file and enter the password that IBM supplied?

4. *Were the host or AS/400 database server and workstation communications configured properly?*
   - There are three configurations that must be considered:
     a. The host or AS/400 database server configuration identifies the application requester to the server. The host or AS/400 server database management system will have system catalog entries that will define the requestor in terms of location, network protocol and security.
     b. The DB2 Connect workstation configuration defines the client population to the server and the host or AS/400 server to the client.
     c. The client workstation configuration must have the name of the workstation and the communications protocol defined.
   - Problem analysis for not making an initial connection includes verifying for SNA connections that all the LU (logical unit) and PU (physical unit) names are complete and correct, or verifying for TCP/IP connections that the correct port number and hostname have been specified.
   - Both the host or AS/400 server database administrator and the Network administrators have utilities available to diagnose problems.

5. *Do you have the level of authority required by the host or AS/400 server database management system to use the host or AS/400 server database?*
   - Consider the access authority of the user, rules for table qualifiers, the anticipated results.

6. *If you attempt to use the command line processor to issue SQL statements against a host or AS/400 database server, are you unsuccessful?*

- Did you follow the procedure to bind the command line processor to the host or AS/400 database server?

Problems Encountered after an Initial Connection

The following questions are offered as a starting point to assist in narrowing the scope of the problem.

1. *Are there any special or unusual operating circumstances?*
   - Is this a new application?
   - Are new procedures being used?
   - Are there recent changes that might be affecting the system? For example, have any of the software products or applications been changed since the application or scenario last ran successfully?
   - For application programs, what application programming interface (API) was used to create the program?
   - Have other applications that use the software or communication APIs been run on the user's system?
   - Has a PTF recently been installed? If the problem occurred when a user tried to use a feature that had not been used (or loaded) on their operating system since it was installed, determine IBM's most recent PTF level and load that level *after* installing the feature.

2. *Has this error occurred before?*

   Is there any documented resolution to previous error conditions?

   Who were the participants and can they provide insight into possible course of action?

3. *Have you explored using communications software commands that return information about the network?*

   Is there a verification tool available for your SNA software?

   If you are using TCP/IP there may be valuable information retrieved from using TCP/IP commands and daemons.

4. *Is there information returned in the SQLCA (SQL communication area) that can be helpful?*

   Problem handling procedures should include steps to examine the contents of the SQLCODE and SQLSTATE fields.

   SQLSTATEs allow application programmers to test for classes of errors that are common to the DB2 family of database products. In a distributed relational database network this field may provide a common base. For more information see the *Message Reference*

5. *Was DB2START executed at the Server?* Additionally, ensure that the DB2COMM environment variable is set correctly for clients accessing the server remotely.

6. *Are other machines performing the same task able to connect to the server successfully?* The maximum number of clients attempting to connect to the server may have been reached. If another client disconnects from the server, is the client previously not able to connect, now able to connect?

7. *Does the machine have the proper addressing?* Verify that the machine is unique in the network.

8. *When connecting remotely, has the proper authority been granted to the client?* Connection to the instance may be successful, but authorization not granted at the database or table level.

9. *Is this the first machine to connect to a remote database?* In distributed environments routers or bridges between networks may block communication between the client and the server. For example, when using APPC, ensure that a session can be established. When using TCP/IP, ensure that you can PING the remote host.

Diagnostic Tools

When you encounter a problem, you can use the following:

- The first failure service log, where diagnostic information is consolidated and stored in a readable format. For more information, see the *Troubleshooting Guide.* For information about messages found in the log, see the *Message Reference.*

- `db2diag.log`

  This file is located in `/u/db2/sqllib/db2dump/db2diag.log` on UNIX systems, where `db2` is the instance name.

  This file is located in `x:\sqllib\db2\db2diag.log` on Intel systems, where `x:` is the logical drive, and `db2` is the instance name.

- `db2alert.log` (Same file locations as `db2diag.log`).

- The trace utility, as described in "Trace Utility (ddcstrc)" on page 43.

- For UNIX-based systems, the **ps** command, which returns process status information about active processes to standard output.

- For UNIX-based systems, the core file that is created in the current directory when severe errors occur. It contains a memory image of the terminated process, and can be used to determine what function caused the error.

- For Windows NT systems, use the Event Viewer.

For more information on troubleshooting TCP/IP connections (or other topics), please refer to *Troubleshooting Guide*, or search for "Technotes" on the DB2 Product and Service Technical Library (see "Using the World Wide Web" on page 95.

Most Common DB2 Connect Problems

This section lists the most common symptoms of connection problems encountered when using DB2 Connect. In each case, you are provided with:

- A combination of a message number and a return code (or protocol specific return code) associated with that message. Each message and return code combination has a separate heading, and the headings are ordered by message number, and then by return code.
- A symptom is provided, usually in the form of a sample message listing.
- A suggested solution is provided, indicating the probable cause of the error. In some cases more than one suggested solution may be provided.

Notes:

1. Refer to your product Quick Beginnings manual and the latest Release Notes for the most up-to-date information about recommended software fix levels.

2. For message and return code combinations specific to APPC communications, an SNA sense code may also be indicated. At present, any SNA sense code information associated with a particular message must be obtained from the SNA subsystem.

   Sometimes SNA sense codes can be viewed by looking through system logs. Whether this is the case or not depends on the SNA subsystem being used, and in some situations you may have to recreate the problem with an SNA trace active in order to obtain the sense code information.

3. The term gateway refers to DB2 Connect Enterprise Edition.

SQL0965 or SQL0969

Symptom

Messages SQL0965 and SQL0969 can be issued with a number of different return codes from DB2 Universal Database for AS/400, DB2 Universal Database for OS/390, DB2 for MVS/ESA, and DB2 for VM & VSE.

When you encounter either message, you should look up the original SQL code in the documentation for the database server product issuing the message.

Solution

The SQL code received from the host database cannot be translated. Correct the problem, based on the error code, then resubmit the failing command.

SQL1338 During CONNECT

Symptom / Cause

The symbolic destination name was not defined, or it is not properly defined.

For example, this can happen when an APPC node is used and the symbolic destination name specified in the DB2 node directory does not match a CPI-C entry in the local APPC communications subsystem configuration.

Another cause can be that there is more than one SNA stack installed on your machine. You may need to check PATH and LIBPATH to ensure that the stack you want to use is referenced first.

Solutions
1. Make sure the CPIC Side Information profile name specified in the DB2 Node directory entry matches the SNA configuration (it is case sensitive).
2. You may need to check PATH and LIBPATH to ensure that the SNA stack that you want to use is referenced first.

SQL1403N During CONNECT

Symptom

SQL1403N The username and/or password supplied is incorrect.

Solution
1. User fails to authenticate at the DB2 Connect workstation. Determine whether the user is supposed to be authenticated at the DB2 Connect workstation.

   If yes, make sure that the correct password is provided on the CONNECT statement if necessary.

   If no, the system database directory entry must have been incorrectly cataloged using AUTHENTICATION SERVER (this is the default if AUTHENTICATION is not specified explicitly). If this is the case, then recatalog the entry using AUTHENTICATION DCS or CLIENT.
2. Password not available to send to the target server database. If the system database directory entry is cataloged using AUTHENTICATION DCS, then a password has to be flowed from the DB2 Client to the target server database. On certain platforms, for example AIX, the password can only be obtained if it is provided on the CONNECT statement.

SQL5043N

Symptom

Support for one or more communications protocols failed to start successfully. However, core database manager functionality started successfully.

Perhaps the TCP/IP protocol is not started on the DB2 Connect gateway. There may have been a successful client connection previously.

If `diaglevel = 4`, then `db2diag.log` may contain a similar entry, for example:
```
1997-05-30-14.09.55.321092   Instance:svtdbm5   Node:000
PID:10296(db2tcpcm)   Appid:none
common_communication  sqlcctcpconnmgr_child   Probe:46
DIA3205E Socket address "30090" configured in the TCP/IP
services file and
required by the TCP/IP server support is being used by another
process.
```

Solution

This warning is a symptom which signals that DB2 Connect, acting as a gateway for remote clients, is having trouble handling one or more client communication protocols. These protocols can be TCP/IP, APPC and others, and usually the message indicates that one of the communications protocols defined to DB2 Connect is not configured properly.

Often the cause may be that the DB2COMM profile variable is not defined, or is defined incorrectly. Generally, the problem is the result of a mismatch between the DB2COMM variable and names defined in the database manager configuration (for example, svcename, nname, or tpname).

One possible scenario is having a previously successful connection, then getting the SQL5043 error message, while none of the configuration has changed. This could occur using the TCP/IP protocol, when the remote system abnormally terminates the connection for some reason. When this happens, a connection may still appear to exist on the client, and it may become possible to restore the connection without further intervention by issuing the commands shown below.

Most likely, one of the clients connecting to the gateway still has a handle on the TCP/IP port. On each client machine that is connected to the gateway:
1. db2 terminate
2. db2stop

SQL30020

Symptom

SQL30020N Execution failed because of a Distributed Protocol Error that will affect the successful execution of subsequent commands and SQL statements.

Solutions

Service should be contacted with this error.

Check the db2dump directory for an ffdc dump (pid.000). Then, format this dump file with db2fdump and look in the result file for ″ERROR″. An MVS ABEND may be listed here. In this case check the MVS console for further information, and look up the abend code in the DB2 for MVS Messages and Codes manual.

SQL30060

Symptom

SQL30060N ″<authorization-ID>″ does not have the privilege to perform operation ″<operation>″.

Solution

When connecting to DB2 for MVS or DB2 for OS/390, the Communications Database (CDB) tables have not been updated properly. Refer to:
- DB2 Connect Quick Beginnings, or
- DB2 Connectivity Supplement

SQL30061

Symptom

Connecting to the wrong host or AS/400 database server location - no target database can be found.

Solution

The wrong server database name may be specified in the DCS directory entry. When this occurs, SQLCODE -30061 is returned to the application.

Check the DB2 node, database, and DCS directory entries. The target database name field in the DCS directory entry must correspond to the name of the database based on the platform. For example, for a DB2 Universal Database for OS/390 database, the name to be used should be the same as that used in

the Boot Strap Data Set (BSDS) ″LOCATION=locname″ field, which is also provided in the DSNL004I message (LOCATION=location) when the Distributed Data Facility (DDF) is started. See also "The Database Concept" on page 3, and "Chapter 3. Updating Database Directories" on page 19.

Your DB2 Connect `Quick Beginnings` manual also contains examples showing how to update the DB2 catalogs. See the ″Update the DB2 Directories″ section in each chapter describing SNA configuration, or see the ″Configuring Host and AS/400 Databases for DB2 Connect″ chapter, and the ″Configuring the TCP/IP Connection″ section.

The correct commands for an APPC or APPN node are:

```
db2 catalog appc node <node_name> remote <sym_dest_name> security program
db2 catalog dcs database <local_name> as <real_db_name>
db2 catalog database <local_name> as <alias> at node <node_name>
            authentication dcs
```

The correct commands for a TCP/IP node are:

```
db2 catalog tcpip node <node_name> remote <host_name_or_address>
            server <port_no_or_service_name>
db2 catalog dcs database <local_name> as <real_db_name>
db2 catalog database <local_name> as <alias> at node <node_name>
            authentication dcs
```

To connect to the database you then issue:

```
db2 connect to <alias> user <user_name> using <password>
```

SQL30073 with Return Code 119C During CONNECT

Symptom

Message SQL30073 is issued with return code 119C. This happens when the target server database does not support the code page used by the DB2 client (going through DB2 Connect). The code page is derived from the configuration of the operating environment in which the DB2 client is running.

See the *Administration Guide* for further information.

Solution

This problem can often be resolved by installing a fix at the target server database system. Contact the appropriate service organization and obtain and apply any fix which might be recommend for this symptom.

As a temporary workaround, the user can override the default code page by setting the DB2CODEPAGE environment variable. Check the locale or set DB2CODEPAGE=850.

On UNIX platforms, the user may be able to switch to a different code page by setting the LANG environment variable to a different value.

SQL30081N with Return Code 1

Symptom

Symptom is the following message plus an SNA sense code:

```
db2 connect to <database name> user <userid>
Enter password for <userid>:
SQL30081N  A communication error has been detected.
Communication protocol
being used: "APPC".  Communication API being used: "CPI-C".
Location where
the error was detected: "".  Communication function detecting
the error:
"cmallc".  Protocol specific error code(s): "1", "*",
"0x10030021".
SQLSTATE=08001
```

Solution(s)

In this sample the sense code is 10030021.

The most common sense codes associated with this error message, and the suggested solution in each case, are as follows:

1.
```
SQL30081N with return code 1 and sna sense code 0877002C
```

   Wrong network name has been specified.

2.
```
SQL30081N with return code 1 and SNA sense code ffff0003
```

   The wrong MAC address has been specified or the SNA link is not active.

3.
```
SQL30081N with return code 1 and SNA sense code 10030021
```

   There is an LU type mismatch.

4.
```
SQL30081N with return code 1 and SNA sense code 084B6031
```

   The MAXDBAT in DSNZPARM (at a DB2 for MVS or DB2 for OS/390 host) is set to 0

Other suggestions:

1. When creating the Local LU profile, define the LU as the default LU. For example, in the SNA Feature list panel in CM/2, either:
   - Place a checkmark in the checkbox 'Use this local LU as your default local LU alias', or
   - Set the profile or environment variable APPCLLU on the DB2 Connect Enterprise Edition gateway system to the Local LU name. You can do this on OS/2 systems, for example, by editing CONFIG.SYS, or on Windows NT systems via the Control Panel.

2. Check that SNA is started on the DB2 Connect gateway

3. If you are using DB2 for MVS or DB2 for OS/390, check that the Distributed Data Facility (DDF) address space is started and that DB2 is running.

SQL30081N with Return Code 2

Symptom

Message SQL30081N is received with Return Code 2 and SNA Sense Code 08120022.

Solution

The NUMILU parameter at the NCP (host end of the link) may be set to the default (0). Check this. Modify the NCP definition if necessary before retrying, after putting the change into effect.

SQL30081N with Return Code 9

Symptom

Symptom is the following message (the SNA sense code is not required in this case):

```
db2 connect to <database> user <userid>
SQL30081N  A communication error has been detected.
Communication protocol
being used: "APPC".  Communication API being used: "CPI-C".
Location where
the error was detected: "".  Communication function detecting
the error:
"cmsend".  Protocol specific error code(s): "9", "*",
"0x10086021".
SQLSTATE=08001
```

Solution

The problem is that the Transaction Program name (TPNAME) is not defined correctly on the DB2 Connect system. For example, you may have updated

your SNA configuration, but not yet verified it at the DB2 Connect gateway. Refer to the *DB2 Connect Enterprise Edition for OS/2 and Windows NT Quick Beginnings*, or *DB2 Connect Personal Edition Quick Beginnings* manual for further details.

SQL30081N with Return code 10

Symptom

Symptom is the following message (the SNA sense code is not required):

```
SQL30081N  A communication error has been detected.
Communication protocol
being used: "APPC".  Communication API being used: "CPI-C".
Location where
the error was detected: "".  Communication function detecting
the error:
"cmrcv".  Protocol specific error code(s): "10", "*", "*".
SQLSTATE=08001
```

Solution

Check to make sure that DB2 is correctly installed.

If you are using a DB2 Connect for OS/2 gateway, you may see the following if the TP name is not defined properly:

```
Protocol specific error code(s): "10", "*", "0x084C0000".
SQLSTATE=08001
```

For example, in CM/2, in this case it should be defined as follows:

```
Transaction program name        = 'tpname'    (user defined)
OS/2 program path and file name = notused
```

and (on the next CM/2 configuration screen)

```
Presentation type - background
Operation type - Queued, operator preloaded
```

SQL30081N with Return Code 20

Symptom

```
SQL30081N  A communication error has been detected.
Communication protocol
being used: "APPC".  Communication API being used: "CPI-C".
Location where
the error was detected: "".  Communication function detecting
the error:
"xcstp".  Protocol specific error code(s): "20", "*", "*".
SQLSTATE=08001
```

Solution

Ensure that the SNA subsystem is started on the DB2 Connect system.

SQL30081N with Return code 27

Symptom

Message SQL30081N is received with Return Code 27 and SNA Sense Code 800Axxxx.

Solution

The VTAM Path Information Unit (PIU) is too large.

SQL30081N with Return Code 79

Symptom

```
SQL30081N  A communication error has been detected.
Communication protocol
being used: "TCP/IP".  Communication API being used: "SOCKETS".
Location
where the error was detected: "".  Communication function
detecting the error:
"connect".  Protocol specific error code(s): "79", "*", "*".
SQLSTATE=08001
```

Solution(s)

This error can occur in the case of a remote client failing to connect to a DB2 Connect gateway. It can also occur when connecting from the DB2 Connect gateway to a host.

1. The `DB2COMM` profile variable may set incorrectly on the DB2 Connect gateway. Check this. For example, the command `db2set db2comm=tcpip` should appear in `sqllib/db2profile` when running DB2 Extended Enterprise Edition on AIX.

2. There may be a mismatch between the TCP/IP service name and/or port number specifications at the DB2 client and the DB2 Connect gateway. Verify the entries in the TCP/IP `services` files on both machines.

3. Check that DB2 is started on the DB2 Connect gateway. Set the Database Manager Configuration `diaglevel` to 4, using the command:

       db2 update dbm cfg using diaglevel 4

   After stopping and restarting DB2, look in the `db2diag.log` file to check that DB2 TCP/IP communications have been started. You should see output similar to the following:

```
1998-02-03-12.41.04.861119   Instance:svtdbm2   Node:00
PID:86496(db2sysc)   Appid:none
common_communication  sqlcctcp_start_listen   Probe:80
DIA3000I "TCPIP" protocol support was successfully started.
```

SQL30081N with Protocol Specific Error Code 10032

Symptom

```
SQL30081N  A communication error has been detected.
Communication protocol
being used: "TCP/IP".  Communication API being used: "SOCKETS".
Location
where the error was detected: "9.21.85.159".  Communication
function detecting
the error: "send".  Protocol specific error code(s): "10032",
"*", "*".
SQLSTATE=08001
```

Solution

This error message may be received when trying to disconnect from a machine where TCP/IP communications have already failed. Correct the problem with the TCP/IP subsystem.

On most machines, simply restarting the TCP/IP protocol for the machine is the way to correct the problem. Occasionally, recycling the entire machine may be required.

Chapter 11. DB2 Connect Performance

The DB2 Connect product interacts with many different products including DRDA application server products, client products, and communication products. Its performance depends on all these pieces working together efficiently.

Note: This chapter assumes that all applications are run on the DB2 Connect workstation. If you have remote clients, you must also consider the performance of the remote connections to the DB2 Connect workstation.

Other Information Sources

This section lists additional sources of information.

Other Publications

The following books contain additional information about performance:
- *DDCS for OS/2 to DB2 Performance Benchmark*
- *SNA Server for AIX and SNA Server Gateway for AIX Performance Guide*

Using the World Wide Web

You can find extensive information about DB2 Connect performance tuning, as well as case studies and examples, on the World Wide Web. Set your Web browser to the following URL:

```
http://www.software.ibm.com/data/db2/performance/
```

Additional Hints and Tips for SNA Users

See "Appendix H. Additional SNA Performance Tuning Hints and Tips" on page 163.

Performance Concepts and Tools

Performance is the way a computer system behaves given a particular workload. It is affected by the resources available and how they are used and shared.

If you want to improve performance, you must first decide what you mean by performance. You can choose many different *performance metrics*, including:

Response time
> The interval between the time that the application sends the database request and the time that the application receives a response.

Transaction throughput
> The number of units of work that can be completed per unit of time. The unit of work could be simple, like fetching and updating a row, or complicated, involving hundreds of SQL statements.

Data transfer rate
> The number of bytes of data transferred between the DB2 Connect application and the host or AS/400 database per unit of time.

Performance will be limited by the available hardware and software resources. CPU, memory and network adapters are examples of hardware resources. Communication subsystems, paging subsystems, `mbuf` for AIX, and `link` for SNA are examples of software resources.

Data Flows

Figure 8 on page 113 shows the path of data flowing between the host or AS/400 database server and the workstation through DB2 Connect.
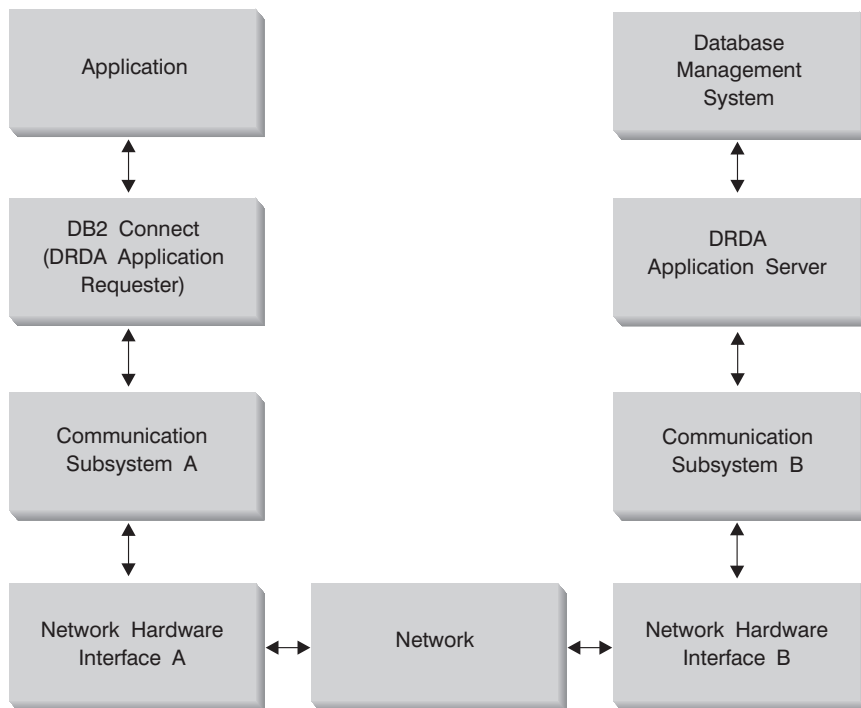
Figure 8. Data Flows in DB2 Connect

- The host or AS/400 database and part of communication subsystem B are usually running on the same system. This system is made up of one or more CPUs, main storage, an I/O subsystem, DASD, and an operating system. Because other programs may share these components, resource contention could cause performance problems.

- The network is composed of a combination of cables, hubs, communication lines, switches, and other communication controllers. For example, the network hardware interface B could be communication controllers such as 3745 or 3172 or a token ring adapter for an AS/400. There could be more than one transmission medium involved between network hardware interfaces A and B.

- Network hardware interface A could be token ring, Ethernet**, other LAN adapter, or an adapter which supports the SDLC or X.25 protocols. Communication subsystem A might be a product such as IBM Communications Server for OS/2, Microsoft SNA Server, IBM SNA Server for AIX, or SNAplus2 for HP-UX.

- The DB2 Connect product and the communication subsystem A are usually located on the same system. In this chapter, we assume that the application is also on the same system.

Bottlenecks

Transaction throughput is dependent on the slowest component in the system. If you identify a performance bottleneck, you can often alleviate the problem by changing configuration parameters, allocating more resources to the problem component, upgrading the component, or adding a new component to offload some of the work.

You can use various tools to determine how much time a query spends in each component. This will give you an idea of which components should be tuned or upgraded to improve performance. For example, if you determine that a query spends 60% of its time in the DB2 Connect machine, you might want to tune DB2 Connect or (if you have remote clients) add another DB2 Connect machine to the network.

For more information about performance tools, see "Performance Tools" on page 115.

Benchmarking

Benchmarking is a way to compare performance in one environment with performance in another.

Benchmarking can begin by running the test application in a normal environment. As a performance problem is narrowed down, specialized test cases can be developed to limit the scope of the function that is tested and observed.

Benchmarking does not need to be complex. Specialized test cases need not emulate an entire application in order to obtain valuable information. Start with simple measurements and increase the complexity only when warranted.

Characteristics of good benchmarks (or measurements) include:
- Each test is repeatable.
- Each iteration of a test is started in the same system state.
- The hardware and software used for benchmarking matches your production environment.
- There are no functions or applications active in the system other that those being measured (unless the scenario includes some amount of other activity going on in the system).

  Note: Applications that are started use memory even when they are minimized or idle. This could cause paging and skew the results of the benchmark.

Performance Tools

The following table lists some of the tools that can help you measure system performance. Because these tools themselves use system resources, you might not want to have them active all the time.

Table 7. Performance Tools

System	Tool	Description.
CPU and memory usage		
AIX	vmstat, time, ps, tprof	Provide information about CPU or memory contention problems on the DB2 Connect workstation and remote clients
HP-UX	vmstat, time, ps, monitor and glance if available	
OS/2	SPM/2, THESEUS/2, pstat	
Win NT	MS Performance Monitor	
Database activity		
All	Database monitor	Determines if the problem originates from the database.
MVS or OS/390	DB2PM (IBM), OMEGAMON/DB2 (Candle), TMON (Landmark), INSIGHT (Goal Systems) and DB2AM (BMC)	
Win NT	MS Performance Monitor	
Network activity		
AIX	netpmon	Reports low level network statistics, including TCP/IP and SNA statistics such as the number of packet or frames received per second.
DOS or OS/2	Token-Ring Network 16/4 Trace and Performance Program	Most network monitors are platform dependent; this tool works for token-ring only.
Network controller such as 3745	NetView Performance Monitor	Reports utilization of communication control and VTAM

Table 7. Performance Tools (continued)

System	Tool	Description.
OS/2	DatagLANce	A trace tool that presents performance-related data to users graphically.
UNIX-based	netstat	Handles TCP/IP traffic

Application Design

When you create an application, you can improve performance in several ways, including:

- Use compound SQL and stored procedures.
- Group requests.
- Use predicate logic to request only the data that you need.
- Use data blocking.
- Use static SQL whenever possible.

Compound SQL and Stored Procedures

For applications that send and receive many commands and replies, network overhead can be significant. Compound SQL and stored procedures are two ways to reduce this overhead.

If an application sends several SQL statements without intervening programming logic, you can use compound SQL. If programming logic is needed within the group of SQL statements, you can use stored procedures.

All executable statements except the following can be contained within a Compound SQL statement:

```
CALL
FETCH
CLOSE
OPEN
Compound SQL
Connect
Prepare
Release
Describe
Rollback
Disconnect
Set connection
execute immediate
```

See *SQL Reference* for more details.

For information about using compound SQL in an application, see "NOT ATOMIC Compound SQL" on page 81. For information about using compound SQL with the import utility, see "Using Import and Export Utilities" on page 41

Stored procedures help to reduce network traffic by placing program logic at the server. In DB2 prior to Version 5.0, a stored procedure could return output parameters only, and a separate commit command had to be issued by the application. This resulted in two network trips. In DB2 Versions 5.0 and above, you can commit automatically when exiting the procedure. You can also return results sets, which minimize application logic at the client.

For information about using stored procedures, see "Stored Procedures" on page 80.

Grouping Requests

Grouping related database requests (SQL statements) into one database request can reduce the number of requests and responses transmitted across the network. For example, grouping the following statements:

```
SELECT COL1, COL2, COL5, COL6 FROM TABLEA WHERE ROW_ID=1
SELECT COL1, COL2, COL5, COL6 FROM TABLEA WHERE ROW_ID=2
```

into

```
SELECT COL1, COL2, COL5, COL6 FROM TABLEA WHERE ROW_ID=1 OR ROW_ID=2
```

sends fewer requests across the network.

You can also use keywords such as IN and BETWEEN to reduce the number of rows returned. In addition, you can use WHERE, IN, and BETWEEN keywords on UPDATE and DELETE statements.

Predicate Logic

You can use predicate logic to request only the rows and columns that are needed. This minimizes the network traffic and CPU overhead for data transmission.

For example, do not use the query:

```
SELECT * FROM TABLEA
```

if only the first row of TABLEA with ROW_ID=1 is really needed or if only column 1 and column 2 are needed.

Data Blocking

You should use data blocking if you expect large amounts of data from the server. Blocking improves the utilization of the network bandwidth and reduces the CPU overhead of both the host or AS/400 database server and the DB2 Connect workstation.

There is fixed amount of CPU and network overhead for each message sent and received regardless of size. Data blocking reduces the number of messages required for the same amount of data transfer.

With blocking, the first row of data from a query will not be delivered to the application until the first *block* is received. Blocking increases the retrieval time for the first row, but improves the retrieval time for subsequent rows.

Another consideration is the amount of memory that is used. The memory working set usually increases when blocking is turned on. For a complete discussion of blocking when using SNA connections, see the *DRDA Connectivity Guide.*

Within DB2 Connect, you can control the amount of data that is transferred within each block, as described in "RQRIOBLK" on page 120.

To invoke blocking, use the BLOCKING option of the prep or bind command. (For more information, see "The BIND Command" on page 38.) Blocking is on, if:

- The cursor is read-only, or
- The cursor is ambiguous and blocking is specified during the prep or bind.

For the definitions of read-only, updateable, and ambiguous cursor, refer to the *Application Development Guide.*

Note: When using dynamic SQL, the cursor is always ambiguous.

SQL Statements with BLOCKING

Updateable SELECT statements (using `UPDATE/DELETE WHERE CURRENT OF` statements) are non-blocking queries, so you should use them only when absolutely necessary.

An updateable SELECT ensures that the row has not changed between the time the `SELECT` is completed and the `UPDATE/DELETE` is issued. If this level of concurrency is not important to your application, an alternative is to use a `DELETE` or `UPDATE` with search criteria based on the values returned from a non-updateable SELECT.

For read-only `SELECT`, specify `FOR FETCH ONLY` (except under VM and VSE, where it is not supported).

Static and Dynamic SQL

Use static SQL as much as possible. It avoids run-time SQL section preparation and ambiguous cursors. If dynamic SQL cannot be avoided, you can do the following to minimize the network traffic and improve performance:

- If the statement is a `SELECT` and must be prepared, perform `PREPARE ... INTO` SQLDA. The SQLDA should be allocated to the full size needed for your settings. If the maximum number of columns is x and is expected to stay that way, allocate an SQLDA with x SQLVARs. If the number of potential columns is uncertain (and memory is not a problem), use the maximum number of SQLVARs (256).

  If the SQLDA allocation is not big enough to store the returning SQLDA, the program must issue another DESCRIBE with a big enough SQLDA to store the result again. This would increase the network traffic.

  Do not use the PREPARE and DESCRIBE sequence. Using the `PREPARE.....INTO` statement provides better performance.

- Execute statically bound SQL COMMIT or ROLLBACK statements instead of dynamic COMMIT or ROLLBACK statements.

- If it is not a SELECT, COMMIT, or ROLLBACK statement, issue EXECUTE IMMEDIATE to execute the statement instead of the PREPARE and EXECUTE sequence.

Other SQL Considerations

Using the command line processor is, in general, slower than having dynamic SQL in the program because the command line processor must parse the input before submitting the SQL to the database engine. The command line processor also formats data when it is received, which may not be necessary for your application.

SQL statements in an interpreted language (such as REXX) are substantially slower than the same SQL statements in a compiled language (such as C).

There are two types of CONNECT statement, called type 1 and type 2. With type 2 connect, connecting to a database puts the previous connection into a dormant state but does not drop it. If you later switch to a dormant connection, you avoid the overhead of loading libraries and setting up internal data structures. For this reason, using type 2 connect may improve performance for applications that access more than one database. For more information about type 2 connects, see the *Administration Guide* and the *SQL Reference.*

DB2 Connect Tuning

Various parameters in the database manager configuration file can be used to tune DB2 Connect. For information about changing these parameters, see the *Administration Guide.*

RQRIOBLK

The RQRIOBLK parameter sets the maximum size of network I/O blocks. A larger block size may improve the performance of large requests. The block size does not usually affect the response time for small requests, such as a request for a single row of data.

A larger block size usually requires more memory on the DB2 Connect workstation. This increases the size of the working set and may cause large amounts of paging on small workstations.

Use the default DRDA block size (32767) if it does not cause too much paging on executing your application. Otherwise, reduce the I/O block size until there is no paging. Once paging begins, a noticeable degradation of performance will occur. Use performance monitor tools (such as the vmstat tool for UNIX-based systems or SPM/2 for OS/2) to determine whether paging is occurring on your system. For other tools, refer to "Performance Tools" on page 115.

DIR_CACHE

The DIR_CACHE parameter determines whether directory information is cached. With caching (DIR_CACHE=YES), directory files are read and cached in memory to minimize the overhead of creating the internal directory structure and reading the directory files every time a connection is established.

Without caching (DIR_CACHE=NO), whenever you connect to a database the appropriate directory is read from a disk and then the search is performed. After the requested entries are found, all memory related to directory searches is freed.

With caching, a shared directory cache is built during **db2start** processing and freed when DB2 stops. This cache is used by all DB2 server processes (db2agent). Also, a private application directory cache is built when an application issues its first connect to a database and freed when the application ends.

Each cache provides an image of the system database directory, the database connection services directory and the node directory. The cache reduces connect costs by eliminating directory file I/O and minimizing directory searches.

If a cached directory is updated, the changes are not immediately propagated to the caches. If a directory entry is not found in a cache, the original directory is searched.

Caching increases the private memory that is needed for the life of an application. Without caching, this memory is needed only when a directory lookup is processed. Overall use of shared memory by DB2 increases slightly because directory information that is shared among database agents is moved to shared memory. The size of the memory required for a cache depends on the number of entries defined in each directory.

Other DB2 Connect Parameters

MAXDARI and NUMDB should be set to their minimum values if there is no local database on the DB2 Connect workstation. This will minimize resource consumption.

AGENTPRI applies only with remote clients. AGENTPRI controls the priority given by the operating system scheduler to agents of a DB2 Connect instance. The DB2 Connect instance is granted more CPU cycles if it has a higher priority (lower number). This reduces the number of CPU cycles left for other processes executing on the DB2 Connect workstation. For example, you could have a high-priority DB2 Connect instance and a low-priority DB2 Connect instance running on the same workstation with different AGENTPRI values.

Every connection from a client machine to a host or AS/400 database server through DB2 Connect requires an agent running on the DB2 Connect workstation. Set MAXAGENTS to a value greater than or equal to the peak number of remote client connections accessing a host or AS/400 database server through the DB2 Connect workstation.

If you decide to use accounting strings, using the sqlesact() API has performance advantages over the DB2ACCOUNT environment variable method. For more information, see "Implementing Chargeback Accounting" on page 84.

If you do not need a tailored SQLCODE mapping file, you can improve performance by using the default SQLCODE mapping or turning off SQLCODE mapping. (The default mapping file is imbedded in the DB2 Connect library; a tailored mapping file must be read from disk, which affects

performance.) For more information about SQLCODE mapping, see
"Chapter 9. SQLCODE Mapping" on page 89.

Database Tuning

System performance will be affected by the performance of the host or AS/400
database server database.

Different database management systems have different performance features.
SQL optimizers of different systems, for example, could behave differently
with the same application. Check your host or AS/400 database server system
performance documentation for more information.

For DB2 Universal Database for AS/400, you may be able to improve
performance by using the uncommitted read (UR) or no commit (NC) bind
options to avoid journaling. Please note that when using UR, unjournalled
data can only be read, not updated, and then only if blocking is set to ALL.

Depending on the Application Server and the lock granularity it provides, the
isolation level used for a query or application may have a significant effect on
performance.

The database should have the appropriate level of normalization, effective use
of indexes, and suitable allocation of database space. Performance can also be
affected by the data types that you use, as described in the following sections.

Tuning DB2 for OS/390

You will find the best performance in OS/390 V2R5 and later versions.
TCP/IP support was introduced in OS/390 R3, so this is the base release level
if you wish to use TCP/IP.

The Distributed Data Facility (DDF) is responsible for connecting distributed
applications to DB2 for OS/390. The DDF should be set up as an application
server. In order to do this, you can either insert the LU name of the remote
system into the `SYSIBM.LUNAMES` table, or insert the `LUNAME`, `SYSMODENAME`,
`USERSECURITY`, `ENCRYPTPSWDS`, `MODESELECT`, and `USERNAMES` values into the
`SYSIBM.SYSLUNAME` table. Then perform a DDF update to the Boot Strap Data
Set (BSDS), for instance as follows:

```
DDF LOCATION=LOC1,LUNAME=LU1,PORT=8000,RESPORT=8001
```

For best performance, you should use the recommended DDF address space
prioritization (slightly lower or equal to DBM1 if you are in `COMPAT` mode).

Use RACF caching of authorizations in VLF, and use V5 package authorizations caching if you can. A value of CACHEPAC=32768 is sufficient for most operations.

Since DDF will try to connect to VTAM, VTAM must be active when DDF starts. A sample VTAM APPL definition is included below:

```
SYD51TC* APPL  AUTH=(ACQ),                              X
               PARSESS=YES,                             X
               HAVAIL=YES,                              X
               EAS=1600,                                X
               APPC=YES,                                X
               DSESLIM=1024,                            X
               DMINWNL=512,                             X
               DMINWNR=512,                             X
               AUTOSES=1,                               X
               SECACPT=ALREADYV,                        X
               SRBEXIT=YES,                             X
               SYNCLVL=SYNCPT,                          X
               MODETAB=DB2MODET,                        X
               VPACING=63                               X
```

You can optimize inactive thread processing in OS/390. In V3, you are allowed up to 10,000 concurrently connected clients, and up to 25,000 in V4 and V5. In all cases, the maximum number that can be concurrently active, however, is 1999. Each workstation client can stay connected when it is inactive; its thread is placed on an inactive chain at each commit.

The DSNZPARM parameters CMSTAT, CONDBAT and MAXDBATM affect thread processing. For best performance, set CMSTAT to INACTIVE, adjust CONDBAT to the maximum number of connected DBATs that provide good performance, and MAXDBAT to the maximum acceptable number of active DBATs.

For a complete discussion on connecting DB2 for OS/390 in a DRDA network, including VTAM configuration, refer to the on-line *Connectivity Supplement.*

Data Conversion

When data is transferred from one environment to another, it may need to be converted. This conversion can affect performance.

Consider the following platforms:
- Intel (OS/2 or Windows NT)
- IEEE (UNIX-based systems)
- System/370 and System/390 (MVS, OS/390, VM, and VSE)
- OS/400.

and the following types of numeric data:
- Packed decimal

- Zoned decimal
- Integer
- Floating point.

Table 8 shows when conversion takes place.

Table 8. Data Conversion

	Intel	IEEE	S/370 & S/390	OS/400
Packed decimal data				
Intel	No	No	No	No
IEEE	No	No	No	No
S/370/390	No	No	No	No
OS/400	No	No	No	No
Zoned decimal data				
Intel	No	No	Yes	Yes
IEEE	No	No	Yes	Yes
S/370/390	Yes	Yes	No	No
OS/400	Yes	Yes	No	No
Integer data				
Intel	No	Yes	Yes	Yes
IEEE	Yes	No	No	No
S/370/390	Yes	No	No	No
OS/400	Yes	No	No	No
Floating point data				
Intel	No	Yes	Yes	Yes
IEEE	Yes	No	Yes	No
S/370/390	Yes	Yes	No	Yes
OS/400	Yes	No	Yes	No

The CPU cost of single-byte character data conversion is generally less than that of numeric data conversion (where data conversion is required).

The data conversion cost of DATE/TIME/TIMESTAMP is almost the same as that of single-byte CHAR. FLOATING point data conversion costs the most. The application designer may want to take advantage of these facts when designing an application based on DB2 Connect.

If a database table has a column defined 'FOR BIT DATA', the character data being transferred between the application and the database does not require any data conversion. This can be used when you are archiving data on the host or AS/400 database server.

Data Types for Character Data

Character data can have either the CHAR or VARCHAR data type. Which data type is more efficient depends on the typical length of data in the field:

- If the size of actual data varies significantly, VARCHAR is more efficient because CHAR adds extra blank characters to fill the field. These blank characters must be transmitted across the network like any other characters.
- If the size of actual data does not vary much, CHAR is more efficient because each VARCHAR field has a few bytes of length information which must be transmitted.

Network Tuning

The best way to improve overall performance in a distributed database environment is to eliminate delays from the network. It is common for network administrators to consider a network to be more efficient if it collects as much data as possible between transmissions. This approach doesn't work for applications such as distributed databases because it builds delays into the network. The end-user doesn't see the efficiency of the network, only the delays.

Most network devices have delay parameters, and most of them default to values that are very bad for distributed databases. To improve performance you should locate these parameters and if possible, set them to zero. In addition you should ensure that the buffer size on the device is large enough to prevent retransmits due to lost data. For instance, UNIX systems typically have a Transmit or Receive queue depth default of 32. For better results, set the queue depth to 150. A corresponding parameter on DLC settings is the Receive Depth, which should also be 150.

The IOBUF parameter is set too low at most sites. It is usually set at 500, but experience has shown that a value of 3992 works best if you are moving large amounts of data, especially for channel connections such as ESCON or 3172.

For SNA connections, you should set the Mode Profile of any workstation software to 63. In general, receive pacing values throughout the network should be set to their highest value, so the VPACING and PACING parameters on the DB2 APPL statement, and the PU/LU for the workstation in a switched major mode should also be set to 63. What this will do is allow the amount of message flows before the sender must wait for a response to increase progressively.

On a LAN system the DLC or LLC transmit and receive window sizes can have a dramatic effect on performance. The send value should be set to seven or more, and for most configurations a receive value of four or less works best.

If you are running Ethernet, you should set the TCP segment size to 1500 bytes. On a token ring or FDDI network this value should be 4400 bytes, and if you are using an ESCON adapter with TCP/IP, the segment size should always be 4096.

Finally, for TCP/IP networks, the TCP Send and Receive buffer sizes should be set higher than 32768. A value of 65536 is generally best.

Note:

> Establishing a connection from the gateway to the server (outbound connection) is much more expensive than establishing a connection from a client to the gateway (inbound connection). In an environment where thousands of clients frequently connect to and disconnect from the server through the gateway, a substantial amount of processing time is spent establishing outbound connections. DB2 Connect provides connection pooling over TCP/IP. When a client requests disconnection from the server, the gateway drops the inbound connection with the client, but keeps the outbound connection to the server in a pool. When a new client comes into the gateway to request a connection, the gateway provides an existing one from the pool thus reducing the overall connection time and saving the high CPU connect cost on the server.
>
> For more information about connection pooling under DB2, refer to the DB2 Administration Guide.

A summary of network performance tuning methods is provided in the following table.

What to Look For	Example	Setting	Notes
Deliberate Delays	Delay parameters on network devices	Set to 0	Defaults are usually higher.
Buffers	IOBUF parameter	Set up to 3992.	Particularly useful for ESCON or other channel adapter.
	RUSIZE	Optimum size is 4096	Setting RUSIZE and RQRIOBLK to same size may give best performance.

	Pacing	VPACING, PACING, and Mode Profiles should be set to 63.	Use adaptive pacing where applicable.
Adapter Settings	Transmit/Receive queue depth	Recommended value is 150.	Default is usually 32.
	DLC Windowing on SNA	Set transmit window size high (>7). Set receive window size low (eg. at 1), test and increment repeatedly to find ideal value.	Every logical device adds delays. Simply network topology as much as possible.
TCP Settings	Segment Sizes	1500 on Ethernet, 4400 on token ring and FDDI.	ESCON adapters used for TCP/IP should always be set to 4096.
	Send/Receive Space Sizes	Should be 64K for both.	Default is only 8192 for Windows. Can be set in the Windows registry.

Network Hardware

The following considerations relate to the hardware:

- Speed of the network or transmission media

  Performance improves with a faster transmission medium. For example, the following are some typical raw data transfer rates:

  Channel-to-channel (fiber optics)
  	4.0 MB/s

  16 Mbps LAN
  	2.0 MB/s

  Channel-to-channel (regular)
  	1.0 MB/s

  4 Mbps LAN
  	0.5 MB/s

  high speed T1 carrier (1.544 Mbps)
  	0.193 MB/s

  fast remote 56 Kbps phone line
  	0.007 MB/s

19.6 Kbps modem
   0.002 MB/s

9600 bps modem
   0.001 MB/s

The data transfer rate is limited by the slowest transmission medium in the path to the host or AS/400 database server.

• Network adapter or communication controller

You should carefully plan the memory usage of the network adapter and communication controller. In addition, you should work with a network specialist to ensure that the controller has the capability to handle the extra traffic generated by DB2 Connect.

• Network topology

If data crosses from LAN to LAN, and from one SNA Network to another SNA Network, consider the travel time. Bridges, routers, and gateways will add to the elapsed time. For example, reducing the number of bridges that are crossed reduces the number of hops required for each request.

The physical distance between nodes should also be considered. Even if a message is transferred by satellite, the transfer time is limited by the speed of light (3 * 10**8 m/s) and the round-trip distance between the sender and receiver.

• Network traffic

If the bandwidth of the network has been fully utilized, both the response time and the data transfer rate for a single application will decrease.

Congestion can occur in the network when data accumulates at a particular part of the network; for example, at an old NCP with a very small buffer size.

• Network reliability

If the error rate of the network is high, the throughput of the network will decrease and this will cause poor performance because of data re-transmission.

Contention for System Resources

Performance could be degraded if many tasks in the system are contending for system resources. Consider the following questions:

• Is the CPU saturated? Consider upgrading the system, reducing the system workload, and tuning the system to reduce processing overhead.

• Is the memory over-committed? Consider upgrading memory, reducing system workload and tuning the system to reduce the memory working set.

- Is the communication adapter/communication controller too busy? Consider upgrading the network or pairing up token-ring cards.
- Is one of the subsystems too busy, and is this subsystem on the data path?
- Are any unnecessary processes or tasks running on the system? The general rule is not to configure or start services unless they are used regularly since they will waste system resources.
- Do a few processes or tasks use most of the resource? Can they be stopped? Can their priorities be reduced? Can they be refined so that they don't use as much resource?

Performance Troubleshooting

If DB2 Connect users are experiencing long response times during large queries from hsot or AS/400 servers, the following areas should be examined for the possible cause of the performance problem:

1. For queries which result in returning large data blocks from the host or AS/400 server (usually 32K of data and above), ensure that the database manager configuration parameter RQRIOBLK is set to 32767. This can be done using the Command Line Processor (CLP) as follows:

   ```
   db2 update database manager configuration using RQRIOBLK 32767
   ```

2. If VTAM is used in the connection to the host or AS/400 server, look under "switched major node" configuration for the value of the PACING parameter. On the DB2 Connect workstation, examine the communication setup of the "LU 6.2 Mode Profile" for IBMRDB mode definition. In this definition, ensure the value for the "Receive pacing window" parameter is less than or equal to the PACING value defined on VTAM. A common value for "Receive pacing window" on the DB2 Connect workstation and "PACING" on VTAM is 8.

3. Ensure the maximum RU size defined in the IBMRDB mode definition is set to a suitable value. We recommend not less than 4K for connections using Token-ring hardware. For connections using Ethernet hardware, note the maximum Ethernet frame size of 1536 bytes, which may be a limiting factor.

4. Consult with the VTAM administrator in your environment to ensure that VTAM is using "adaptive pacing" in LU-LU sessions with your DB2 Connect workstation.

Appendix A. Sending Accounting Information to a DB2 for OS/390 Server

Many mainframe customers consider detailed accounting for resources used by different applications to be an important part of their operational procedures. DB2 for OS/390 provides extensive facilities for producing accounting reports that allow information systems departments to charge individual user departments for the mainframe resources they use. This is often called ″charge-back accounting.″ DB2 Connect products allow for accurate accounting for the host resources used by the PC and UNIX applications using existing accounting reports and procedures.

DB2 Connect implements this feature in a flexible way by providing:

- A default accounting string for all usage generated by a particular DB2 Connect Enterprise Edition server.
- Two mechanisms for individual users or applications to specify the account to which their usage should be charged.

Setting the Accounting String

The default accounting string is set by the DB2 Connect workstation's dft_account_str configuration parameter. This default mechanism is useful for database clients that do not have the capability to forward an accounting string to DB2 Connect (for example, applications developed prior to Version 2 products).

There are two ways for client applications to override the default accounting string set at the DB2 Connect server:

- Using the Set Accounting String API: **sqlesact()**

  The SQLESACT() API is called before the application connects to a database. You should use this method because:
  - Calling an API does not incur the cost of reading a registry value.
  - You do not need to call this API again unless you want to use a new accounting string for future connect requests. If you are using the registry value, you need to end the application process, redefine DB2ACCOUNT, and then restart the process.

  Refer to the *Administrative API Reference* for details on using this API.
- Using the DB2ACCOUNT registry value at the client workstation.

131

If the **sqlesact()** API is not called prior to the first database connect request, the DB2ACCOUNT registry value is read. The accounting string is used for all subsequent database connect requests.

Note: When defining the accounting string, you should observe the following rules:

1. Use the characters A-Z, 0-9, or '_' (underscore).
2. Limit the accounting string to 199 bytes—longer strings are truncated.

Appendix B. Changing Your MVS Password

DB2 Connect now provides an ability to change user passwords. This facility is especialy usefull for situations where host security service such as RACF is used to authenticate users. Previously changing host password would require users to log in to a TSO session to change their password. With the new password maintenace support provided by the DB2 Connect products users can issue SQL CONNECT statement from DB2 Command Line Processor (CLP), use PASSWORD button on DB2 Client Configuration Assistant (CCA), or press CHANGE button on the ODBC login dialog to change their host password.

Changing MVS passwords from DB2 Connect workstations connected to DB2 for OS/390 V5.1 via TCP/IP requires that the DB2 OS/390 Extended Security Field be set to ″Yes″. This field appears in the DB2 OS/390 DSNTIPR panel.

Changing of MVS passwords on host systems conencted via SNA requires that a special password expiration management program be set up on the host and that DB2 Connect workstation be configured to communicate with this host program.

The host password expiration management program is provided as part of the following MVS program products:
* MVS/ESA SP Version 4.2 or higher (password expiration management is a part of the APPC/MVS component)
* CICS/ESA Version 3.3 or higher

and has IBM Resource Access Control Facility (RACF) 1.9.2 installed.

You need to:
1. Configure the host's transaction program to receive your requests for password expiration maintenance.
2. Configure your DB2 Connect workstation for communications with the host transaction program.

Configuring the Host for Password Expiration Management

Setting up Password Expiration Management facilities on MVS hosts is the responsibility of MVS system programmers and is beyond the scope of this book. The following publications provide detailed documentation:
* *MVS/ESA SP V5.x Planning: APPC Management*, GC28-1503

- *MVS/ESA SP V5.x APPC/MVS Handbook for the OS/2 Administrator*, GC28-1504
- *MVS/ESA SP V4.3 Planning: APPC Management*, GC28-1110
- *MVS/ESA SP V4.2 APPC/MVS Handbook for the OS/2 Administrator*, GC28-1133
- *CICS/ESA: RACF Security Guide (CICS 4.1)*, SC33-1185
- *CICS/ESA: CICS-APPC Password Expiration Management Guide (CICS 3.3)*, SC33-0921
- *APPC Security MVS/ESA CICS/ESA OS/2*, GG24-3960.

Configuring the DB2 Connect Workstation for Password Expiration Management

Once the host password expiration management transaction program is configured you will need to configure your DB2 Connect workstation to communicate with the host program. This configuration involves two steps:

1. Define a symbolic destination name for the host password maintenance program in the SNA subsystem on your DB2 Connect workstation.
2. Record this symbolic destination name in the DCS directory for the databases that reside on this host system.

Step 1. Define Symbolic Destination Name

How you define symbolic destination name for the host password expiration management program depends on what SNA subsystem you are using:

- If you are using an SNA subsystem that can be configured by the DB2 Client Configuration Assitant (CCA) then you should use CCA to configure this symbolic destination name. You will need to obtain the LU name for the password expiration management program from your MVS administrator.
- If your SNA subsystem can not be configured using DB2 CCA you should follow documentation provided with your SNA subsystem to configure the symbolic destination name. You will need to obtain the following information from your MVS administrator:
  - Network name for the host that you are connecting to
  - LU name for the host password expiration management program

When configuring symbolic destination name you will also need to specify x'06F3F0F1' (hexadecimal number) for the Transaction Program (TP) name and set security to ″NONE″. You can specify mode such as #INTER or any other mode that your MVS may suggest to you.

Step 2. Record Symbolic Destination Name in the DCS Directory

If you are running DB2 Connect on a platform that provides DB2 CCA then you should use DB2 CCA to update your DCS directory with the symbolic destination name for the host password expiration management program. You should be able to do this regardless of the SNA subsystem on your DB2 Connect workstation.

You can also use `catalog dcs database` command (from DB2 CLP) to record the symbolic destination name in the DCS directory. For example:

```
catalog dcs database db1 as dsn_db_1 parms ",,,,,,,CHGPWD_SDN=pempgm"
```

records ″pempgm″ as the symbolic destination name that is to be used when users request to change passwords for database ″db1″.

Appendix C. Directory Customization Worksheet

Use this worksheet to customize your directories. Refer to "Updating the Directories" on page 30 or the *Command Reference* for command syntax.

Table 9. Node Directory Parameters

Parameter	Example	Your value
Node name	DB2NODE or MVSIPNOD	
Symbolic destination name (APPC node)	DB2CPIC	
Remote hostname (TCP/IP node)	MVSHOST	
Server (TCP/IP service name or port number)	db2inst1c (or port number)	
Security type	PROGRAM for APPC Nodes; NONE for TCP/IP nodes.	

Notes:

1. The default TCP/IP port number for DRDA is 446

2. Unless you know that the host or AS/400 database server supports SECURITY SOCKS, do not specify SECURITY for a TCP/IP node.

Table 10. DCS Directory Parameters

Parameter	Example	Your value
Database name	DB2DB	
Target database name	NEW_YORK3	
Application requester		
Parameter string		

Table 11. System Database Directory Parameters

Parameter	Example	Your value
Database name	DB2DB	
Database alias	NYC3	
Node name	DB2NODE	
Authentication	DCS	

Appendix D. National Language Support Considerations

DB2 Connect has the following national language support considerations:

- DB2 Connect messages are translated into certain languages. For information about accessing translated messages, see the *Quick Beginnings* for your platform.
- DB2 Connect supports a large number of languages and code pages. For a list of these code pages, see the *Administration Guide*.
- When data is transferred between DB2 Connect and a host or AS/400 database server, it is usually converted from a workstation code page to a host CCSID (and vice versa). This is described in the remainder of this appendix.

Further information about using DB2 Connect can be found in the DB2 Connect *Quick Beginnings* manuals, including:

- Date and time formats
- Which languages are supported by DB2 Connect Enterprise Edition and DB2 Connect Personal Edition
- How to customize your DB2 Connect workstation for your particular national language environment
- How to customize your host Coded Character Set Identifier (CCSID) setting.

Conversion of Character Data

When character data is transferred between machines, it must be converted to a form that the receiving machine can use.

For example, when data is transferred between the DB2 Connect workstation and a host or AS/400 database server, it is usually converted from a workstation code page to a host CCSID, and vice versa. If the two machines use different code pages or CCSIDs, code points are mapped from one code page or CCSID to the other. This conversion is always performed at the receiver.

Character data sent **to** a database consists of SQL statements and input data. Character data sent **from** a database consists of output data. Output data that is interpreted as bit data (for example, data from a column declared with the FOR BIT DATA clause) is not converted. Otherwise all input and output character data is converted if the two machines have different code pages or CCSIDs.

139

For example, if DB2 Connect is used to access DB2 Universal Database for OS/390 or DB2/MVS data, the following happens:

1. DB2 Connect sends an SQL statement and input data to OS/390 or MVS.
2. DB2 Universal Database for OS/390 converts the data to an EBCDIC CCSID and processes it.
3. DB2 Universal Database for OS/390 sends the result back to the DB2 Connect workstation.
4. DB2 Connect converts the result to an ASCII or ISO code page and returns it to the user.

The table that follows shows the conversions that are supported between code pages (on the workstation) and CCSIDs (on the host).

For more detailed information about supported code page conversions, refer to the *Administration Guide*.

Table 12. Workstation Code Page to Host CCSID Conversion

Host CCSIDs	Code Page	Countries
037, 273, 277, 278, 280, 284, 285, 297, 500, 871, 1140-1149	437, 819, 850, 858, 860, 863, 1004, 1051, 1252, 1275	Albania, Australia, Austria, Belgium, Brazil, Canada, Denmark, Finland, France, Germany, Iceland, Ireland, Italy, Latin America, Netherlands, New Zealand, Norway, Portugal, South Africa, Spain, Sweden, Switzerland, UK, USA
423, 875	737, 813, 869, 1253, 1280	Greece
870	852, 912, 1250, 1282	Croatia, Czech Republic, Hungary, Poland, Romania, Serbia/Montenegro (Latin), Slovakia, Slovenia
1025	855, 866, 915, 1251, 1283	Bulgaria, FYR Macedonia, Russia, Serbia/Montenegro (Cyrillic)
1026	857, 920, 1254, 1281	Turkey
424	862, 916, 1255	Israel - see note 3 below
420	864, 1046, 1089, 1256	Arabic countries - see note 3 below
838	874	Thailand
930, 939, 5026, 5035	932, 942, 943, 954, 5039	Japan
937	938, 948, 950, 964	Taiwan

Table 12. Workstation Code Page to Host CCSID Conversion (continued)

Host CCSIDs	Code Page	Countries
933, 1364	949, 970, 1363	Korea
935, 1388	1381, 1383, 1386	People's Republic of China
1112, 1122	921, 922	Estonia, Latvia, Lithuania
1025	915, 1131, 1251, 1283	Belarus
1123	1124, 1125, 1251	Ukraine

Notes:

1. Code page 1004 is supported as code page 1252.

2. In general, data can be converted from a code page to a CCSID and back again to the same code page with no change. The following are the only exceptions to that rule:

   - In double-byte character set (DBCS) code pages, some data containing user-defined characters may be lost.

   - For single-byte code pages defined within mixed-byte code pages, and for some newer single-byte code pages, characters that do not exist in both the source and the target may be mapped to substitution characters and then lost when the data is converted back to the original code page.

3. For bidirectional languages, a number of special ″BiDi CCSIDS″ have been defined by IBM and are supported by DB2 Connect Version 6.

   If the bidirectional attributes of the database server are different from those of the client you can use these special CCSIDS to manage the difference.

   Refer to the *Administration Guide* for details of these special CCSIDs. Refer to the Release Notes for DB2 Connect Version 6 for detailed information about how to set them up for DRDA host connections.

Appendix E. Using DCE Directory Services

With DCE Cell Directory Services (CDS), you can store server information in CDS rather than having to store server information on each client. CDS is supported for all DB2 Universal Database clients and DB2 Connect Enterprise Edition on all platforms.

Note: If you want to use DCE Cell Directory Services support in DB2 Connect to connect to DB2 for MVS/ESA over SNA connections, then you must apply the DB2 for MVS/ESA PTF UN73393 that supports the use of DB2DRDA as the name of the remote transaction program name (RTPN).

If you want to use a DCE directory, you must create the following:
- The database object, which contains information about a database
- The database locator object, which contains information about the connection between remote clients and the DB2 Connect workstation
- The routing information object, which matches database objects to database locator objects

For each host or AS/400 database server that you will access, before creating these objects, you should:
- Make sure that the following DCE attributes have been added to the cds attributes file on the workstation from which you created the objects. On an AIX system, the filename is */etc/dce/cds_attributes*. On an OS/2 system, the filename is *x:\opt\dcelocal\etc\cds_attr*, where *x:* is the drive name. On a Windows NT system, the filename is *x:\root\dcelocal\etc\cds_attributes* where *x:* is the drive name and *root* is the directory where you installed dce.

```
1.3.18.0.2.4.30 DB_Comment                 char
1.3.18.0.2.4.31 DB_Communication_Protocol  char
1.3.18.0.2.4.32 DB_Database_Protocol       char
1.3.18.0.2.4.33 DB_Database_Locator_Name   char
1.3.18.0.2.4.34 DB_Native_Database_Name    char
1.3.18.0.2.4.35 DB_Object_Type             char
1.3.18.0.2.4.36 DB_Product_Name            char
1.3.18.0.2.4.37 DB_Product_Release         char
1.3.18.0.2.4.38 DB_Target_Database_Info    char
1.3.18.0.2.4.39 DB_Authentication          char
1.3.18.0.2.4.63 DB_Principal               char
```

- Make sure you have logged into DCE with sufficient authority to create the objects. The following DCE command can be used to login on a UNIX or Windows NT system:

```
dce_login principal-id password
```

143

The following DCE command can be used to login on an OS/2 system:

```
dcelogin principal-id password
```

Note: Before you can connect to databases using these objects, you should also configure communications on the host or AS/400 database server and workstations. This is described in *DB2 Connect Enterprise Edition for OS/2 and Windows NT Quick Beginnings*, and *DB2 Connect Personal Edition Quick Beginnings*.

Creating a Database Object

The database object defines the host or AS/400 database server to DB2 Connect. It must always be defined.

For each host or AS/400 database server that you will access, use the DCE command `cdscp create object` to create a database object. For example:

```
cdscp create object database_global_name
```

Add the following attributes to the object:

DB_Object_Type
> D for database

DB_Product_Name
> The relational database product (for example, DB2_for_MVS, or DB2_for_OS390)

DB_Native_Database_Name
> The database name on the host or AS/400 database server system, as follows:
>
> **MVS or OS/390**
> > The LOCATION value
>
> **VSE or VM**
> > The database name
>
> **OS/400**
> > The relational database name

DB_Database_Protocol
> DRDA

DB_Authentication
> SERVER, CLIENT, or DCE, as described in "Security with DCE Directory Services" on page 150.

DB_Principal

If the Authentication method is DCE, enter the DCE Principal in this attribute.

DB_Communication_Protocol

The following information about the communication protocol between the DB2 Connect server and the host or AS/400 database server:

1. For the communication protocol `APPC`:
   a. The communication protocol (APPC)
   b. The network ID of the host or AS/400 database server
   c. The LU name for the host or AS/400 database server
   d. The transaction program name for connections to the host or AS/400 database server. For DB2 for MVS/ESA, specify DB2DRDA. For any other operating system, specify a valid value that is not in hexadecimal format.
   e. The mode name
   f. The security type, as described in "Security with DCE Directory Services" on page 150. For example:

   ```
   APPC;SPIFNET;NYM2DB2;DB2DRDA;IBMRDB;PROGRAM
   ```

2. For the communication protocol `TCPIP`:
   a. The communication protocol (TCPIP)
   b. The destination TCP/IP hostname (for the host or AS/400 database server).
   c. The TCP/IP port number.
   d. Type of connection (whether using `SOCKS` or `NONE`). This is optional. If not specified, `NONE` is used. For example, the following are the attribute values for the communication protocol TCP/IP:

   ```
   tcpip;jaguar;19713;NONE
   ```

To create a Database object with system security you could put the following instructions into a file:

```
create object /.../cdscell1/subsys/database/DBMVS01
add    object /.../cdscell1/subsys/database/DBMVS01 DB_Object_Type=D
add    object /.../cdscell1/subsys/database/DBMVS01 DB_Product_Name=DB2_for_MVS
add    object /.../cdscell1/subsys/database/DBMVS01 DB_Database_Protocol=DRDA
add    object /.../cdscell1/subsys/database/DBMVS01 DB_Native_Database_Name=\
               NEW_YORK
add    object /.../cdscell1/subsys/database/DBMVS01 DB_Authentication=SERVER
add    object /.../cdscell1/subsys/database/DBMVS01 DB_Communication_Protocol=\
               APPC;SPIFNET;NYM2DB2;DB2DRDA;IBMRDB;PROGRAM
```

Then, enter the command:

```
cdscp < filename
```

Note: In the file, specify a backslash (\) whenever you want a statement to continue to the next line.

To create a Database object with DCE security you could put the following instructions into a file:

```
create object /.../cdscell1/subsys/database/DBMVS02
add    object /.../cdscell1/subsys/database/DBMVS02 DB_Object_Type=D
add    object /.../cdscell1/subsys/database/DBMVS02 DB_Product_Name=DB2_for_MVS
add    object /.../cdscell1/subsys/database/DBMVS02 DB_Database_Protocol=DRDA
add    object /.../cdscell1/subsys/database/DBMVS02 DB_Native_Database_Name=\
                 NEW_YORK
add    object /.../cdscell1/subsys/database/DBMVS02 DB_Authentication=DCE
add    object /.../cdscell1/subsys/database/DBMVS02 DB_Principal=\
            /.../cdscell1/principal_name
add    object /.../cdscell1/subsys/database/DBMVS02 DB_Communication_Protocol=\
                 APPC;SPIFNET;NYM2DB2;DB2DRDA;IBMRDB;NONE
```

Then, enter the command:

```
cdscp < filename
```

Creating a Database Locator Object

The database locator object is used to define a DB2 Connect Enterprise Edition gateway to its clients.

For your DB2 Connect workstation, use the DCE command `cdscp create object` to create a database locator object. For example:

```
cdscp create object object_global_name
```

Add the following attributes to the object:

DB_Object_Type
L for locator object

DB_Communication_Protocol
Following is the information that you need to set up each communication protocol between the DB2 Connect workstation and remote clients.

The following summarizes protocol support by platform:
- On OS/2: APPC, IPX, NETBIOS, and TCP/IP
- On Windows 32-bit operating systems: APPC, IPX, NETBIOS, NPIPE, and TCP/IP
- On UNIX: APPC and TCP/IP

For APPC:
1. The communication protocol (APPC)

2. The network ID of the DB2 Connect workstation
3. The LU name for the DB2 Connect workstation
4. The transaction program name for connections from remote clients
5. The mode name
6. The security type, as described in "Security with DCE Directory Services" on page 150.

For TCP/IP:

1. The communication protocol (`TCPIP`)
2. The host name of the DB2 Connect workstation
3. The connection port used by the DB2 Connect workstation to accept connections from remote clients
4. Type of connection (whether using SOCKS or NONE). This is optional. If not specified, NONE is used.

For IPX/SPX:

1. The communication protocol (`IPXSPX`)
2. The file server name. Use `*` for direct addressing
3. The object name. Use the internetwork address for direct addressing.

For Named Pipes:

1. The communication protocol (`NPIPE`)
2. The `computer name` of the DB2 Connect workstation.
3. The instance name.

For NETBIOS:

1. The communication protocol (`NETBIOS`)
2. The `NNAME` for the server or DB2 Connect Enterprise Edition gateway.

For example, you could put the following lines into a file:

```
create object /.../cdscell1/subsys/database/DBAIX01
add    object /.../cdscell1/subsys/database/DBAIX01 DB_Object_Type= L
add    object /.../cdscell1/subsys/database/DBAIX01 DB_Communication_Protocol=\
         TCPIP;AIX001;3700
add    object /.../cdscell1/subsys/database/DBAIX01 DB_Communication_Protocol=\
         APPC;SPIFNET;NYX1GW01;NYSERVER;IBMRDB;NONE
```

Then, enter the command:

```
cdscp < filename
```

On Windows 32-bit operating systems you can specify a named pipe in a similar manner. For example:

```
add     object /.../cdscell1/subsys/database/DBAIX01 DB_Communication_Protocol=\
        NPIPE;computer_name;instance_name
```

On OS/2 you can specify the protocol in the DB_Communication_Protocol attribute. For example:

- IPXSPX;fileserver;objectname
- NETBIOS;nname

Creating a Routing Information Object

This object is required to be defined in DCE and it is retrieved by the DB2 client.

Use the DCE command `cdscp create object` to create a routing information object. For example:

```
cdscp create object object_global_name
```

Add a **DB_Object_Type** attribute of R.

For each database object, add one **DB_Target_Database_Info** attribute. Each **DB_Target_Database_Info** attribute consists of the following parameters:

Database
> The name of a database object, including the full path. Specify `*OTHERDBS` to indicate all other databases that are not specified explicitly.

Outbound protocol
> The database protocol for host or AS/400 database server connections (DRDA)

Inbound protocol
> The database protocol for remote client connections (DB2RA),

Authenticate at Gateway
> 0 (for No) or 1 (for Yes), as described in "Security with DCE Directory Services" on page 150.

Parameter String for Gateway
> The string containing the parameters to be used in the gateway. Its content is gateway specific. For DB2 Connect gateway specific strings refer to "DCS Directory" on page 21.

Database Locator
> The name of a database locator object representing the DB2 Connect workstation

For example, you could put the following lines into a file:

```
create object /.../cdscell1/subsys/database/ROUTE1
add    object /.../cdscell1/subsys/database/ROUTE1 DB_Object_Type=R
add    object /.../cdscell1/subsys/database/ROUTE1 DB_Target_Database_Info=\
/.../cdscell1/subsys/database/DBMVS01;DRDA;DB2RA;0;;\
/.../cdscell1/subsys/database/DBAIX01
add    object /.../cdscell1/subsys/database/ROUTE1 DB_Target_Database_Info=\
*OTHERDBS;DRDA;DB2RA;0;;\
/.../cdscell1/subsys/database/DBAIX02
```

Then, enter the command:

```
cdcsp < filename
```

Setting Configuration Parameters

Update the Database Manager configuration of the client as follows:

```
DB2 UPDATE DATABASE MANAGER CONFIGURATION USING
[DIR_PATH_NAME path]
DIR_OBJ_NAME loc_obj
DIR_TYPE DCE
[ROUTE_OBJ_NAME route_obj]
[DFT_CLIENT_COMM protocol]
[DFT_CLIENT_ADPT 0-15]
```

where:
- *path* is the default path used to form the complete name of target databases (default /.:/subsys/database/)
- *loc_obj* is used to identify the client in the DCE name space
- DIR_TYPE DCE specifies that DCE directories are used by the client application
- *route_obj* is the name of the routing information object (e.g. ROUTE1).
- *protocol* is the communication protocol between the client and the DB2 Connect workstation (APPC or TCPIP for UNIX; APPC, IPXSPX, NETBIOS, or TCP/IP for OS/2; for Windows 32-bit operating systems, APPC, TCPIP, IPXSPX, NETBIOS, and NPIPE).
- Default Client Adapter 0 through 15 for NETBIOS. If the protocol is NETBIOS and the client adapter number is not the default value 0, specify the client adapter number.

Note: The following environment variables can respectively overwrite those listed above.
  - DB2DIRPATHNAME can overwrite DIR_PATH_NAME
  - DB2ROUTE can overwrite ROUTE_OBJ_NAME
  - DB2CLIENTCOMM can overwrite DFT_CLIENT_COMM

- DB2CLIENTADPT can overwrite DFT_CLIENT_ADPT

Cataloging the Database

If a database is in a different path than the default, or if you want to use an alias that is different than the database name, you can catalog the global database. You can use the command line processor CATALOG GLOBAL DATABASE command as follows:

```
db2 CATALOG GLOBAL DATABASE database_global_name
    AS alias
    USING DIRECTORY DCE
```

The alias will be used by any application program that accesses the database.

For example:

```
db2 CATALOG GLOBAL DATABASE /.../cdscell2/subsys/database/dbmvs12 AS NYC3
    USING DIRECTORY DCE
```

Security with DCE Directory Services

As DB2 Connect administrator, you can determine where user names and passwords are validated. With DCE directories, you do this by setting the following:
- The security type of the communication protocol in the database locator object representing the DB2 Connect workstation. Use security type NONE.
- The authentication type of the database object.
- The security type of the communication protocol in the database object
- The `authenticate at gateway` parameter in the routing information object

Table 13 on page 151 and Table 14 on page 151 show the possible combinations of these values and where validation is performed for each combination. Only the combinations shown in these tables are supported by DB2 Connect with DCE Directory Services.

Table 13. Valid Security Scenarios with DCE using APPC connections

	Database object of the Server		Routing object	Validation
Case	Authentication	Security	Authentication at DB2 Connect Gateway (1=true, 0=false)	
1	CLIENT	SAME	0	Remote client (or DB2 Connect workstation)
2	CLIENT	SAME	1	DB2 Connect workstation
3	SERVER	PROGRAM	0	host or AS/400 database server
4	SERVER	PROGRAM	1	DB2 Connect workstation and host or AS/400 database server
5	DCE	NONE	NOT APPLICABLE	At the DCE security server

Note: If a remote client is connected to the DB2 Connect Enterprise Edition gateway workstation via an APPC connection, specify a security type of NONE in the DCE locator object of the gateway.

Table 14. Valid Security Scenarios with DCE using TCP/IP connections

Case	Database object of the Server	Routing object	Validation
	Authentication	Authentication at DB2 Connect Enterprise Edition Gateway (1=true, 0=false)	
1	CLIENT	0	Remote client (or DB2 Connect workstation)
2	CLIENT	1	DB2 Connect workstation
3	SERVER	0	host or AS/400 database server
4	NOT APPLICABLE	NOT APPLICABLE	None

Table 14. Valid Security Scenarios with DCE using TCP/IP connections (continued)

Case	Database object of the Server	Routing object	Validation
	Authentication	Authentication at DB2 Connect Enterprise Edition Gateway (1=true, 0=false)	
5	DCE	NOT APPLICABLE	At the DCE security server

Each combination is described in more detail below:

- In the first case, the user name and password are validated only at the remote client. (For a local client, the user name and password are validated only at the DB2 Connect workstation.)

  The users are expected to be authenticated at the location they first sign on to. The user ID is sent across the network, but not the password. Use this type of security only if all client workstations have adequate security facilities.

- In the second case, the user name and password are validated at the DB2 Connect workstation only. The password is sent across the network from the remote client to the DB2 Connect server but not to the host or AS/400 database server.

- In the third case, the user name and password are validated at the host or AS/400 database server only. The password is sent across the network from the remote client to the DB2 Connect server and from the DB2 Connect workstation to the host or AS/400 database server.

- In the fourth case, the user name and password are validated at both the DB2 Connect workstation and the host or AS/400 database server. The password is sent across the network from the remote client to the DB2 Connect server and from the DB2 Connect server to the host or AS/400 database server.

  Because validation is performed in two places, the same set of user names and passwords must be maintained at both the DB2 Connect server and the host or AS/400 database server.

- In the fifth case, a DCE ticket is obtained from the DCE security server.

Notes:

1. For AIX systems, all users using security type SAME must belong to the AIX **system** group.
2. For AIX systems with remote clients, the instance of the DB2 Connect product running on the DB2 Connect server must belong to the AIX **system** group.

3. Access to a host or AS/400 database server is controlled by its own
   security mechanisms or subsystems; for example, the Virtual
   Telecommunications Access Method (VTAM) and Resource Access Control
   Facility (RACF). Access to protected database objects is controlled by the
   SQL **GRANT** and **REVOKE** statements.

Appendix F. Binding Utilities for Back-Level Clients

If you have remote clients from a previous release, you may need to bind utilities on these clients to the host or AS/400 database server:

- If the old client was used with a previous release of DB2 Connect against the same host or AS/400 database server, you do not need to do any additional steps.
- If the old client was not used with DB2 Connect (for example, if several OS/2 machines were connected with no connection to a host or AS/400 database server), do the following:

  1. If you have any DB2 for OS/2 Version 1.0 or 1.2 clients, create a bind list file with the following lines:

     sqlabind.bnd+
     sqlueiwi.bnd+
     sqluigsi.bnd+
     sqluiici.bnd+
     sqluiict.bnd+
     sqluexpm.bnd+
     sqluimpm.bnd+
     sqlurexp.bnd+
     sqlarxcs.bnd+
     sqlarxrr.bnd+
     sqlarxur.bnd

     and copy each of these bind files from one of your clients to the DB2 Connect workstation.

  2. If you have Client Application Enabler Version 1.0 or 1.2, create a bind list file with the following lines:

     db2ajgrt.bnd+
     db2clics.bnd+
     db2clpcs.bnd+
     db2clprr.bnd+
     db2clpur.bnd+
     db2ueiwi.bnd+
     db2uigsi.bnd+
     db2uiici.bnd+
     db2uiict.bnd+
     db2uexpm.bnd+
     db2uimpm.bnd+
     db2urexp.bnd

and copy each of these bind files from one of your clients to the DB2 Connect workstation.

3. At the DB2 Connect server, bind each bind list file to each host or AS/400 database server database. Issue commands similar to the following:

```
db2 connect to DBALIAS user USERID using PASSWORD
db2 bind path@bindfile.lst blocking all
      sqlerror continue messages bindfile.msg grant public
db2 connect reset
```

where *DBALIAS*, *USERID*, and *PASSWORD* apply to the host or AS/400 database server database, *bindfile* is the name of the bind list file, and *path* is the location of the bind list file.

You can use the `grant` option of the `bind` command to grant EXECUTE privilege to PUBLIC or to a specified user name or group ID. If you do not use the `grant` option of the `bind` command, you must GRANT EXECUTE (RUN) individually for each package.

To find out the package names for the bind files, enter the following command:

```
ddcspkgn @bindfile.lst
```

Appendix G. Tuning CLI/ODBC Application Performance with the CLISCHEMA Keyword

This section contains new information to help you tune the performance of your ODBC∕CLI applications using the CLISCHEMA initialization keyword. It does not contain general information about tuning network or database performance (see "Chapter 11. DB2 Connect Performance" on page 111). Here is a summary of the information that follows:

- "Target Environment"
- "What CLI∕ODBC Is"
- "How CLI∕ODBC Works" on page 158
- "The DB2 CLISCHEMA Initalization Keyword" on page 158
- "Suggested Approach" on page 161
- "Additional Hints and Tips" on page 161
- "db2ocat Catalog Optimizer Tool" on page 162
- "Additional Information Sources" on page 162

Target Environment

The information presented here is intended primarily for users of DB2 Universal Database for OS∕390, and the target environment comprises:
- A CLI∕ODBC application running with a DB2 Universal Database client
- DB2 Connect Version 5 or higher (Personal Edition or Enterprise Edition)
- DB2 Universal Database for OS∕390 Version 5.1 or higher (or DB2 for MVS∕ESA except where indicated otherwise).

What CLI/ODBC Is

CLI∕ODBC is an SQL application programming interface that is callable from your database applications. It passes dynamic SQL statements as database function calls. Unlike embedded SQL it does not require host variables or a precompiler.

157

How CLI/ODBC Works

When an application program calls CLI/ODBC, the first thing that it must do is to make SQL calls to some of the system catalog tables on the target database in order to obtain information about other database contents. CLI/ODBC applications always access the system catalog tables in this way. There are ten API calls that may be made in order to gather information about the database that is being connected to. These API calls include:

- SQLTables
- SQLColumns
- SQLSpecialcolumns
- SQLStatistics
- SQLPrimarykeys
- SQLForeignkeys
- SQLTablePrivileges
- SQLColumnPrivileges
- SQLProcedures
- SQLProcedureColumns.

Further information about these API calls and the tables they reference is published in *CLI Guide and Reference*.

By default, when you connect to a database, your CLI/ODBC application will query the system catalog tables for information about ALL the database tables in that database. Especially on a large system this can result in a lot of network traffic and considerable delays when starting an application.

The DB2 CLISCHEMA Initalization Keyword

DB2 Universal Database provides several CLI/ODBC initialization keywords that can be used to limit the amount of data that will be returned by the initial API calls during the "information gathering" stage after the database is first connected to. These keywords can be set by:

1. Manually editing the db2cli.ini file

2. Changing ODBC/CLI settings for the database using the Client Configuration Assistant (on those platforms which support it)

3. Updating the database CLI configuration using the DBA Command Line Interface.

The keywords are:

- DBNAME
- TABLETYPE
- SCHEMALIST
- SYSSCHEMA
- CLISCHEMA

With the exception of the information for CLISCHEMA, these keywords are documented in the CLI/ODBC help and *CLI Guide and Reference.* The rest of this discussion relates only to the use of CLISCHEMA.

Here is the documentation for CLISCHEMA that will be added to *CLI Guide and Reference* in due course:

```
db2cli.ini Keyword Syntax:  CLISCHEMA = clischema

Default Setting:            No alternatives specified.

DB2 CLI/ODBC Settings Tab:  Not present.
```

Usage Notes

The CLISCHEMA option indicates an alternative schema, tables, and index set to be searched instead of the SYSIBM (or SYSTEM, QSYS2) schemas when the DB2 CLI and ODBC Catalog Function calls are issued to obtain catalog information.

For example, if you specify CLISCHEMA='SERGE', the internal CLI/ODBC API calls that normally reference the system tables will reference the following user tables instead:

- SERGE.TABLES
- SERGE.COLUMNS
- SERGE.SPECIALCOLUMNS
- SERGE.TSTATISTICS
- SERGE.PRIMARYKEYS
- SERGE.FOREIGNKEYS
- SERGE.TABLEPRIVILEGES
- SERGE.COLUMNTABLES
- SERGE.PROCEDURES
- SERGE.PROCEDURESCOLUMNS.

These user tables must be built by the database administrator before CLISCHEMA can be used.

Note: DataPropagator provides support for CLISCHEMA, so that the Database Administrator can perform this task in three possible ways:
1. Using db2cli.exe on the client
2. Automatically on the server using DataPropagator
3. Manually on the server.

The information which follows explains how this taks can be performed on the client.

db2cli and bldschem Utilities

A utility to set up the user tables required by CLISCHEMA is provided in the form of the previously undocumented bldschem support command of the CLI Command Line Interface, which can be found as: /samples/cli/db2cli.exe. Documentation for db2cli.exe, not including the bldschem support command, can found in /samples/cli/INTCLI.DOC.

For example, to build the set of user tables that is required to work with CLISCHEMA='SERGE' for the table name STAFF owned by the schema owner (creator) USERID, in the database SAMPLE, you would run the following command after issuing db2start and after registering the database to ODBC/CLI:

```
db2cli < addstaff.txt
```

Where "addstaff.txt" contains the following script:

```
opt callerror on
opt echo on
quickc 1 1 sample userid password
#
# Repeat next line for each table to add.
#
bldschem 1 SERGE USERID STAFF
#
# Exit
#
killenv 1
```

This will result in the creation of the set of tables SERGE.* as listed above, with indexes, populated using the system catalog table data for the table USERID.STAFF. For example, SERGE.TABLES will be populated with a new row for each entry that is matched. Additional bldschem calls result in appends to the existing SERGE.* tables, with replacement of existing rows.

In summary, the syntax of the bldschem support command is:

```
bldschem <handle_number> <value_of_CLISCHEMA> <schema_owner> <table_name>
```

Where:

```
- <handle_number> should be 1

- <value_of_CLISCHEMA> should be the same as the schema name specified with
    the CLISCHEMA keyword

- <schema_owner> is the creator of the table

- <table_name> can be the name of a user table, a view, an alias, a synonym,
    or a system table name.
  (Wildcard characters are allowed).
```

A second example is as follows. If you subsequently run the following through db2cli.exe, then you will append to the user tables SERGE.* created in the previous example, adding rows which reflect the data in the system catalog tables for each table for which FRED and BERT are the schema owners.

```
bldschem 1 SERGE FRED %
bldschem 1 SERGE BERT %
```

When the CLISCHEMA CLI/ODBC keyword is subsequently set to SERGE, processing by ODBC/CLI applications against the SAMPLE database will reference the SERGE.* set of tables instead of the system catalog tables.

Suggested Approach

In most production environments the default search of the system catalog tables can return a very large amount of data, so that each time a CLI/ODBC application opens a database there can be a considerable delay. Even on a typical test database, the delay can easily be of the order of 25 seconds or so.

Measure this delay initially without having any of the above CLI keywords set, while remembering to discount connect time and especially the long delays that can occur when a DB2 client issues its first ever connection to a new database - autobinding can often take several minutes.

How you proceed next depends on the structure of your data and your organization. In some cases you can use DBNAME, SCHEMALIST, and TABLETYPE in combination to limit the search for use by a particular application or group of applications. For example, if production DBA clients usually access tables under a given DBNAME and schema then this is easy to specify.

CLISCHEMA provides the best performance advantages for most users. For this reason we would generally recommend using CLISCHEMA in a production environment, since it is much easier to set up and modify CLISCHEMA's user tables through the CLI command Line Interface (db2cli.exe) and its bldschem support command. See also "db2ocat Catalog Optimizer Tool" on page 162.

Additional Hints and Tips

The CLISCHEMA keyword has to be added to the db2cli.ini file within a section for the DSN name or the common section. A section is text in square brackets. The COMMON section is indicated by the text ″COMMON″ in square brackets. Note that keywords and section names are not case-sensitive.

At connect, each possible keyword is checked for first under the DSN name, then if not found, under the COMMON section. This allows for both DSN specific keywords and global (client) keywords.

Also, the DBALIAS keyword can be used to create different DSN (ODBC Data Sources) that map to the same database. (A DSN name can be up to 255 characters in length, and it is mapped to the 8 char dbname).

In the example below, any time a user connects to TESTDB or any DSN that is not listed in the file, they will use clischema=ODBCCAT. If they connect to TestDBcar2, they will use clischema=odbccat2, but still connect to the testdb database.

Example db2cli.ini file:

```
[TESTDB]

[COMMON]
clischema=odbccat

[TestDBcat1]
DBALIAS=testdb
clischema=odbccat1

[TestDBcat2]
DBALIAS=testdb
clischema=odbccat2
```

db2ocat Catalog Optimizer Tool

A new tool db2ocat is provided on Windows 32-bit operating systems in order to help you optimize system catalog searches for ODBC and JDBC applications.

You can obtain the db2ocat point-and-click catalog optimizer utility by downloading db2ocat.zip from: ftp://ftp.software.ibm.com/ps/products/db2/tools.

Additional Information Sources

The following additional information sources may be of interest:
• Paper on automatic catalog propagation using DataPropagator:
  http://www.software.ibm.com/data/db2/os390/odbcatlg.html
• Paper on manual approach not using DataPropagator:
  http://www.software.ibm.com/data/db2/os390/odbcmanu.html

Appendix H. Additional SNA Performance Tuning Hints and Tips

This appendix contains additional SNA performance tuning hints and tips for use with DB2 Connect.

General Performance Information for DB2 Connect

The performance characteristics of DB2 Connect are that it predominantly uses the processor and performs very little I/O. In general, the faster the processor speed, the faster DB2 Connect will run. DB2 Connect fully exploits SMP processor configurations.

A fast DB2 Connect Enterprise Edition gateway can handle an SQL request/reply pair in less than five milliseconds, not counting client time, network time, and processing time at the host or AS/400 server. A simple SQL statement or query with a few rows of data could be completed end-to-end in less than 0.1 seconds (from client to the host or AS/400 server and back).

When there are more than four or five SQL statements in a query, then the use of stored procedures will help to ensure high OLTP performance and to avoid increases in lock contention due to network delays between SQL statements.

Performance problems are usually caused by to the type of host attachment in use, network routing and tuning chracteristics, and application design. Some general DB2 Connect performance information can be found in "Other DB2 Connect Performance Information Sources" on page 164.

Selection and Tuning of the Network Attachment

In order of likely best performance when using DB2 Connect, various types of network attachment include:

1. Channel attachment card
2. IBM 3172 Model 3, or newer models, or equivalent
3. IBM 2216
4. Open System Adaptor Card (OSA-2, not OSA-1)
5. IBM 3745 with Network Control Program (NCP)
6. IBM 3174 Terminal Controllers, or equivalent

The last option is not recommended - see below.

The recommended best way to connect to the host is to use ESCON channel attachment cards for AIX or Windows NT. The IBM 3172 Model 3 and 2216 also perform well, but they tend to deliver throughput inferior to ESCON.

When using AIX with ESCON cards, please apply the PTFs related to MPC (Multi Path Channel). Without these PTFs the AIX SNA ESCON driver may deliver worse performance. See "Multi Path Channel Support for SNA over ESCON" on page 165 for more details. Further information can also be found at:

```
http://www.networking.ibm.com.cms/cmsnew01.html
```

See "How to Tune DB2 Connect Connections via NCP" on page 165 for a checklist of which Communications Server, NCP, and VTAM parameters to tune to optimize DB2 Connect performance. All the non-NCP specific recommendations are applicable to all types of DB2 Connect and client/server attachments.

The OSA-2 card on System 390 might not deliver throughput as high as a 3272 Model 3 when there is a high volume of small transactions, owing to its lower frames-per-second capability. See "Information about OSA-2 Enhancements" on page 169 for details of some recent enhancements.

3145 with NCP is usually tuned specifically for existing network traffic. Consequently it might not perform as well for database client/server applications. Most DB2 Connect performance problems are due to the time delay between the NCP and VTAM and/or between NCPs. See "How to Tune DB2 Connect Connections via NCP" on page 165, which provides a tuning checklist.

In general, we recommend avoiding the use of 3174 Terminal Controllers because their packet size (RU size) of 256 bytes is too small. 3174 microcode level C is required in order to provide Independent LU support for APPC database connections. Some OEM 3174 equivalents may have similar dependencies.

Other DB2 Connect Performance Information Sources

- Search the DB2 Technical Library web site at http://www.software.ibm.com/data/db2/library. Search the DB2 Universal Database library for ″Technotes″ with the keywords ″DB2CONNECT″ and ″Performance″ to find the latest information with a section on DB2 Connect considerations on the World Wide Web.
- There are also a number of DB2 client/server performance reports involving DB2 Connect in http://www.software.ibm.com/data/performance.

Multi Path Channel Support for SNA over ESCON

Multi Path Channel (MPC) support for SNA over ESCON allows a system running IBM eNetwork Communications Server to use an ESCON adapter to create an MPC linkstation to the host. MPC is typically faster than CDLC because:

1. MPC uses separate subchannels for read and write
2. MPC is not limited by IOBUF size. Frames are 4K and may be blocked together.

Tests have shown as much as a threefold improvement on an MPC link compared to an ESCON Channel Data Link Control (CDLC) link with an IOBUF size less than 1K. AIX SNA MPC requires ESCON and MVS VTAM V4R4 or later and feature code 4024 of Communications Server for AIX (5765-652). Windows NT systems must use IBM eNetwork Communications Server for Windows NT Version 6.

The following are the Communications Server for AIX PTFs required for MPC:

```
APAR #     PTF #     LPP name

IX67032    U449693   sna.books.chdoc
IX67032    U449693   sna.books.escdoc
IX67032    U449300   sna.rte
IX67032    U450027   sna.msg.en_US.rte
IX65820    U447759   sna.dlcchannel
IX67618    U449691   mpc.rte
IX65813    U447758   devices.mca.8fc3.rte
```

How to Tune DB2 Connect Connections via NCP

A typical network configuration might be:

SDLC lines

3270 traffic etc.

VTAM

NCP

LAN

DB2 Connect
Enterprise Edition

DB2 CAE
for DOS

DB2 CAE
Windows*

DB2 CAE
for OS/2

DB2 CAE
for UNIX

DB2 CAE
for Mac
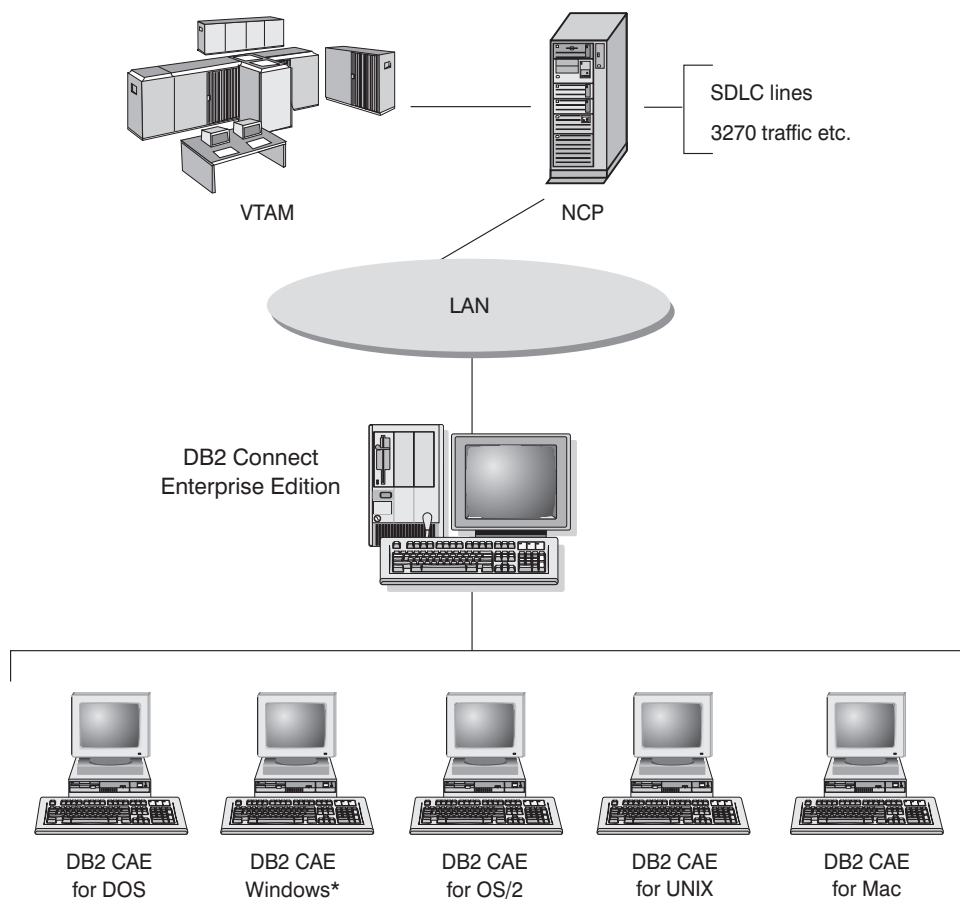
*16 and 32-bit Windows operating systems.

Figure 9. DB2 Connect Enterprise Edition gateway SNA network scenario

This scenario focuses on the throughput and response time between the host or AS/400 database server to the DB2 Connect Enterprise Edition gateway and various parameters that could affect this.

Tuning Criteria

The suggested order in which to make these changes is:

```
1 - DELAY on PCCU macro*
2 - DLC/LLC Tuning*
3 - PIU size*
4 - Pacing window changes*
5 - DELAY on LINE macro*
6 - MAXBFRU changes
```

```
  7 - LAN Frame sizes
```

```
  * Major improvement in throughput is possible
```

PIU size (RU + 29 bytes)

The RU size at the host and the DB2 Connect server should be maximized. This implies that the RU size should be large enough to contain the API crossing (both SEND and RECEIVE data for the transaction where possible) in order to minimize the number of times VTAM program stack must be traversed. Also, the network frame size may limit the maximum RU size if RU segmentation is not desired.

It is a good idea to set the DB2 Connect block size (RQRIOBLK), RU and pacing values such that RU * pacing >= RQRIOBLK. For example, the default RQRIOBLK size of 32K is a good value for most situations, and to exploit this you would set RU = 4K and receive window pacing to 8.

- RU size and pacing are set by the mode table which is defined on both DB2 Connect workstation and in VTAM. The mode table definitions should be the same in both places.
- RQRIOBLK is set using the DB2 UPDATE DBM CFG command.
- Network frame size I-frame is set in DLC configuration on the DB2 Connect workstation and in NCP.

Pacing window sizes

The session and VR pacing windows should be maximized: the largest value that does not cause network congestion or VR-held conditions, and so on, should be used. For a test environment set pacing to 0 (no pacing) or set it to the maximum value X'3F'.

Coat-tailing values (DELAY)

Coat-tailing is controlled by the DELAY parameter. The DELAY Parameter in the PCCU macro controls outbound coat-tailing (outbound with reference to the host). The DELAY value in the LINE definition statement for the NCP controls inbound coat-tailing (inbound with reference to the host). The DELAY value determines how long a PIU is held in the queue (NCP or VTAM) before it is transmitted. The purpose of this wait is to increase the possibility that other PIUs will arrive in the interim and all of these can be transmitted on a single channel program. For the lowest latency, the DELAY value should be set to 0. Changing the value of the outbound coat-tailing delay value to 0 should have no noticeable effect on the host except for improved performance for outbound traffic. Some improvement in inbound traffic performance will also be realized. Changing the DELAY on the NCP to 0 should be done with a little more care. The value can be set to 0 if the NCP is not overloaded and

the inbound traffic does not consist of a significant percentage of small frames. Setting the values of DELAY to 0 may improve response time significantly, especialy under light loads or test/benchmark environments.

```
VTAMB7   PCCU  CUADDR=CAF,
                              AUTODMP=NO,
                              AUTOIPL=NO,
                              AUTOSYN=YES,
                              BACKUP=YES,
                              DELAY=0,
                              VFYLM=YES,
                              CHANCON=UNCOND,
                              MAXDATA=32768,
                              DUMPDS=NCPDUMP,
                              OWNER=HOSTB7,
                              SUBAREA=17


LNCTLS        GROUP   LNCTL=CA,CA=TYPE6,DELAY=0.0,TIMEOUT=500.0
CA0                   LINE        ADDRESS=00
PUCHAN0   PU          PUTYPE=5,TGN=1
CA1                   LINE        ADDRESS=01
PUCHAN1   PU          PUTYPE=5,TGN=1
```

DELAY considerations are documented in the ″VTAM Network Implementation Guide″.

MAXBFRU

The MAXBFRU value should be set to a value two ro three times larger than the largest PIU size.

DLC/LLC layer Tuning

Ensure that the LLC2 window sizes (DLC send and receive window counts) between the NCP and the DB2 Connect Enterprise Edition gateway are the same. This has a significant effect specially when the server is DB2 Connect for AIX. It is recommended that the send window count be set higher than the receive window count.

In general, for any SNA connection across a Token-ring the LLC2 timers/windows should be optimized. In some cases, this change led to a six-fold improvement in throughput and response time.

LAN Frame sizes

The token ring maximum frame size should be as large as possible.

Information about OSA-2 Enhancements

The following information is reproduced from the IBM WSC Flash document number 9718.

```
TITLE: WSC  FLASH 9718:  OSA-2 ENHANCEMENTS AVAILABLE
DOCUMENT ID G023691 UNCLASSIFIED

Open Systems Adapter 2 (OSA-2) Systems Network Architecture (SNA)
enhancements are being made available earlier than previously announced.
The enhancements are:

o    SNA/APPN enhancements for OS/390, MVS/ESA, VM/ESA, and VSE/ESA

     -    Enhanced availability: load balancing, redundancy, and overflow
     -    Enhanced connectivity: increased Physical Unit (PU) support
          (from 255 PUs per port to 2047 PUs per port).

o    Support for ACF/VTAM for VSE/ESA networks

NOTE: These enhancements do not pertain to OSA-1.


LOAD BALANCING, REDUNDANCY, AND OVERFLOW
_____

LOAD BALANCING: A single Medium Access Control (MAC) address can now be
defined for attached OSA-2 SNA/APPN Physical Units (PUs), even though
connections may be via multiple physical ports.  This support is offered for
source-route bridged environments only (Token-Ring and FDDI). The number of
sessions established through a port is monitored, and user session loads are
evenly distributed across the equally configured ports.

REDUNDANCY:  A secondary path between the LAN workstation and the host system
can now be configured. If the primary path becomes unavailable, the secondary
path will receive the LAN traffic. This increases system availability and
simplifies network management.

OVERFLOW:  User sessions flow through the primary OSA-2 port until the
session capacity has been reached. Additional user sessions will
automatically flow to the next OSA-2 port. Since all user workstations are
identically configured, network administration is simplified and the network
becomes more scalable. New users can be added non-disruptively.

Load balancing, redundancy, and overflow support is provided by PTFs for
OSA/SF as follows:

o    OS/390 and MVS  - OW20205/UW34618          03/31/97
o    VM/ESA          - OW23952/UW37028          03/31/97
o    VSE/ESA         - Provided with VSE/ESA V2.2.1      04/29/97


INCREASED PHYSICAL UNIT (PU) SUPPORT (VIA OSA/SF):
_____
```

The architecture has been changed to allow up to a maximum of 2047 PUs per physical port to be defined for OSA-2 Ethernet, Token-Ring and FDDI features instead of the current 255 PUs per port. This enhancement is available for currently installed features, as well as new installations. Actual connectivity may vary based upon user workloads.

Increased Physical Unit (PU) Support is provided by PTFs for OSA/SF as follows:

```
o   OS/390 and MVS  - OW23429/UW37210        03/31/97
o   VM/ESA          - OW24952/UW37028        03/31/97
o   VSE/ESA         - PQ03091/UQ04224        04/29/97
```

Increased Physical Unit (PU) Support is provided by PTFs for ACT/VTAM as follows:

o   ACF/VTAM for OS/390 and MVS

```
    -   VTAM 4.1 OW14043/UW24904
    -   VTAM 4.2 OW14043/UW24905
    -   VTAM 4.3 OW14043/UW24906
```

o   ACF/VTAM VM/ESA

```
    -   VM60877/UV59834
```

o   ACF/VTAM VSE/ESA

```
    -   DY44347/UD50254
```

VSE/ESA - SNA SUPPORT

_____

OSA-2 and OSA/SF support is delivered via VSE/ESA Version 2 Release 2.1. This announcement of VSE/ESA support satisfies the Statement of General Direction contained in Hardware Announcement 196-194, and Hardware Announcement 196-193, dated September 10, 1996.

The OSA-2 feature provides ACF/VTAM for VSE/ESA host applications with direct access to Ethernet, Token-Ring, and FDDI LANs and Asynchronous Transfer Mode (ATM) Forum-compliant LAN emulation networks.

OSA/SF is available:

o   As a non-exclusive element of OS/390 Release 1 or above (5645-001)
o   As a separate program product, S/390 Open Systems Adapter Support Facility Version 1 Release 2 for MVS/ESA 4.3 or above (5655-104)
o   As a facility of VM/ESA Version 2 Release 2.0 (5654-030)
o   As a component of VSE Central Functions 6.1.1 in VSE/ESA Version 2 Release 2.1 (5690-VSE).

MORE INFORMATION

_____

Announcements 297-043, 297-040

Appendix I. Notices

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

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

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

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

> IBM Canada Limited
> Office of the Lab Director
> 1150 Eglinton Ave. East
> North York, Ontario
> M3C 1H7
> CANADA

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

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

Trademarks

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

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

Trademarks of Other Companies

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

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

HP-UX is a trademark of Hewlett-Packard.

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

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

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

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

UNIX is a registered trademark in the United States, other countries or both and is licensed exclusively through X/Open Company Limited.

Other company, product, or service names, which may be denoted by a double asterisk (**), may be trademarks or service marks of others.

Index

Special Characters

Contacting IBM

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

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

For information or to order any of the DB2 Universal Database products contact an IBM representative at a local branch office or contact any authorized IBM software remarketer.

Telephone

If you live in the U.S.A., call one of the following numbers:
- 1-800-237-5511 to learn about available service options.
- 1-800-IBM-CALL (1-800-426-2255) or 1-800-3IBM-OS2 (1-800-342-6672) to order products or get general information.
- 1-800-879-2755 to order publications.

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

then performing a search using the keyword "handbook".

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

World Wide Web
> http://www.software.ibm.com/data/
> http://www.software.ibm.com/data/db2/library/

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

Anonymous FTP Sites
> ftp.software.ibm.com

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

Internet Newsgroups
　　　　comp.databases.ibm-db2, bit.listserv.db2-l
These newsgroups are available for users to discuss their experiences with
DB2 products.

CompuServe
　　　　GO IBMDB2 to access the IBM DB2 Family forums
All DB2 products are supported through these forums.

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

IBM ®