

IBM[®] DB2[®] Universal Database



What's New

Version 7.2

IBM[®] DB2[®] Universal Database



What's New

Version 7.2

Before using this information and the product it supports, be sure to read the general information under "Appendix B. Notices" on page 67.

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

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

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

© **Copyright International Business Machines Corporation 1997, 2001. All rights reserved.**

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

Contents

Welcome to DB2 Universal Database and

DB2 Connect Version 7. vii

Who Should Read This Book viii

How This Book Is Structured viii

DB2 Version 7.2 Enhancements ix

Chapter 1. Introducing Version 7. 1

Ease of Use 1

Launchpads 1

First Steps 1

Quick Tour. 2

Information Center 2

Wizard Enhancements 2

Information Feedback 2

New Products and Packages 3

DB2 Warehouse Manager 3

QMF for Windows 3

DB2 Query Patroller. 4

DB2 Relational Connect 4

Workgroup Edition on UNIX. 5

DB2 OLAP Starter Kit 5

DB2 Spatial Extender 5

DB2 Universal Database Text Information

Extender 5

DB2 Net Search Extender 6

Application Development Client 6

DB2 Life Sciences Data Connect. 6

Migration to Version 7 6

Migrating from Visual Warehouse Version

3.1 or Version 5.2. 7

Chapter 2. e-business Enhancements 9

XML Extender Support. 9

Web Services 9

Net.Data Enhancements 9

Stored Procedures 10

Stored Procedure Builder. 10

MQSeries Integration 10

MQSeries Assist Wizard 10

OLE DB Assist Wizard 11

Nested Stored Procedure Support 11

SQL Procedure Language 11

Visual Basic Stored Procedures 11

Calling Stored Procedures 11

Stored Procedure Samples 11

Java Stored Procedure Class Loader 12

DB2 Universal Database on Linux. 12

DB2 Universal Database Enterprise Edition

for Linux on S/390. 12

DB2 Universal Database Enterprise -

Extended Edition on Linux 12

JDBC 2.0 for Linux and HP-UX 13

Raw I/O for DB2 Universal Database for

Linux 2.4 13

Chapter 3. Business Intelligence

Enhancements. 15

Data Warehouse Center 15

Star Schema Builder 15

Process Modeler 15

Enhanced Support for Creating and

Changing Warehouse Target Tables 16

Name and Address Cleansing Support 16

Mandatory Fields 16

OLAP Capability 16

OLAP Starter Kit 16

Relational OLAP Capabilities 16

Business Intelligence Tutorial 17

e-Video Central Tutorial 17

New Warehouse Sources 18

DB2 Warehouse Manager Connector for i2

TradeMatrix BPI 18

DB2 Warehouse Manager Connector for

SAP R/3 18

DB2 Warehouse Manager Connector for

the Web 19

MQSeries Message Queues 19

Microsoft OLE DB and Data Transaction

Services Support 19

Metadata Interchange 20

Common Warehouse Metamodel Support 20

IBM ERwin Metadata Extract Program 20

Interim Commits when Importing

Warehouse Metadata 20

Additional Metadata Templates 20

Heterogeneous Distributed Query. 21

Access to New Database Management

Systems 21

New Support for Accessing Heterogeneous

Data from SQL Procedures 21

Query Patroller	22
QMF	22
Global Snapshot.	22
Convert ODBC (Dynamic SQL) to Static SQL	23
Temporary Table Support	23
Object Relational Features	23
Structured Types	23
Transform Functions	24
SQL-Bodied Functions	24
Dynamic Compound Statements	24
Variables and Control Flow in Triggers and	
SQL Functions	24
Typed Tables.	25
User-Defined Extended Index Types	25
Replication	26
Data Links	26
Capture and Apply	26
Data Link Replication Support on the	
Solaris Operating Environment.	27
Updateable Partitioning Keys	27
Additional Enhancements	27

Chapter 4. Data Management Enhancements. 29

DB2 Data Links Manager	29
Native OLE DB Support	29
SQL Assist	29
Closing Log After Backup	30
On Demand Log Archiving	30
Identity Column Support	30
Larger Index Key	30
Externalized Savepoint Support	31
Increased Log Limit to 32GB	31
Block Transactions When Log Directory Is	
Full	31
Rename a Table Space.	31
Database Buffer Pool	31
Windows 2000 Address Windowing	
Extensions (AWE)	31
64-Bit Support on AIX, HP-UX, and the	
Solaris Operating Environment.	32
Increased Locklist Limit	32
Unicode Enhancements	32
Encrypting and Decrypting String Data	33
Database Recovery.	33
Suspended I/O	33
Dual Logging	34
Parallel Recovery	34
Backup from Split Image.	35

Named Pipe Support for Backup and	
Restore Operations.	35
Incremental and Delta Backups	35
Sequence Support	36
Statement-Level Isolation Levels	36
New SQL Built-in Scalar Functions	37
Parallel Container Creation	37

Chapter 5. DB2 Family Enhancements . . . 39

Mobile Computing.	39
DB2 Everywhere.	39
>8 Character UserID Support	39
Pre-Version 7 Servers	40
Replication	40
DB2 Universal Database for OS/390	40
DB2 Universal Database for AS/400	40
Import and Export	40
Stored Procedures	40
LOAD Authority	40
USE OF TABLESPACE Authority	41
BIND Options	41
OS/390 in the Control Center	41
Generate DDL	41
OS/390 Utilities.	41
Windows 32-bit Operating Systems	42
Windows 2000	42
Visual C++ Application Development	42
Visual Studio Integration.	43
Administering DB2 UDB Workgroup Edition	
and DB2 UDB Personal Edition as Satellites	43
Running Compiled SQL Procedures	43
Cross-Platform Backup and Restore	43
DB2 Data Links Manager for the Solaris	
Operating Environment	44
TSM Integration with DB2 Data Links	
Manager	44

Chapter 6. DB2 Connect Enhancements. . . 45

DB2 Connect Managed Access	45
Improved Support for Microsoft	
Transaction Server (MTS) and COM+	
Technologies	45
Multisite Update Improvements	46
DB2 Connect Web Starter Kit	46
DB2 Connect for Linux	47
Support DCL SNA Product	47

Appendix A. Using the DB2 Library . . . 49

DB2 PDF Files and Printed Books.	49
DB2 Information	49

Printing the PDF Books	58	Appendix B. Notices.	67
Ordering the Printed Books	59	Trademarks	70
DB2 Online Documentation	60	Index	73
Accessing Online Help	60	Contacting IBM	79
Viewing Information Online.	62	Product Information	79
Using DB2 Wizards	64		
Setting Up a Document Server	65		
Searching Information Online	66		

Welcome to DB2 Universal Database and DB2 Connect Version 7

Database management software has evolved into the core of enterprise computing. As companies move into an Internet age of broad-band communications, databases must be able to store and serve huge multimedia files, manage ever increasing volumes of data, handle the tremendous growth in the number of users, deliver steadily improved performance, and support the next generation of applications. With its virtually unlimited ability to scale, its multimedia extensibility, its industry-leading performance and reliability, and its platform openness, DB2 Universal Database has helped lead this evolution.

DB2 Universal Database is the first multimedia, Web-ready relational database management system, strong enough to meet the demands of large corporations and flexible enough to serve medium-sized and small businesses. With DB2 Universal Database Version 7, IBM continues its role as the database innovator in:

- **e-business**

DB2 Universal Database powers the most demanding e-business applications, such as electronic commerce, enterprise resource planning, customer relationship management, supply-chain management, Web self-service, and business intelligence. It's about a scalable, industrial-strength database that can be the data management foundation for your evolution into e-business.

- **Business Intelligence**

DB2 Universal Database means using data assets to make better business decisions. It's about data access, data analysis, and decisions that help control costs, uncover new opportunities, boost market share, and increase customer loyalty.

- **Data Management**

DB2 Universal Database data management is more than simply running queries and applications. It's about where to store data, how to access it quickly, how to protect it against loss, and how to administer databases for optimum performance on your hardware and for your mix of applications.

- **DB2 Family**

DB2 Universal Database continues to meet the demands of today's heterogeneous computing environments. It's about open solutions that can access and integrate data from multiple, geographically separated sources on different platforms.

Universal Database is an ambitious name. It suggests a product that is designed to be used for a variety of purposes and in a variety of environments, and that is a good description of DB2 Universal Database. Each new release of DB2 Universal Database builds on the strong foundation of the previous version. This book describes the new DB2 functions and enhancements available with DB2 Universal Database Version 7.

Who Should Read This Book

This book is for current DB2 and DB2 Connect owners, or owners of related products such as DataJoiner or Visual Warehouse, who are upgrading or considering upgrading to DB2 Version 7.

If you are considering DB2 Universal Database or DB2 Connect for the first time, you should start by reading the appropriate *Quick Beginnings* book for your operating system to get basic DB2 information.

How This Book Is Structured

This book begins with an overview of some of the major DB2 enhancements for Version 7 and then describes these new features and enhancements.

Chapter 1. Introducing Version 7 describes the main DB2 product enhancements available with DB2 Universal Database Version 7 such as additional platform support for existing products, and describes the migration support provided to help you migrate from previous versions of DB2 to Version 7.

Chapter 2. e-business Enhancements describes the enhancements for building and running e-business applications that leverage existing data and help to transform business into e-business.

Chapter 3. Business Intelligence Enhancements describes how the ability to copy, consolidate, and analyze operational data has been enhanced for better business decisions that can often lead to competitive advantage.

Chapter 4. Data Management Enhancements describes the enhancements in DB2 function and performance.

Chapter 5. DB2 Family Enhancements describes how DB2 continues to expand its support for users from mainframe systems through to the newest generation of handheld devices.

Chapter 6. DB2 Connect Enhancements describes the main enhancements available with Version 7.2.

Appendix A. Using the DB2 Library describes the DB2 library; including books and online help.

Appendix B. Notices contains notice and trademark information.

DB2 Version 7.2 Enhancements

Responding to customer needs and the demands of the marketplace ensures that DB2 Universal Database continues to set the pace in the ever-changing enterprise computing environment with even further enhancements.

This book provides an overview of these enhancements. You will find that the descriptions of the enhancements are marked with a vertical bar to the left of the change. DB2 Version 7.2 includes:

- e-Business Enhancements, such as:
 - MQ-Assist wizard
 - MQ Series Integration
 - DB2 Enterprise - Extended Edition on Linux
 - DB2 Universal Database Enterprise Edition for Linux on S/390
 - DB2 Connect Enterprise Edition for Linux on S/390
 - DB2 Connect Web Starter Kit
 - JDBC 2.0 support on HP-UX and Linux
 - Exploitation of the Linux 2.4 kernel
- Business Intelligence Enhancements, such as:
 - New warehouse sources, such as i2, SAP R/3, web clickstream data, MQSeries messages queues, and Microsoft OLE DB objects
 - The Data Warehouse Center expanding its role as a tool integration hub for data quality and cleansing tools by adding support for Trillium
 - Enhanced support for creating and changing warehouse target tables
 - Distributed query access to new data sources, such as Sybase and Microsoft SQL Server, and to Oracle data from additional DB2 systems, on platforms such as the Solaris Operating Environment and Linux
 - Distributed query access from SQL procedures
 - Additional metadata interchange, including importing ERwin metadata and support for Common Warehouse Metamodel XML objects
 - Datalink Replication on the Solaris Operating Environment
 - Procedural logic in stored procedures, triggers, and SQL functions for enhanced cleansing and transformation capabilities
 - Dynamic compound SQL statements
 - Updateable partitioning keys in partitioned database environments
- Data Management Enhancements, such as:

- On-demand log archiving
- Alter identity columns
- 64-bit support on AIX, HP-UX, and the Solaris Operating Environment
- Support for Windows 2000 Address Windowing Extensions (AWE)
- Unicode enhancements
- Encryption and decryption of string data through user-defined functions (UDFs)
- Incremental and delta backup
- Sequence support
- Parallel container creation
- DB2 Family Enhancements, such as:
 - The ability to run SQL procedures on servers without compilers
 - DB2 Data Links Manager for the Solaris Operating Environment

The *DB2 Release Notes* describe how these enhancements are implemented. For your convenience and ease of reference, the enhancements are documented in the *DB2 Release Notes* by the appropriate DB2 book, chapter, and section. For example, changes to a command are found in the *Command Reference* section.

The Version 7.2 *DB2 Release Notes* are available in PDF format and can be found on the Publications CD-ROM with the DB2 manuals.

Note: The *DB2 Release Notes* are also available in ASCII and HTML on the product CDs.

Chapter 1. Introducing Version 7

This section provides a brief summary of the enhancements for the latest version of DB2 Universal Database, including the support DB2 provides for migrating to Version 7 from previous versions of the product.

Ease of Use

This section describes some of the Version 7 enhancements that make DB2 Universal Database even easier to use.

Launchpads

DB2 Universal Database provides launchpads to assist you with some of your DB2 tasks.

Installation Launchpad

When you insert the DB2 Universal Database Version 7 product CD-ROM, the auto-run feature no longer takes you to the first screen in the DB2 setup program. Instead, it initiates the DB2 Launchpad. The DB2 Launchpad is a window that gives you the following options:

- Release Notes
- Installation Prerequisites
- Quick Tour
- Install
- Exit

The DB2 Launchpad gives you access to needed information from the CD-ROM before starting a DB2 installation.

Warehouse Launchpad

The first time you open the Data Warehouse Center, the Warehouse Launchpad also opens. This launchpad simplifies the job of populating a warehouse by leading you through the related tasks. You can use the launchpad whenever you need to create warehouse objects by selecting it from the Data Warehouse Center menu.

First Steps

In Version 7, the First Steps tool that is used to create and view the SAMPLE database on a DB2 server has been extended to allow client options. On a client, you have the option to connect to the SAMPLE database (using the Client Configuration Assistant) or create a SAMPLE database on a remote DB2 server.

You can also create sample databases for warehousing and online analytical processing (OLAP). You will need these databases for the Business Intelligence tutorial (see “Business Intelligence Tutorial” on page 17).

As First Steps is a Java application, it can be used on all DB2-supported operating systems, provided that the appropriate Java Run-Time Environment (JRE) is installed on the workstation. Refer to the DB2 Universal Database *Quick Beginnings* manuals for your operating system for more information.

Quick Tour

The DB2 Universal Database Quick Tour is a high-level presentation that, in a short time (20 to 25 minutes), introduces you to changes, features, benefits, capabilities, and sources of information associated with DB2 Universal Database. It has a special emphasis on e-business, business intelligence (BI), and data management. The Quick Tour navigation is flexible, giving you control over your learning experience.

Information Center

The Information Center, accessible by selecting Help → Information Center from the Control Center, is now easier to use. In the Information Center you can now view information by subject or by category. The Information Center now contains a link to search the entire DB2 library.

You can also access help with the **db2ic** command on UNIX, or from the Start menu on Windows operating systems.

Wizard Enhancements

All wizards (called SmartGuides in previous versions) have a new look with a Table of Contents (TOC) on the left side. You can resize the TOC to make it larger so that all text can be displayed or hidden altogether. (Some wizards contain an Introduction page that introduces the wizard and a Summary page that summarizes the information that has been entered).

Help or instructions can be found on each page of the wizard, in the infopops for each field or control, and in the links that are identified as underlined text. A *Hints* feature is also available.

Information Feedback

When you have a comment about the information you’re reading, you can now send it directly to the DB2 Universal Database team.

From any HTML DB2 document (books or online help), you can click the Feedback link in the top frame of your browser. This takes you to a form on the World Wide Web where you can submit issues regarding the DB2 Universal Database documentation.

Note: The DB2 Universal Database team will read all your feedback but cannot respond directly to you. Use normal channels to reach Customer Support or Customer Service.

New Products and Packages

This section contains information on products and packages that have joined the DB2 family or have changed for Version 7.

DB2 Warehouse Manager

This new offering brings together the tools to build, manage, govern, and access DB2 data warehouses. The DB2 Warehouse Manager simplifies and speeds warehouse prototyping, development, and deployment. It gives the data center the control for governing queries, analyzing costs, managing resources, and tracking usage. It helps satisfy user requirements for finding, accessing, and understanding information. It provides flexible tools and techniques for building, managing, and accessing the warehouse. And it meets the most common reporting needs for enterprises of any size.

The DB2 Warehouse Manager adds to the basic warehouse and analytical functions available in DB2 Universal Database by providing:

- Additional warehouse scalability through warehouse agents co-located with the database. Warehouse agents manage the flow of data between warehouse sources and warehouse targets.
- Advanced transformations using Java stored procedures and user-defined functions including cleaning data, pivoting tables, generating keys, and more.
- An integrated business information catalog to guide users to relevant information that they can use for decision making.
- Sophisticated query governing and workload distribution.
- Query reporting that satisfies the common reporting needs of most enterprises.
- DB2 Warehouse Manager Connectors, which extend your access to web clickstream data, and to data that is managed by SAP R/3 and i2 TradeMatrix BPI. See “New Warehouse Sources” on page 18 for more information.

The offering applies to DB2 Universal Database Enterprise Edition and Enterprise - Extended Edition only.

QMF for Windows

QMF for Windows is included with the DB2 Warehouse Manager. It provides a multipurpose query tool for business reporting, data sharing, server resource protection, robust application development, and native connectivity to all of the DB2 workstation platforms.

QMF for Windows provides native support for TCP/IP connectivity to Version 5 and later. It is architecturally coupled with DB2 using Distributed Relational Database Architecture (DRDA). Support for a heterogeneous database environment is available through DB2 DataJoiner, which allows users to access other vendor and non-relational workstation databases. See “QMF” on page 22 for more Version 7 information.

DB2 Query Patroller

Query Patroller for DB2 Universal Database is now available for Enterprise Edition and provides support for HP-UX and NUMA-Q, in addition to Windows 32-bit operating systems, AIX, and Solaris platforms. This allows you to manage, govern, and schedule business analysis queries and to more fully utilize the resources available on those DB2 servers. See “Query Patroller” on page 22 for other Version 7 enhancements.

DB2 Query Patroller is now part of the Warehouse Manager offering.

DB2 Relational Connect

DB2 Relational Connect is used in a federated system to query and retrieve data located on other DBMSs, such as Oracle, Sybase, and Microsoft SQL Server. SQL statements can refer to multiple DBMSs or to individual databases in a single statement. For example, you can join data located in a DB2 UDB table, an Oracle table, and a Sybase view. See “Heterogeneous Distributed Query” on page 21 for more information on distributed queries.

In DB2 Version 7.1, DB2 Relational Connect was available for Oracle on the Windows NT and AIX platforms. In DB2 Version 7.2, the Oracle support has been enhanced to include additional operating systems, and support for Sybase and Microsoft SQL Server data sources has been added. Supported DBMSs include Oracle, Sybase, Microsoft SQL Server and members of the DB2 Universal Database family (such as DB2 for OS/390, DB2 for OS/400, and DB2 for Windows). DB2 Relational Connect can also be used with the Data Warehouse Center to improve select performance from Oracle, Sybase, and Microsoft SQL Server sources for loading a DB2 data warehouse (see “Native OLE DB Support” on page 29).

DB2 Relational Connect operates under some restrictions. Distributed requests are limited to read-only operations in DB2 Version 7. In addition, you cannot execute utility operations (such as LOAD, REORG, REORGCHK, IMPORT, and RUNSTATS) against nicknames. You can, however, use a pass-through facility to submit DDL and DML statements directly to DBMSs using the SQL dialect associated with that data source. For more information about federated database concepts, see the *Administration Guide: Planning*.

DB2 Relational Connect is an optional product that can be used with DB2 Universal Database Enterprise Edition, DB2 Universal Database Enterprise -

Extended Edition, DB2 Connect Enterprise Edition, and DB2 Connect Unlimited Edition for Windows NT, AIX, and Linux operating systems, and the Solaris Operating Environment.

Workgroup Edition on UNIX

In Version 7, a Workgroup Edition of DB2 is available for all DB2 supported UNIX-based platforms (except for NUMA-Q).

DB2 OLAP Starter Kit

DB2 Universal Database includes the DB2 OLAP Starter Kit which is functionally based on DB2 OLAP Server Version 7. The DB2 OLAP Server Starter Kit lets users of DB2 Universal Database explore and apply online analytical capabilities before acquiring the full product, DB2 OLAP Server, for more complex analysis or deployment to large groups of users. Applications developed by the DB2 OLAP Starter Kit can continue to be extended and used in the full product. Both the DB2 OLAP Starter Kit and DB2 OLAP Server are based on Essbase, the market-leading OLAP technology from Hyperion Solutions. See “OLAP Capability” on page 16 for more information.

DB2 Spatial Extender

The DB2 Spatial Extender can now be hosted directly by DB2 Universal Database. This facilitates the inclusion of spatial attributes, such as distance or driving time, into business analyses. This integration allows spatial data to exploit the performance benefits available with a database management system (DBMS). DB2 Spatial Extender conforms to the OpenGIS Consortium (OGS) and ISO standards for storing, indexing, and querying spatial data. See the *Spatial Extender User's Guide and Reference* for more information.

DB2 Universal Database Text Information Extender

DB2 Universal Database Text Information Extender provides a new way to search text documents using an SQL query. Supporting HTML and XML formats, it combines the power of a fast text search engine with the intelligent strategies of the DB2 Optimizer. Built on the success of the current extenders, DB2 Universal Database Text Information Extender augments, and begins to combine, the key functionality of the DB2 Text Extender and the DB2 Net Search Extender into a single, integrated product for today's information-orientated marketplace.

Using the DB2 Net Search Extender text search engine and the DB2 Text Extender search interface, DB2 Universal Database Text Information Extender provides the baseline for a new generation of text retrieval products. The result is an easy-to-use and highly integrated way to search documents in your database, or external documents using DB2 Data Links Manager. For more information see <http://www.software.ibm.com/data/db2/extendere>.

DB2 Net Search Extender

DB2 Universal Database Net Search Extender provides a high performance method of searching text stored in DB2 databases. Searching with it can be particularly advantageous in Internet applications where performance is an important factor.

Net Search Extender can add the power of fast full-text retrieval to Net.Data, Java, and CLI applications. Its features let you store unstructured text documents of up to 2 gigabytes in databases. It offers application developers a fast, versatile, and intelligent method of searching through such documents.

Application Development Client

The Software Developer's Kit is now called the DB2 Application Development Client.

DB2 Life Sciences Data Connect

When life sciences users need information to do their jobs, they often combine data from multiple sources to get what they need. However, in the life sciences and biotechnology industries, vast amounts of data reside in specialized data sources with proprietary query capabilities. These differences in function make it difficult for users to collect the data in a single source and format to work with it.

IBM DB2 Life Sciences Data Connect is a database middleware system that allows you to work with a virtual database, whose underlying data can be stored in multiple life sciences industry data sources. Using DB2 Life Sciences Data Connect, you run a single query against the virtual database, and are able to make connections that you might not otherwise find by querying each data source individually.

DB2 Life Sciences Data Connect is part of the IBM Sciences Solution's DiscoveryLink offering.

Migration to Version 7

DB2 provides migration support to allow you to convert back-level DB2 databases and instances to a format usable by DB2 Universal Database Version 7. Database migration is not required to move from DB2 Version 7.1 to Version 7.2.

The following DB2 releases are supported by the DB2 Universal Database Version 7 migration process:

Migration to DB2 Universal Database Version 7

- DB2 for AIX Version 5.x and Version 6
- DB2 for HP-UX Version 5.x and Version 6

- DB2 for OS/2 Version 5.x and Version 6
- DB2 for Solaris Version 5.x and Version 6
- DB2 for Windows NT Version 5.x and Version 6
- DB2 for NUMA-Q Version 6
- DB2 for Linux Version 6

Migration to DB2 Enterprise - Extended Edition Version 7

- DB2 for AIX Version 5.x and Version 6
- DB2 for HP-UX Version 6
- DB2 for Solaris Version 5.x and Version 6
- DB2 for Windows NT Version 5.x and Version 6
- DB2 Enterprise - Extended Edition for AIX Version 5.x and Version 6
- DB2 Enterprise - Extended Edition for Solaris Version 5.x and Version 6
- DB2 Enterprise - Extended Edition for Windows NT Version 5.x and Version 6
- DB2 for NUMA-Q Version 6

Note: DB2 migration does not support migration across operating systems. You can only migrate from a previous version of DB2 on one operating system to DB2 Version 7 on the same operating system. For example, if your current database server is DB2 Version 5 for Windows 32-bit operating systems, you can only migrate to DB2 Version 7 for Windows 32-bit operating systems.

Some incompatibilities exist between DB2 Version 7 and previous releases of DB2. An incompatibility is a part of DB2 that works differently than it did in a previous release. If used in an existing application, the incompatibility will produce an unexpected result, necessitate a change to the application, or reduce performance. For a list of incompatibilities, see the *Incompatibilities Between Releases* appendix in the *Administration Guide: Planning*. This appendix describes the incompatibilities and their symptoms or effects, and suggests possible resolutions.

Depending on your platform, there may be certain pre-installation and post-installation tasks that you must perform to migrate successfully to the Version 7 format. For more information on migrating to DB2 Version 7, refer to the *Quick Beginnings* manuals for your platform.

Migrating from Visual Warehouse Version 3.1 or Version 5.2

The function provided by the Visual Warehouse product has been integrated into DB2 Universal Database and DB2 Warehouse Manager. This function is accessed in DB2 Universal Database through the Data Warehouse Center. For detailed planning and migration information, see the book: *Migrating to DB2 UDB Version 7.1 in a Visual Warehouse Environment*, SG24-6107.

Chapter 2. e-business Enhancements

Behind every significant e-business application is a database, and the robustness of that database and its associated components will determine your online success. With DB2 Universal Database, e-business is really enterprise computing, transformed through the power of the Internet.

XML Extender Support

DB2 now allows you to store eXtended Markup Language (XML) documents as a new column data type. You also have the ability to decompose and store XML in its component parts as columns in multiple tables. In both cases, indexes can be defined over the element or attribute of an XML document for fast retrieval. In addition, text search and section search can be enabled on the XML column or its decomposed part using text extenders. You can also formulate an XML document from your existing DB2 tables for data interchange in a business-to-business environment. Net.Data and XML Extender can be used to generate XML documents from DB2 and distributed to consumers for viewing with a browser.

Note: XML documents may be received from files or via MQSeries messages. XML documents may also be generated as files or as MQSeries messages. See “MQSeries Message Queues” on page 19 for more information.

Web Services

Web services, via Simple Object Access Protocol (SOAP), make it possible for heterogeneous applications to request services over the web. The standards-based IBM approach provides access to DB2 data and stored procedures without programming. These web services are deployed in a WebSphere runtime environment.

For more information on this support and related tools, see <http://www.ibm.com/software/data/db2/extenders/xmlext/>. For more information on SOAP, see “e-Video Central Tutorial” on page 17.

Net.Data Enhancements

Net.Data, which is a web application builder designed to access DB2 data, now provides XML output, XHTML compatibility, file upload capability, SQL statement nesting, and more. See <http://www.ibm.com/software/data/net.data> for more information and complete documentation.

Stored Procedures

Application performance can be improved with these enhancements to the stored procedure capabilities of DB2 Universal Database.

Stored Procedure Builder

Stored Procedure Builder has been enhanced to simplify the process of building your stored procedures and also to make it easier to debug and deploy them in your enterprise:

- AIX and Solaris clients are now supported.
- The Stored Procedure Builder can be launched from the Control Center.
- Java stored procedures can be built on DB2 for OS/390.
- SQL Procedure Language support is available across the entire DB2 Family. See “SQL Procedure Language” on page 11 for more information.

For more information on the Stored Procedure Builder see the *Online Help*.

MQSeries Integration

DB2 Version 7.2 provides a basic set of functions for integrating MQSeries messaging operations within SQL statements. These functions support a variety of messaging models, including simple, one-way messages, request/reply, and publish/subscribe. Using these basic functions, it is possible to support a wide range of applications, ranging from simple event notification to operational data store creation.

More information is available in the DB2 7.2 Release Notes and at www.software.ibm.com/software/data/integration/MQSeries. Additional support for semi-structured messages is provided by the MQSeries Assist wizard (see “MQSeries Assist Wizard” for more information). Additional support for XML messages is provided by new functionality in the DB2 XML Extender (see <http://www.ibm.com/software/data/db2/extenders/xmlext/>).

The new DB2 Version 7.2 Data Warehouse Center also provides MQSeries Integration. Further details on this new functionality may be found at “MQSeries Message Queues” on page 19.

A copy of MQSeries is now provided to support the use of these exciting new capabilities on Windows 2000, Windows NT, AIX, HP-UX, and the Solaris Operating Environment.

MQSeries Assist Wizard

DB2 Version 7.2 provides a new MQSeries Assist wizard. This wizard creates a table function that reads from an MQSeries queue using the MQSeries user-defined functions, which are also new in Version 7.2.

The wizard can treat each MQSeries message as a delimited string or a fixed length column string depending on your specification. The created table function parses the string according to your specifications, and returns each

MQSeries message as a row of the table function. The wizard also allows you to create a view on top of the table function and to preview an MQSeries message and the table function result.

OLE DB Assist Wizard

DB2 Version 7.2 provides a new OLE DB Assist wizard. This wizard helps you to create a table function that reads data from another database provider that supports the Microsoft OLE DB standard. You can optionally create a DB2 table with the data that is read by the OLE DB table function, and you can create a view for the OLE DB table function.

For more information on the OLE DB Assist wizard, see the *DB2 Release Notes*.

Nested Stored Procedure Support

DB2 now delivers nested stored procedure support that not only simplifies the migration of applications from other database vendors, but also allows application developers to use and exploit stored procedures in a more effective way. See the *Application Development Guide* for more information.

SQL Procedure Language

SQL Procedure Language support is now available for DB2 on Windows, OS/2, and UNIX, in addition to DB2 Universal Database for OS/390 and DB2 Universal Database for OS/400. It supports building DB2 stored procedures that employ SQL in a manner consistent with the Persistent Stored Module definition of the ANSI SQL99 standard. See the *Application Development Guide* and the *Application Building Guide* for more information.

Visual Basic Stored Procedures

In addition to Java and the stored procedure language, DB2 Universal Database now enables you to write stored procedures using Microsoft Visual Basic. Stored procedures, written using Microsoft Visual Basic, can be deployed on DB2 Universal Database servers running on Windows 32-bit operating systems. You can specify OLE as the language in which the stored procedure is implemented and call it from a client application written in any of the languages supported by DB2. See “Native OLE DB Support” on page 29 for more information on DB2 OLE support.

Calling Stored Procedures

Stored procedures can now be called from the command line processor (CLP), using the CALL command. See the *SQL Reference, Volume 1* for more information.

Stored Procedure Samples

An enhanced set of stored procedure samples is available with the Application Development Client. The samples are written using embedded SQL in C/C++, CLI in C, JDBC in Java, and SQL procedures. They demonstrate the capability of client applications in one language to call a stored procedure written in

another language. Separate source files for CREATE PROCEDURE statements, client applications, stored procedures, and DROP PROCEDURE statements make it easy to copy and modify the sample code to create your own stored procedures.

Java Stored Procedure Class Loader

With the Java class loader, DB2 Universal Database enables you to improve the performance of your Java stored procedures and increase the availability of your production servers. The Java class loader enables you to perform online replacement of your Java stored procedure classes. When you issue the CALL SQLJ.REFRESH_CLASSES() statement, DB2 continues to run, and the Java Virtual Machine (JVM) remains loaded in memory, ready to execute the refreshed stored procedure classes. Stored procedures can use the new classes without the need to stop the database manager. See the *Application Development Guide* for more information.

DB2 Universal Database on Linux

The DB2 Universal Database Version 7.2 offering enhances the capabilities of internet computing on the Linux operating system with the following:

- DB2 Universal Database Enterprise Edition for Linux on S/390
- DB2 Universal Database Enterprise - Extended Edition on Linux

DB2 Universal Database Enterprise Edition for Linux on S/390

DB2 Universal Database Enterprise Edition for Linux on S/390 provides you with the ability to exploit the benefits of DB2 and Linux on the IBM zSeries (390) hardware. You can perform administrative functions, such as backup and recovery, with reduced cost and overhead. DB2 also provides the reliability that high-speed internet transactions require. DB2 Universal Database Enterprise Edition is available on 32-bit Intel-based Linux and Linux on S/390.

DB2 Universal Database Enterprise - Extended Edition on Linux

DB2 Universal Database Enterprise - Extended Edition (DB2 EEE) on Linux makes query performance unparalleled and more powerful than before. DB2 EEE on Linux adds the power of DB2's shared-nothing architecture to Linux, allowing you to cluster Linux machines for scalability.

Configuring DB2 EEE to execute in a shared-nothing architecture means that each Linux machine has exclusive access to its own disks and memory, and does not compete with other machines for resources. Being able to partition a database across multiple machines with DB2 EEE allows you to perform complex parallel processing operations on platforms such as AIX, HP-UX, Windows NT, Linux, and the Solaris Operating Environment.

The cost-based query optimizer, provided by DB2 EEE, determines the most efficient way to handle multiple, complex queries, and selects an execution plan that has the lowest overhead. Response time for both decision support and online transaction processing (OLTP) is fast and cost effective.

For further information on how DB2 Universal Database Enterprise - Extended Edition can help you, see the *DB2 for UNIX Quick Beginnings*.

JDBC 2.0 for Linux and HP-UX

JDBC 2.0 support is now available on Linux, Linux on S/390, and HP-UX, in addition to the previously supported AIX and Windows platforms, and the Solaris Operating Environment. For more information on using JDBC 2.0 to access your databases, see the *DB2 Release Notes*.

Raw I/O for DB2 Universal Database for Linux 2.4

Table spaces created in raw disk partitions offer increased performance improvements. DB2 Version 7.2 provides the ability to create table spaces and logs on raw partitions if you are using the Linux 2.4 kernel.

Note: DB2 does not support raw devices on Linux for OS/390.

Chapter 3. Business Intelligence Enhancements

DB2 Universal Database includes the built-in foundation for creating, populating, storing, and maintaining data warehouses and OLAP cubes. DB2 Warehouse Manager extends those basic functions with greater scalability, manageability, and accessibility.

Data Warehouse Center

The power of Visual Warehouse and the simplicity of the DB2 Control Center have been merged to provide a single, new user interface for business intelligence customers. You can use the Data Warehouse Center to register and access data sources, define data extraction and transformation steps, populate data warehouses, automate and monitor warehouse management processes, and manage and interchange metadata. The Warehouse launchpad simplifies the job of populating a warehouse by leading you through the related tasks. The Data Warehouse Center supports full refresh and incremental update data movement options including leveraging the power of IBM's integrated data replication functions. Integrated SQL Assist capabilities (see "SQL Assist" on page 29) help customers define over 100 transformations using the rich SQL available with DB2. Users can also build custom transformation routines using Stored Procedure Builder (see "Stored Procedure Builder" on page 10).

See the *Data Warehouse Center Administration Guide* and the *Data Warehouse Center Application Integration Guide* for more warehousing information.

Star Schema Builder

Within the Data Warehouse Center, the warehouse schema modeler is a specialized tool for generating and storing schema associated with a data warehouse. Any schema resulting from this process can be passed easily as metadata to the OLAP Integration Server, which is part of the OLAP Starter Kit and the DB2 OLAP Server. See the *Data Warehouse Center Administration Guide* for more information.

Process Modeler

The process modeler allows users to graphically link the steps needed to build and maintain data warehouses and dependent data marts. Dependency relationships, conditional processing, and notifications can all be included in the model. Processes can be scheduled for one-time or repeating execution, or triggered by internal or external processes.

Enhanced Support for Creating and Changing Warehouse Target Tables

With DB2 Version 7.2, the Data Warehouse Center provides additional flexibility for creating and changing warehouse target tables. When a change to a warehouse step creates additional columns, you can have the Data Warehouse Center automatically create the new columns in the target table and map the results of the step to these new columns. You can also now choose the name and table space for the target table.

Name and Address Cleansing Support

The Data Warehouse Center expands its role as a tool integration hub for data quality and cleansing tools by adding support for Trillium. Trillium can be used to perform name and address cleansing, matching, merging, and demographic augmentation processes. Using the Data Warehouse Center, you can seamlessly include name and address cleansing operations in warehouse processes, which can be scheduled or run on demand. The Data Warehouse Center already provides integration with tools from Vality and Evolutionary Technologies, Inc.

Mandatory Fields

The Data Warehouse Center now displays red borders on required fields. The red borders alert you to information such as database names, user IDs, or passwords that are necessary to define the Data Warehouse Center objects. When you enter the required information, the borders disappear.

OLAP Capability

DB2 Universal Database provides powerful tools for business analysis, including online analytical processing (OLAP) tools.

OLAP Starter Kit

DB2 Universal Database provides you with integrated online analytical processing capabilities through the DB2 OLAP Starter Kit. It allows users to easily define multidimensional applications. The applications can be linked and populated with warehouse data, and shared among a work group. The Starter Kit prepares you for more complex OLAP applications which require the full function DB2 OLAP Server product.

The Starter Kit is based on the DB2 OLAP Server product and Hyperion's industry-leading Essbase technology.

Relational OLAP Capabilities

New SQL functions for moving aggregates, such as moving average and moving sum (for example, the balance column of a bank statement), are now available for analytical queries. These additional computational functions expand your OLAP capabilities and improve the performance of these typically challenging queries.

Business Intelligence Tutorial

The DB2 Universal Database tutorial for business intelligence is an online HTML document that provides an end-to-end guide to typical business intelligence tasks. The lessons in the tutorial provide step-by-step instructions for data warehousing and OLAP tasks, using the sample databases that come with DB2. The tutorial is accessible from the First Steps window, the Information Center, and the Data Warehouse Center and OLAP components of DB2 Universal Database.

e-Video Central Tutorial

A new DB2 tutorial, called eVideoCentral, is available for download at <http://www.ibm.com/software/data/developer/samples/evideo>. The purpose of the DB2 eVideoCentral tutorial is to give an example solution for a company that provides business services to another company over the web. This concept is known as a Business-to-Business application. The e-VideoCentral tutorial demonstrates the integration of DB2, WebSphere, and MQSeries.

Many technologies are used to design and build eVideoCentral, including: the IBM WebSphere Application Server, VisualAge for Java, and the DB2 XML Extender interaction tool. A set of centralized services for individual retail video stores is accessed via SOAP (see *Simple Object Access Protocol* at <http://www.w3.org/2000/xml/>).

SOAP is a lightweight protocol for exchanging information in a decentralized, distributed environment. It is an XML-based protocol that consists of three parts:

- An envelope that defines a framework for describing what is in a message and how to process it
- A set of encoding rules for expressing instances of application-defined datatypes
- A convention for representing remote procedure calls and responses.

SOAP services can provide a mechanism for one company to access DB2 services from another company.

The eVideoCentral tutorial demonstrates the design, development, and implementation of simple data repository (insertion/modification) and query services. The query services use the existing DB2 XML Extender. The first version of eVideoCentral is provided via the web, and it includes: Java Servlets, JSP (Java Server Pages), schema for the DB2 database objects, and the accompanying tutorial documentation.

New Warehouse Sources

With DB2 Version 7.2, you can access many new and diverse warehouse sources. The DB2 Warehouse Manager Connectors, three new business intelligence offerings, extend your access to data that is managed by applications like SAP R/3 and i2 TradeMatrix BPI. The DB2 Warehouse Manager Connectors also provide the capability to bring web clickstream data into the data warehouse. Other new warehouse data sources include Microsoft OLE DB objects, Microsoft Data Transaction Services targets, and MQSeries message queue data, including messages that are formatted as XML documents.

DB2 Warehouse Manager Connector for i2 TradeMatrix BPI

The DB2 Warehouse Manager Connector for i2 TradeMatrix BPI allows you to:

- Import changes to the metadata of a deployed Business Process Intelligence (BPI) business model.
- Update the BPI distributed datamart and an OLAP cube for the deployed business model.

The DB2 Warehouse Manager Connector for i2 TradeMatrix BPI generates the steps that manage the loading of the i2 dimension and fact tables from existing input files (in standard BPI comma-separated variable (CSV) format), as well as loading of the OLAP cube. After you import the metadata for i2 sources, the Data Warehouse Center creates and populates the Processes folder for the i2 load steps. Because the i2 load steps are controlled by the Data Warehouse Center, you can control when and how often they run, or run them on demand.

The DB2 Warehouse Manager Connector for i2 TradeMatrix BPI runs on Microsoft Windows NT or Windows 2000. The i2 TradeMatrix business-model source can be on any supported platform.

DB2 Warehouse Manager Connector for SAP R/3

The DB2 Warehouse Manager Connector for SAP R/3 allows you to access and bring SAP business objects stored in SAP R/3 systems into a DB2 data warehouse. Business objects, along with business components, provide an object-oriented view of R/3 business functions. You can then use the power of DB2 and the Data Warehouse Center for data analysis, data transformation, or data mining.

You define the data extraction step using the Data Warehouse Center by simply clicking and dropping an SAP object onto the process modeler. When you define an SAP source, you see all the metadata about the SAP object, including key fields, parameter names, data types, precision, scale, length, and mandatory parameters. You also see all basic and detailed parameters associated with the SAP business object.

The DB2 Warehouse Manager Connector for SAP R/3 runs on Microsoft Windows NT or Windows 2000. (For Windows 2000, you should use Microsoft's service pak 1.) The SAP server can be on any platform.

DB2 Warehouse Manager Connector for the Web

The DB2 Warehouse Manager Connector for the Web allows you to extract data from an IBM WebSphere Site Analyzer (WSA) database, or webmart, into a data warehouse. The DB2 Warehouse Manager Connector for the Web includes a polling step that checks whether WSA has copied web traffic data from its data imports (log files, tables, and clickstream data) to the webmart. When this check is successful, an SQL step copies the web traffic data from the webmart into a warehouse target. You can then use the power of DB2 and the Data Warehouse Center for data analysis, data transformation, or data mining. You can also incorporate WebSphere Commerce data with the web traffic data for a more complete analysis of your web site.

After defining a web traffic source, you can define the web traffic polling step from the Data Warehouse Center by simply clicking and dropping a web object onto the process modeler.

The DB2 Warehouse Manager Connector for the Web runs on the same platforms as the DB2 Version 7.2 warehouse agent: Windows NT, Windows 2000, AIX, and the Solaris Operating Environment.

For more information about the DB2 Warehouse Manager Connectors, see the *DB2 Version 7.2 Warehouse Manager Installation Guide*.

MQSeries Message Queues

The Data Warehouse Center now enables you to access data from an MQSeries message queue as a DB2 database view. A wizard is provided to create a DB2 table function and the DB2 view through which you can access the data. (See "MQSeries Assist Wizard" on page 10, for more information.) Each MQSeries message is treated as a delimited string, which is parsed according to your specification, and returned as a result row.

In addition, MQSeries messages that are XML documents can be accessed as a warehouse source. Using the Data Warehouse Center, you can import metadata about XML documents from an MQSeries queue and a DB2 XML Extender Document Access Definition (DAD) file. The Data Warehouse Center uses this metadata to automatically create the warehouse target definition and the warehouse step to retrieve the XML documents from the queue.

Microsoft OLE DB and Data Transaction Services Support

The Data Warehouse Center now enables you to access data from an OLE DB provider as a DB2 database view. A wizard is provided to create a DB2 OLE DB table function and the DB2 view through which you can access the data. See "MQSeries Assist Wizard" on page 10, for more information.

Because Data Transaction Services (DTS) packages can be accessed as OLE DB sources, the wizard also allows you to create the view for a DTS package. When you access the view at run time, the DTS package executes and the target table of the DTS package is exposed as the created view.

Metadata Interchange

DB2 Version 7.2 enhances your ability to interchange metadata with the existing support of the Common Warehouse Metadata Interchange standard, and the addition of a number of new tools and offerings.

Common Warehouse Metamodel Support

With DB2 Version 7, metadata interchange was expanded to support the industry's Common Warehouse Metadata Interchange standard from the Object Management Group (OMG), facilitating solution integration among heterogeneous tools. The OMG standard has the support of industry leaders including IBM, Oracle, NCR, and Hyperion.

With DB2 Version 7.2, support has been added for importing and exporting Common Warehouse Metamodel XML objects. When you export a process containing a step that has a relationship to a step in another process (a "shortcut"), both processes are exported and the relationship is maintained. With the addition of new source support tags, you can now export SAP and WebSphere Site Analyzer (WSA) information saved in your Warehouse Center.

IBM ERwin Metadata Extract Program

The IBM ERwin Metadata Extract Program is a DB2 Version 7.2 time-saving tool that you can use to import metadata to IBM products, such as the Data Warehouse Center and the Information Catalog Manager (DataGuide). With the metadata extract program tag language files, you can create warehouse targets and schema, including star schema, from imported metadata. The metadata extract program also provides both graphical and command prompt interfaces.

Interim Commits when Importing Warehouse Metadata

With DB2 Version 7.2, warehouse metadata is imported and exported with interim commits. This feature is designed to improve performance when warehouse metadata is imported. As warehouse metadata is exported, a COMMIT tag will be written to the metadata file when a metadata object and all its required relationships have been exported. When the metadata file is imported, changes are committed to the warehouse control database when the COMMIT tag is detected and all objects since the last COMMIT have been validated.

Additional Metadata Templates

New templates have been added to communicate with the Data Warehouse Center. They are **primarykey.tag**, **primarykeyadditional.tag**, **foreignkey.tag**,

and **foreignkeyadditional.tag**. In addition, you can use the **commit.tag** to improve performance and monitor commit scope.

Heterogeneous Distributed Query

DB2 Universal Database and DB2 Connect users now have the power of distributed query across any DB2 family database or OLE DB source. This means that users and applications can use the DB2 Universal Database SQL syntax and APIs to access data that resides at heterogeneous data sources. With this functionality, users and applications have the capability of referencing multiple data sources in a single SQL statement. With DB2 Relational Connect, distributed queries can also include Oracle and, with DB2 Version 7.2, Sybase or Microsoft SQL Server databases (see “DB2 Relational Connect” on page 4).

This is the first phase of the integration of DB2 DataJoiner into DB2 Universal Database. DataJoiner is IBM’s middleware product for integrating heterogeneous data sources. See the *Administration Guide: Planning* and the *Administration Guide: Implementation* for more information.

Access to New Database Management Systems

DB2 Relational Connect Version 7.2 adds data access to these database management systems:

- Oracle databases from DB2 and DB2 Connect on the Solaris Operating Environment and Linux
- Sybase databases from DB2 and DB2 Connect on AIX and the Solaris Operating Environment
- Microsoft SQL Server databases from DB2 and DB2 Connect on AIX and Windows NT

Access to Oracle databases from DB2 on AIX and Windows NT continues to be supported.

With this new support, a single distributed query can now access data in any DB2 family database or OLE DB source, as well as any Oracle, Sybase, and Microsoft SQL Server for Windows NT database.

New Support for Accessing Heterogeneous Data from SQL Procedures

DB2 Version 7.2 allows you to query tables in heterogeneous databases from stored procedures. With Version 7.2, SQL procedures can use nicknames. You can define a nickname for a table in a heterogeneous database, and use that nickname in an SQL procedure.

Query Patroller

DB2 Query Patroller, now packaged inside the DB2 Warehouse Manager, can trap the SQL going to a DB2 Server by integrating the trap into the client code. This allows all dynamic SQL, regardless of operating system, to be managed, scheduled, and governed by Query Patroller.

A query retry mechanism allows jobs that have been aborted for various reasons to be resubmitted and run again to completion.

It is possible to issue a global start command that starts Query Patroller on all nodes. This gives you a single point of control for starting and stopping Query Patroller.

Query Patroller is available for DB2 Enterprise - Extended Edition and Enterprise Edition, and provides support for HP-UX and NUMA-Q, in addition to Windows 32-bit operating systems, AIX, and the Solaris Operating Environment.

QMF

An intuitive quick start interface allows Query Management Facility (QMF) for Windows users to build their queries and reports, or use the new Java-based query capability to launch queries from their favorite browser. Query results can be easily integrated with any OLE 2 desktop tools, including spreadsheets, charting and analysis tools, and desktop databases.

The QMF for Windows Administrator component prevents waste and abuse through detailed permissions organized on each server by group, schedule, or a combination of the two.

With QMF for Windows, application developers gain a robust Windows application programming interface that lets them rapidly build data access and update applications. It provides full exploitation of DB2 performance, SQL syntax, and advanced database performance techniques (for example, static SQL).

Global Snapshot

The database system monitor now allows you to monitor your DB2 Universal Database Enterprise - Extended Edition system from a single partition. It collects data and aggregates values across all partitions and returns a single result. This provides database administrators with a single point of control for monitoring their entire data warehouse. The database system monitor collects information on the operation and performance of database activities ranging

from reads and writes through to locks and deadlocks. See the *System Monitor Guide and Reference* for more information.

Convert ODBC (Dynamic SQL) to Static SQL

You can now convert your ODBC queries into static SQL. Running the queries as static SQL gives performance improvements if you repeatedly run the same ODBC queries.

For more information on converting ODBC queries into static SQL, see the **db2cap** command in the *DB2 Release Notes*.

Temporary Table Support

DB2 now delivers temporary table support. DB2 will create and use temporary tables within a single connection, providing you with performance gains for the type of query workloads that can make use of intermediary tables. See the *Administration Guide* for more information.

Object Relational Features

Abstract or structured data types are *type* mechanisms for modeling and storing complex objects in a relational database. Structured types may have multiple fields embedded within them, for example, a geometric shape (a list of Cartesian coordinates), or an employee (name, address, gender, birth date, and employee number) can be modeled and stored in DB2 Universal Database.

Structured Types

Structured type support has been extended to provide the ability to create tables with structured type columns. Additionally, structured types can be nested within a structured type. This means that the attributes of a structured type are no longer restricted to the base SQL types, they can now be of another structured type.

In Version 7, you can define functions with input parameters or parameters on the RETURNS clause that are structured types. You can also define *methods* for each structured data type, which permit the encapsulation of behavior with data. A method is defined very much like a function, but its use is strictly associated with structured types. It is essentially a routine with a structured type instance as its implicit first argument.

The Reorganize Table (REORG) and the db2look utilities can be used on tables with structured type columns. See the *Command Reference* for more

information on REORG and the *Administration Guide: Implementation and Administration Guide: Performance* for more information on structured types and db2look.

Transform Functions

Transform functions enable the use of structured type columns with user written programs. Transform functions convert the complex structure within a structured data type into an ordered set of its base SQL types. They also convert the base attributes back to their structured types. These transforms are required to move structured types in and out of a database. See the *Administration Guide: Implementation* for more information.

SQL-Bodied Functions

An SQL-bodied function contains simple SQL procedure statements within its body, which is embedded into the calling SQL (similar to a macro). This allows the query compiler to optimize the whole SQL statement including the SQL-bodied function. Structured data types make use of SQL-bodied functions for their transform functions (see “Transform Functions”) and methods (see “Structured Types” on page 23).

Dynamic Compound Statements

DB2 Version 7.2 provides dynamic compound statements, a new type of compound SQL, to help in reducing the database manager overhead and improve performance in handling requests over the network. Dynamic compound statements are ideal for short scripts involving little control flow logic, but significant data flow.

Within the dynamic compound statement, you can:

- Declare SQL variables, conditions associated with SQLSTATEs, and SQL procedural statements in sub-statements.
- Use several control flow logic statements, such as FOR, IF, ITERATE, and WHILE.

DB2 compiles the dynamic compound statement as a single statement. See the *DB2 Release Notes* for more information.

Variables and Control Flow in Triggers and SQL Functions

DB2 Version 7.2 provides the ability to perform procedural logic in stored procedures, triggers, and SQL functions through a number of SQL-controlled statements.

Prior to Version 7.2, triggers provided only simple sequencing – there was no conditional logic or loops. This enhancement to triggers allows you to migrate your applications to DB2 more easily. You will also benefit from the enhancement to SQL functions, such as SCALAR, TABLE, or ROW. For

example, you can either use SQL table functions with control logic in the Data Warehouse Center or SQL scalar functions to map other database management systems to DB2.

Version 7.2 enables the following control statements in triggers and SQL functions:

- Atomic compound statements
- SQL control statements:
 - FOR
 - GET DIAGNOSTICS
 - IF
 - ITERATE
 - LEAVE
 - WHILE
- SQL local variables

The control statements are integrated into the DB2 query compiler. For more information on SQL control statements and variables, see the *SQL Reference*. For more information on this Version 7.2 enhancement, see the *DB2 Release Notes*.

Typed Tables

You can now use referential integrity constraints and triggers with typed tables.

User-Defined Extended Index Types

User-defined extended index types allow you to create and apply your own logic for the three primary components that make up how an index works. The three components are index maintenance, index search, and index exploitation. You define index maintenance and index search with the CREATE INDEX EXTENSION statement. Index exploitation occurs in the evaluation of the search method.

You can now create indexes on spatial objects stored in a database. Only a structured type or a distinct type column can use the index extension to create a user-defined extended index type on these objects.

See the *Administration Guide: Implementation* and the *SQL Reference* for more information.

Replication

Replication is a process of maintaining a defined set of data in more than one location. DB2 DataPropagator can replicate changes between any DB2 relational databases, or by using DB2 DataJoiner, between DB2 and supported non-DB2 relational databases. Replication involves copying designated changes from one location (a source) to another (a target), and synchronizing the data in both locations.

Data Links

A data link is a link to a file with data like drawings, reports, and sound and other multimedia data. The size of these files can be very large. Accessing these types of files over a remote network would cause delays and increase network traffic. If the data does not change frequently and accessing the most current data is not a requirement, DB2 DataPropagator provides a good solution to this problem. However, since these files are not part of the database, a mechanism is required to replicate the files and ensure that the data and files are in sync at the target site. Data links replication can solve this problem.

With a data type of DATALINK, the database server has control over the remote files, and the files are treated as if they are part of the database. Data links replication replicates not only database data, but also files in a synchronized way.

In Version 7, DB2 Data Propagator replicates DATALINK columns and invokes a user exit routine to replicate the external files to which the DATALINK columns point. The user exit program maps the source data link value to a link in the target system and then copies the referenced file to the target file system. A sample user exit program is shipped with the product and it uses FTP to copy the files. See the *Replication Guide and Reference* for more information.

For more information on data links, see the *DB2 Data Links Manager Quick Beginnings*.

Capture and Apply

On Windows 32-bit operating systems, you can use the Capture and Apply programs on demand using the ASNSAT command. This command provides more flexibility when using DB2 DataPropagator functions and in handling error conditions.

You can start both the Capture and Apply programs from within an application by using the new asnCapture and asnApply application programming interfaces.

See the *Replication Guide and Reference* for more information.

If you are using DB2 DataPropagator Relational (DPropR) Version 1, you must migrate to Version 5 before migrating to Version 6 or Version 7.

Data Link Replication Support on the Solaris Operating Environment

Data link replication is now available on the Solaris Operating Environment as part of DB2 Version 7.2. It requires an FTP daemon that runs in the source and target DATALINK file system and supports the MDTM (modtime) command, which displays the last modification time of a given file. If you are using Version 2.6 of the Solaris Operating Environment, or any other version that does not include FTP support for MDTM, you need additional software such as WU-FTPD.

For more information on data links, see the *DB2 Data Links Manager Quick Beginnings*.

Updateable Partitioning Keys

DB2 Version 7.2 allows you to update the columns in a partitioning key. (A partitioning key exists only in a table spread over multiple database partitions.) Before DB2 Version 7.2, if you wanted to change the partitioning key, you needed two steps:

1. Delete the row.
2. Insert the row with the new key.

Each step impacted the log space requirements on both the database partition losing the data and the one gaining the data.

With DB2 Version 7.2, this can be done in one step with an update statement. In an online transaction processing (OLTP) environment, updateable partitioning keys provide performance improvements in data redistribution.

Additional Enhancements

For customers interested in additional Business Intelligence enhancements, you should also refer to the following Data Management Enhancements:

- “Closing Log After Backup” on page 30
- “Increased Log Limit to 32GB” on page 31
- “Identity Column Support” on page 30
- “Database Recovery” on page 33
- “Parallel Container Creation” on page 37.

Chapter 4. Data Management Enhancements

With each new release, DB2 Universal Database extends its capabilities, which in turn expands the abilities of the end user, administrator, and application programmer. As businesses flock to the Web and more workers take business on the road, DB2 Universal Database continues to enhance its performance, reliability, and scalability to meet the demand for powerful database solutions.

DB2 Data Links Manager

DB2 Data Links Manager capabilities have been expanded to allow you to use it in a more heterogeneous environment. DB2 Data Links Manager can now manage files that are stored in a DFS file system on AIX. Also, it now supports the use of the Veritas' XBSA interface for backup and restore using NetBackup. Data Links Manager is also available for the Solaris operating system.

For more information see the *DB2 Data Links Manager Quick Beginnings*. For information on replicating Data link files see, "Data Links" on page 26.

Native OLE DB Support

DB2 is now both an OLE DB provider and an OLE DB consumer. This support gives customers with OLE DB-based applications the ability to extract or query DB2 data using this native OLE interface. In addition, you can load data into DB2 or simply access OLE DB data using OLE DB table functions. OLE DB table functions combined with DB2's distributed query capability enable distributed query on OLE DB, DB2, and other data sources.

SQL Assist

With SQL Assist and some knowledge of SQL, you can create SELECT, INSERT, UPDATE, and DELETE statements. SQL Assist is an online tool that uses a notebook to help you organize the information that you need to create an SQL statement. SQL Assist is available in the Control Center, the Stored Procedure Builder, and the Data Warehouse Center. See the *Online Help* for more information.

Closing Log After Backup

After an online backup is complete, DB2 forces the currently active log to be closed and as a result it will be archived off. This ensures that your online backup has a complete set of archived logs available for recovery. See the *Administration Guide: Implementation* for more information.

On Demand Log Archiving

DB2 Version 7.2 allows you to close and, if the user exit option is enabled, archive the active log of a recoverable database whenever you want. With this new DB2 API, you can collect a complete set of log files up to a point in time when you initiate the **db2ArchiveLog** API, then use these log files to update a standby database.

Note: You can also issue the **ARCHIVE LOG** command from the Command Line Processor to force a log archive.

Identity Column Support

DB2 now allows you to generate a unique numeric value for each row in a table. For example, an identity column can be used to generate unique primary key values. You can also port applications that use identity columns from other RDBMS vendors to DB2. See the *Administration Guide: Implementation* for more information.

DB2 Version 7.2 enhances identity column support in its use with load. After loading data into a table with identity columns, you can use the ALTER TABLE statement to restart the sequence of values assigned to the identity column at the correct value. See ALTER TABLE in the *DB2 Release Notes* for the new syntax of the ALTER TABLE statement that supports this enhancement.

Larger Index Key

DB2 Version 7.2 allows columns with a length greater than 255 bytes to be specified as part of an index key. A larger index key has a positive effect on the usefulness of an index. The DB2_INDEX_2BYTEVARLEN registry variable allows for the use of two bytes instead of one to store the length of an index key. To convert existing indexes to use larger index keys: drop the indexes, set the DB2_INDEX_2BYTEVARLEN registry variable to ON, then recreate the indexes (using the larger columns).

For more information on larger index keys, see the *DB2 Release Notes*.

Externalized Savepoint Support

You can now programmatically roll back to a specified savepoint rather than rolling back an entire unit of work. This support gives application developers more control over rollback processing. See the *Administration Guide: Implementation* and the *SQL Reference, Volume 2* for more information.

Increased Log Limit to 32GB

The size of log files has been increased from 4GB to 32GB. This enables you to perform extremely large amounts of work within a single transaction. See the *Administration Guide: Implementation* for more information.

Block Transactions When Log Directory Is Full

DB2 Version 7.2 provides `DB2_BLOCK_ON_LOG_DISK_FULL`, a new registry variable that you can set to prevent "disk full" errors from being generated when DB2 cannot create a new log file in the active log path. DB2 attempts to create the log file every five minutes and writes a message to the `db2diag.log` file after each attempt. You can use the `db2diag.log` file to confirm if your application is hanging.

For more information on this registry variable, see the *DB2 Release Notes*.

Rename a Table Space

When you create a table space you specify its name. In previous releases, there was no way of changing that name during the lifetime of the table space object. Now, you can rename a table space with the `RENAME TABLESPACE` statement. See the *Administration Guide: Implementation* for more information.

Database Buffer Pool

The maximum size of database buffer pools has been increased in Version 7.2 to exploit the capabilities of increased memory and address space. Larger buffer pools mean DB2 can put more data in memory, making data access and sorting faster.

Support for the increase in the maximum size of database buffer pools is enabled through:

- Windows 2000 Address Windowing Extensions (AWE)
- 64-Bit Support on AIX, HP-UX, and the Solaris Operating Environment.

Windows 2000 Address Windowing Extensions (AWE)

DB2 Version 7.2 on Windows 2000 can now allocate buffer pools that use up to 64 GB of memory through the Microsoft Windows Address Windowing

Extensions (AWE) API, which is supported on all versions of Windows 2000. With AWE support, the total size of DB2 buffer pools can reach up to 64 GB.

While all versions of Windows 2000 support this functionality, only the Windows 2000 Advanced Server and Datacenter Server provide the increased memory support.

- Windows 2000 Advanced Server provides support for up to 8 GB of memory.
- Windows 2000 Datacenter Server provides support for up to 64 GB of memory.

To support AWE buffer pools, Windows 2000 must be configured correctly. This includes having the DB2_AWE registry variable set, associating the "lock pages in memory" privilege with the user under which DB2 is installed, and allocating the physical pages and the address window pages. See the *DB2 Release Notes* for more information.

64-Bit Support on AIX, HP-UX, and the Solaris Operating Environment

DB2 Version 7.2 offers 64-bit support on AIX, HP-UX, and the Solaris Operating Environment to exploit the increased memory address space, and provides application libraries that can access DB2 databases from 64-bit applications.

Applications running on 64-bit operating systems benefit from the increased memory address space that these systems provide. For DB2 Universal Database, the 64-bit address space allows for the creation of larger buffer pools, sort heap, package caches, and other resources that require large amounts of memory. Additional memory improves the performance of many operations, especially sort and I/O operations.

For additional information, see the manual entitled *Using DB2 Universal Database on 64-bit Platforms*. This manual is available in PDF format on the Publications CD-ROM with the DB2 Manuals.

Increased Locklist Limit

With DB2 Version 7.2, the maximum value allocated to the *locklist* configuration parameter has increased. This has a positive impact on database performance and the query optimizer. See the *DB2 Release Notes* for more information.

Unicode Enhancements

DB2 Version 7.2 provides a number of enhancements, increasing DB2 compliance with the Unicode standard, as follows:

- Casting between data types has been extended for Unicode databases

- A new set of functions as well as new C and SQL data types have been provided for the DB2 CLI Unicode applications
- New CLI configuration keywords have been added to avoid extra overhead when Unicode applications connect to a database.

For more information on these enhancements, see the *Unicode Updates* section in the *DB2 Release Notes*. For more information on using Unicode with DB2, see the *Administration Guide: Planning*.

Encrypting and Decrypting String Data

Securing and controlling access to your database is important to the reliability and integrity of the data. With the built-in encryption and decryption functions provided by DB2 Version 7.2, you can encrypt data to add an additional layer of security.

- The ENCRYPT function encrypts data using a password-based encryption method. The encryption function also allows for a password hint to be stored and another function is provided to get the hint without using the password.
- The DECRYPT_BIN and DECRYPT_CHAR functions decrypt data using a password-based decryption method.
- The GETHINT function returns an encapsulated password hint that a data owner defines as a hint.

For more information on data encryption, see the *DB2 Release Notes*.

Database Recovery

As databases increase in size and queries place demands on the continued availability of the operating system, the time and hardware resources required to back up and recover databases are also growing substantially. Backing up an entire database or the table spaces of large databases can strain your system resources because the storage requirements to maintain multiple copies of a database or of table spaces are considerable.

DB2 Version 7.2 provides a number of new functions that will help you perform online or offline backups, while allowing your data to be accessed as needed. See the *DB2 Release Notes* for more information on each of the following new functions.

Suspended I/O

Many storage vendors now provide storage solutions that enhance the availability of data. One such offering is the ability to split a mirrored copy of

data and make that mirrored copy available for processing or available to another server. To exploit this storage capability, DB2 has delivered two new features in Version 7.2:

- **Suspended I/O**

Suspended I/O supports continuous system availability, while providing for online split mirror handling of the database. By momentarily suspending I/O to disk, DB2 will ensure that the split mirror copy maintains its integrity

- **db2inidb utility**

The **db2inidb** utility operates on the mirrored copy and can do the following:

- Perform crash recovery, which gives you a duplicate copy of the database for reporting purposes
- Put the copied database in a roll-forward pending state and roll forward the database, which keeps the mirrored copy synchronized with the primary database
- Allow the mirrored copy of the database to be backed up, thus having a backup that does not impact the performance of the primary database server. See “Backup from Split Image” on page 35 for additional information.

Dual Logging

Active logs help database administrators perform database recovery. DB2 now has the capability to mirror the active log files to protect databases from:

- Accidental deletion of an active log
- Data corruption caused by hardware failure

A new registry variable, `DB2_NEWLOGPATH2`, allows the database to write an identical copy of the log files to a different path on a physically separate disk.

Parallel Recovery

DB2 now uses multiple agents to perform both crash recovery and database rollforward recovery. You can expect better performance during these operations, particularly on symmetric multi-processor (SMP) machines. Using multiple agents during database recovery takes advantage of the extra CPUs that are available on SMP machines.

Note: You cannot use multiple agents to perform table space rollforward recovery.

See the *DB2 Release Notes* for more information on parallel recovery.

Backup from Split Image

DB2 Version 7.2 allows you to perform a backup from a split image, following the use of the **db2inidb** utility.

- When you use the **snapshot** option, the database will perform crash recovery, start a new log chain and will not be able to rollforward through any of the logs from the original database. The database is available for any operation, including backup.
- When you use the **standby** or **mirror** option, crash recovery is not performed. The database remains in an inconsistent state, with inflight transactions still outstanding. The database is placed in rollforward recovery.

The backup from split image feature allows you to perform an offline backup of the entire database.

Note: Online backup is not supported and is not necessary because the database is in rollforward pending state and is unavailable.

For DB2 Version 7.2, this support is limited to databases that contain DMS table spaces only.

Named Pipe Support for Backup and Restore Operations

You can use named pipes for backing up and restoring databases on UNIX-based systems. To use this function, the writer and reader of the named pipe must be on the same machine. You must create the named pipe on a local file system before backing up the database.

Incremental and Delta Backups

DB2 support for incremental backup responds to the need to manage system resources while database backup and restore operations occur. Two types of incremental backup are supported:

- Incremental backup
A backup image of all database data that contains only pages that have been updated since the previous backup of the database or table space was taken. The image also contains the initial database metadata that is normally stored in full backup images, such as the database configuration, table space definitions and database history. Each successive backup contains the entire contents of the previous incremental image, plus any data that is new or changed since the previous incremental backup.

Note: All LOB and LF data types that have been modified in any way since the preceding full, incremental or delta backup are copied to the backup image.

- Delta backup

A delta image of all data that has changed since the last successful backup of a table space. The last backup could have been a full, incremental, or delta backup. Each successive delta backup also contains the database metadata that is contained in full backup images.

You can use either incremental or delta backup in either online or offline backup operations. Incremental backup can help you by:

- Reducing the size of backup images
- Improving the speed of backup if:
 - The table space is striped across multiple containers
 - You are backing up to slow media, such as a small number of tape devices.

DB2 tracks the database updates with the new database configuration parameter, *trackmod*. See the *DB2 Release Notes* for more information on incremental and delta backups.

Sequence Support

DB2 Version 7.2 provides a standalone, efficient, and recoverable (automatic) number generator through a new type of database object called a SEQUENCE. A sequence object enables the database manager to automatically generate a new numeric value for each call made to the sequence's NEXTVAL expression.

Applications can use sequences to avoid possible concurrency and performance problems that can result from generating a unique counter outside the database.

Note: Sequence data types are not yet supported in DB2 Enterprise - Extended Edition.

See the *DB2 Release Notes* for more information on sequence support.

Statement-Level Isolation Levels

Before DB2 Version 7.2, you could only specify the isolation level at the package level when preparing or binding an application. With Version 7.2, you can now define isolation levels at the statement level for better granularity and improved performance and concurrency. See the *DB2 Release Notes* for more detailed information on statement-level isolation levels.

New SQL Built-in Scalar Functions

Three new SQL built-in scalar functions are now available with the SYSIBM schema for processing columns of data defined with a decimal data type:

- ABS or ABSVAL

This function returns the absolute value of the argument. The result of the function has the same data type and length attribute as the argument.

- MULTIPLY_ALT

This function returns the product of two arguments as a decimal value. It is especially beneficial when the sum of the precisions of the arguments exceeds 31 (that is, more than 31 decimal places are required for an accurate result).

- ROUND

This function returns *expression1* rounded to *expression2*. The result of the function is to round the value of *expression1* to the next higher positive number if *expression1* is positive, or to the next lower negative number if *expression1* is negative.

See the *DB2 Release Notes* for more information on these new built-in scalar functions.

Parallel Container Creation

DB2 Version 7.2 allows you to create table space containers in parallel. This enhancement helps to improve the performance of I/O operations if you are using database-managed space (DMS) device containers for your table spaces.

Starting with Version 7.2, you can create or resize several DMS containers in parallel, with the degree of parallelism equal to the number of available prefetchers plus one. For example, if you have 10 prefetchers, you could resize 11 containers at the same time.

Note: Prefetchers retrieve data from disks and move the data into the database buffer pools before applications need the data.

For more information on creating or resizing containers in parallel, see the *DB2 Release Notes*. For more information on managing device containers, see the *Administration Guide: Performance*.

Chapter 5. DB2 Family Enhancements

It's a heterogeneous world and you need a database server that can integrate fully with a multi-platform environment, so your enterprise can enjoy the full benefits of data warehousing, business intelligence, and e-business.

Mobile Computing

DB2's mobile computing offering is designed to give mobile workers using portable devices access to corporate data and applications. This gives these workers the power of the enterprise from any location at any time.

DB2 Everyplace

The DB2 Everyplace Sync Server extends the power of your DB2 enterprise to handheld devices. Designed for Personal Digital Assistants (PDAs), Handheld Personal Computers (HPCs), and soon to be enabled for cellular phones, DB2 Universal Database Everyplace makes your DB2 data completely mobile. DB2 Everyplace is a slimmed down database that empowers handheld computing devices with enterprise data. With the DB2 Everyplace Sync Server, it mobilizes professionals with e-business information anywhere, anytime. It extends the power of your enterprise DB2 to a wide range of handheld devices, such as those running the Palm Operating System and Windows CE.

>8 Character UserID Support

The userID length supported by DB2 Universal Database has been increased from 8 characters to 30 characters for some operating systems. The following list shows the Version 7 level of support for:

- userID
All DB2 Universal Database Version 7 servers running Windows 32-bit operating systems support userIDs of up to 30 characters. All DB2 Universal Database Version 7 clients support userIDs of up to 30 characters. For example, any Version 7 client can connect to a Version 7 Windows 32-bit server with a 30 byte userID. But a Version 7 client with a 30 byte userID cannot connect to a non-Windows 32 bit server.
- authID
All DB2 Universal Database Version 7 Servers support authIDs of up to 30 characters.
- schema
All DB2 Universal Database Version 7 Servers support schema names of up to 30 characters.

Note: Many versions of the Windows operating system limit userIDs to 20 characters

See the *Quick Beginnings* for your operating system for more information. You should also review the following sections.

Pre-Version 7 Servers

Pre-Version 7 servers do not support userIDs, authIDs, and schemas that are longer than 8 characters. Version 7 applications that are designed to exploit the >8 character support will fail when connecting with down-level DB2 servers.

Replication

The length of userID and schema names supported by replication has been increased from 8 characters to 18 characters. See the *Replication Guide and Reference* for more information.

DB2 Universal Database for OS/390

DB2 Universal Database for OS/390 supports 8-character userIDs and authIDs. Similar length restrictions apply to the schema name in any statements being sent to the host.

DB2 Universal Database for AS/400

DB2 Universal Database for AS/400 supports 10-character userIDs and authIDs. Similar length restrictions apply to the schema name in any statements being sent to the host.

Import and Export

Version 7 databases with a schema name greater than 8 characters cannot be imported or exported with pre-Version 7 code because truncation occurs.

Stored Procedures

Existing stored procedures that assume that userIDs, authIDs, and schema names are limited to 8 characters should be examined. Unexpected behavior might occur if these stored procedures are used in a Version 7 environment.

LOAD Authority

The LOAD authority, in previous releases only available with DB2 Universal Database for OS/390, is now available for the DB2 Universal Database family. Users granted LOAD authority can run the LOAD utility without the need for SYSADM or DBADM authority. This allows users to perform more DB2 functions, and gives database administrators more granular control over the administration of their database. See the *Data Movement Utilities Guide and Reference* for more information.

USE OF TABLESPACE Authority

The USE OF TABLESPACE authority, only available in previous releases with DB2 Universal Database for OS/390, is now available for the DB2 Universal Database family. USE OF TABLESPACE authority allows users to create tables only in those table spaces that they have been granted access to. It also provides the administrators with more control over their database. See the *Administration Guide: Implementation* for more information.

BIND Options

The SQLERROR(CONTINUE) and VALIDATE(RUN) bind options, only available in previous releases with DB2 Universal Database for OS/390, are now available for the DB2 Universal Database family. This allows you to port your DB2 Universal Database for OS/390 applications that make use of these options to the rest of the DB2 Universal Database family. See the *Application Development Guide* for more information.

OS/390 in the Control Center

New DB2 Universal Database for OS/390 functions have been incorporated into the Control Center.

Generate DDL

The Generate DDL function allows the user to selectively generate the DDL statements used to create database objects and optionally, dependent objects. As an example, a user can request that the DDL be generated for a table, including all its indexes and all views based on the table. The generated statements can then be saved either on OS/390 as a data set or in a local workstation file. When running Control Center as an applet, the workstation file resides on the Web server. See the *Control Center Online Help* for more information.

OS/390 Utilities

The Control Center gives you the ability to restart OS/390 utilities that have stopped. Any utilities whose status is active or stopped can be displayed using the display utility command. Utilities whose status is stopped can be restarted. There are two ways to restart utilities: from the last committed point (Current) or the last committed phase (Phase). Stopped utilities that were not started from Command Center cannot be restarted.

Each time a utility is run, a Utility ID is generated. In Version 7, you can modify the ID generated by the Control Center, and replace it with a default Utility ID that the user can specify on the Tools Settings window.

After a utility has been run, you may need to delete or rename the data sets that remain. You can use the Command Center's new OS/390 DATASET object to manage these data sets.

You can use wildcards and dynamic allocation to develop DB2 for OS/390 utility control statements to manage a large number of database objects. Utilities have the ability to dynamically build lists of database objects and dynamically allocate the data sets that they use or produce. LIST, TEMPLATE, and JOB STEP objects have been added to the Control Center in support of DB2 for OS/390.

See the *Control Center Online Help* for more information on OS/390 utility support.

Windows 32-bit Operating Systems

The following enhancements have been made in Version 7 of DB2 Universal Database for Windows.

Windows 2000

DB2 Universal Database supports the following Windows 2000 features:

- DB2 services are published in the Active Directory containing the protocol configuration information that allows client applications to connect to the DB2 database server.
- DB2 authentication can be accomplished using the single sign-on facilities of Kerberos. DB2 authentication is unchanged for environments without Kerberos.
- The DB2 Universal Database Control Center can be launched from Microsoft Management Console (MMC).

Visual C++ Application Development

DB2 Universal Database for Windows 32-bit operating systems includes two components to ease your DB2 application development efforts with Microsoft Visual C++ using embedded SQL. The DB2 Visual C++ Add-Ins provide you with graphical user interfaces that are consistent with the Visual C++ development environment.

- The DB2 Visual C++ Tools Add-In is a toolbar that launches some useful DB2 administration and development tools from within the Visual C++ integrated development environment (IDE).
- The DB2 Visual C++ Project Add-In provides management tools and wizards that plug into the Visual C++ IDE that help you to develop, package, and deploy client applications and stored procedures for DB2 servers on Windows 32-bit operating systems.

Visual Studio Integration

DB2 Universal Database provides a collection of tools and wizards to simplify the building and deploying of applications for DB2 Universal Database for Windows. They employ embedded SQL from within the Visual C++ Integrated Development Environment (IDE).

Administering DB2 UDB Workgroup Edition and DB2 UDB Personal Edition as Satellites

As of DB2 Version 7.2, you can administer any DB2 Universal Database Workgroup Edition or DB2 Universal Database Personal Edition system that runs on a Windows-based platform as a satellite.

Note: To prevent a FixPak from having to be applied on the Version 6 or Version 7.1 Enterprise Edition server that is being used as the DB2 control server, the Version 7.2 Workgroup Edition and Personal Edition satellites will report themselves as being Version 6 satellites at the DB2 control server. This does not affect the functionality of the Version 7.2 DB2 system.

For details on how to set up DB2 UDB Workgroup Edition or DB2 UDB Personal Edition as a satellite, see the *DB2 Release Notes*. For information on setting up and maintaining a satellite environment, see the *Administering Satellites Guide and Reference*.

Running Compiled SQL Procedures

You can run compiled SQL procedures on DB2 servers on the same operating system without compiling the procedure on each destination server with DB2 Version 7.2. This enhancement to DB2 is cost-effective, eliminates the need to install a compiler on each server, and in the long run, saves time.

For more information on distributing compiled SQL procedures, see the *DB2 Release Notes*.

Cross-Platform Backup and Restore

With DB2 Version 7.2, you can back up and restore databases between HP-UX and the Solaris Operating Environment. This is beneficial to customers who maintain databases on HP-UX and the Solaris Operating Environment. For more information on this enhancement, see the *DB2 Release Notes*. For more information on database backup and restore, see “Database Recovery” on page 33 and the *Administration Guide: Implementation*.

DB2 Data Links Manager for the Solaris Operating Environment

Before Version 7, DB2 Data Links Manager was only available on the Windows NT and AIX operating systems. DB2 Version 7.2 extends the availability of the DB2 Data Links Manager to the Solaris Operating Environment. You can now take advantage of the DB2 Data Links technology that provides referential integrity, access control, and recovery capabilities for files that physically reside on file systems that are external to DB2 Universal Database.

DB2 Data Links Manager for the Solaris Operating Environment runs in 32-bit mode. For more information on the DB2 Data Links Manager, see the *DB2 Data Links Manager Quick Beginnings*.

TSM Integration with DB2 Data Links Manager

DB2 Data Links Manager will now be able to use the functionality of Tivoli Space Manager (TSM) and its virtual filesystem, called FSM, which layers itself on top of native filesystems like JFS. FSM can be accessed and configured in the same manner as JFS.

This new feature benefits customers who have filesystems with large files that are required to be moved to tertiary storage periodically, and the space of the filesystem needs to be managed on a regular basis. For many customers, TSM currently provides the means to manage their tertiary storage. The new DB2 Data Links Manager support of TSM provides greater flexibility in managing the space for DATALINK files. Rather than pre-allocating enough storage in the DB2 Data Links Manager file system for all files which may be stored there, TSM allows allocations of the Data Links managed file system to be adjusted over a period of time without the risk of inadvertently filling up the file system during normal usage.

Chapter 6. DB2 Connect Enhancements

DB2 Connect enables easier access to enterprise data stored in relational databases on MVS, OS/400, OS/390, VM and VSE hosts, and in non-relational databases such as IMS. DB2 Connect products share technology with the DB2 UDB products and, as a result, deliver many of the new features and functions provided by the DB2 UDB Version 7.2.

This section provides a brief summary of the features that are new to DB2 Connect products and refers to more detailed descriptions to other places in this book:

- Enhanced access to non-DB2 data delivered by DB2 Relational Connect support for Microsoft SQL Server and Sybase databases as well as support for new platforms. For additional details, see “Access to New Database Management Systems” on page 21.
- Exploitation of 64-bit architectures on AIX, HP-UX, and the Solaris Operating Environment platforms (see “64-Bit Support on AIX, HP-UX, and the Solaris Operating Environment” on page 32).
- Static SQL Profiling support for converting ODBC, JDBC and SQLJ calls to Static SQL (see “Convert ODBC (Dynamic SQL) to Static SQL” on page 23).
- Enhanced support and integration with Windows 32-bit operating systems (see “Windows 32-bit Operating Systems” on page 42).
- Native OLE DB provider for improved support of application development using Microsoft technologies (see “Native OLE DB Support” on page 29).
- Support for UNICODE in the ODBC and DB2 CLI drivers (see “Unicode Enhancements” on page 32).
- Control Center enhancements for improved management and operation of the DB2 for OS/390 database servers (see “OS/390 in the Control Center” on page 41).

DB2 Connect Managed Access

DB2 Connect-managed access to database servers that implement Distributed Relational Database Architecture (DRDA) provides the following enhancements for Version 7.

Improved Support for Microsoft Transaction Server (MTS) and COM+ Technologies

DB2 Connect Version 7.2 implements enhancements to the distributed transaction processing. This improved support greatly reduces the possibility of deadlocks that can occur when multiple COM+ components, participating

in the same global transaction, attempt to access the same DB2 for OS/390 data. The support is implemented in conjunction with the lock space sharing support for DB2 for OS/3900 version 6.1 database servers delivered through the following PTFs:

- PQ39416
- PQ28487
- PQ27022
- PQ32387

As a result of these enhancements, DB2 for OS/390 V6.1 database servers are now able to recognize multiple COM+ components participating in a transaction and to share the lock space between these COM+ objects. These enhancements ensure that resource requirements of one object do not prevent another object from acquiring the resources it needs. A limitation still exists where lock space cannot be shared between multiple members in a data sharing group in a Sysplex environment.

Multisite Update Improvements

The *test connection* feature of the Configure Multisite Update wizard has been expanded to permit testing of remote instances. In addition, you can associate a different user and password with each database entry in the test list. See the *DB2 Connect User's Guide* for more information.

DB2 Connect Web Starter Kit

DB2 Connect Web Starter Kit provides an economical way for customers to evaluate solutions and develop applications that use the latest web technologies with DB2 for OS/390 and DB2 Server for VM and VSE databases. DB2 Connect Web Starter Kit contains all of the DB2 Connect Enterprise Edition servers and DB2 Connect Personal Edition products, and allows unlimited use of these products for a nine month license period.

This product should be considered for projects that require a proof of concept or a pilot implementation of a web-based application, and may not be ready to justify licensing of the DB2 Connect Unlimited Edition product. The nine month license period provides unlimited access to DB2 Connect technologies, and gives ample time for evaluation and proof of concept projects. At the end of the nine month license period, you can do one of the following:

- Stop using the product completely.
- Upgrade to DB2 Connect Unlimited Edition using the MSU entitlement purchased.
- Purchase DB2 Connect Enterprise Edition using Registered Users licensing.

DB2 Connect for Linux

DB2 Connect allows you to access enterprise data stored on mainframes and midrange systems from Windows, OS/2, and UNIX desktop systems. In Version 7.2, DB2 Connect for Linux is available in the following editions:

- DB2 Connect Unlimited Edition for Linux/390
- DB2 Connect Enterprise Edition for Linux on Intel processors
- DB2 Connect Personal Edition for Linux on Intel processors

Support DCL SNA Product

Users of the Solaris Operating Environment have a new way to communicate with DB2 databases on the Solaris Operating Environment, or on host and AS/400 servers. With Version 7, DB2 Connect now supports DCL SNAP-IX V6.1.0 for SPARC Solaris. Previously, DB2 Connect only supported the SUNLINK SNA product. This new DCL SNA product is very similar to CS/AIX V5 in its implementation.

For SUNLINK, outbound connections using CPIC APIs and inbound connections using APPC APIs are now supported. For DCL, in addition to the above, SPM high level description is also supported. See the *DB2 Connect Enterprise Edition for UNIX Quick Beginnings* for more information.

Appendix A. Using the DB2 Library

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

To access product information online, you can use the Information Center. For more information, see “Accessing Information with the Information Center” on page 63. You can view task information, DB2 books, troubleshooting information, sample programs, and DB2 information on the Web.

DB2 PDF Files and Printed Books

DB2 Information

The following table divides the DB2 books into four categories:

DB2 Guide and Reference Information

These books contain the common DB2 information for all platforms.

DB2 Installation and Configuration Information

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

Cross-platform sample programs in HTML

These samples are the HTML version of the sample programs that are installed with the Application Development Client. The samples are for informational purposes and do not replace the actual programs.

Release notes

These files contain late-breaking information that could not be included in the DB2 books.

The installation manuals, release notes, and tutorials are viewable in HTML directly from the product CD-ROM. Most books are available in HTML on the product CD-ROM for viewing and in Adobe Acrobat (PDF) format on the DB2 publications CD-ROM for viewing and printing. You can also order a printed copy from IBM; see “Ordering the Printed Books” on page 59. The following table lists books that can be ordered.

On OS/2 and Windows platforms, you can install the HTML files under the `sql1ib\doc\html` directory. DB2 information is translated into different

languages; however, all the information is not translated into every language. Whenever information is not available in a specific language, the English information is provided

On UNIX platforms, you can install multiple language versions of the HTML files under the `doc/%L/html` directories, where `%L` represents the locale. For more information, refer to the appropriate *Quick Beginnings* book.

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

- “Viewing Information Online” on page 62
- “Searching Information Online” on page 66
- “Ordering the Printed Books” on page 59
- “Printing the PDF Books” on page 58

Table 1. DB2 Information

Name	Description	Form Number PDF File Name	HTML Directory
DB2 Guide and Reference Information			
<i>Administration Guide</i>	<i>Administration Guide: Planning</i> provides an overview of database concepts, information about design issues (such as logical and physical database design), and a discussion of high availability.	SC09-2946 db2d1x70	db2d0
	<i>Administration Guide: Implementation</i> provides information on implementation issues such as implementing your design, accessing databases, auditing, backup and recovery.	SC09-2944 db2d2x70	
	<i>Administration Guide: Performance</i> provides information on database environment and application performance evaluation and tuning.	SC09-2945 db2d3x70	
	You can order the three volumes of the <i>Administration Guide</i> in the English language in North America using the form number SBOF-8934.		
<i>Administrative API Reference</i>	Describes the DB2 application programming interfaces (APIs) and data structures that you can use to manage your databases. This book also explains how to call APIs from your applications.	SC09-2947 db2b0x70	db2b0

Table 1. DB2 Information (continued)

Name	Description	Form Number PDF File Name	HTML Directory
<i>Application Building Guide</i>	Provides environment setup information and step-by-step instructions about how to compile, link, and run DB2 applications on Windows, OS/2, and UNIX-based platforms.	SC09-2948 db2axx70	db2ax
<i>APPC, CPI-C, and SNA Sense Codes</i>	Provides general information about APPC, CPI-C, and SNA sense codes that you may encounter when using DB2 Universal Database products.	No form number db2apx70	db2ap
	Available in HTML format only.		
<i>Application Development Guide</i>	Explains how to develop applications that access DB2 databases using embedded SQL or Java (JDBC and SQLJ). Discussion topics include writing stored procedures, writing user-defined functions, creating user-defined types, using triggers, and developing applications in partitioned environments or with federated systems.	SC09-2949 db2a0x70	db2a0
<i>CLI Guide and Reference</i>	Explains how to develop applications that access DB2 databases using the DB2 Call Level Interface, a callable SQL interface that is compatible with the Microsoft ODBC specification.	SC09-2950 db2l0x70	db2l0
<i>Command Reference</i>	Explains how to use the Command Line Processor and describes the DB2 commands that you can use to manage your database.	SC09-2951 db2n0x70	db2n0
<i>Connectivity Supplement</i>	Provides setup and reference information on how to use DB2 for AS/400, DB2 for OS/390, DB2 for MVS, or DB2 for VM as DRDA application requesters with DB2 Universal Database servers. This book also details how to use DRDA application servers with DB2 Connect application requesters.	No form number db2h1x70	db2h1
	Available in HTML and PDF only.		

Table 1. DB2 Information (continued)

Name	Description	Form Number PDF File Name	HTML Directory
<i>Data Movement Utilities Guide and Reference</i>	Explains how to use DB2 utilities, such as import, export, load, AutoLoader, and DPROF, that facilitate the movement of data.	SC09-2955 db2dmx70	db2dm
<i>Data Warehouse Center Administration Guide</i>	Provides information on how to build and maintain a data warehouse using the Data Warehouse Center.	SC26-9993 db2ddx70	db2dd
<i>Data Warehouse Center Application Integration Guide</i>	Provides information to help programmers integrate applications with the Data Warehouse Center and with the Information Catalog Manager.	SC26-9994 db2adx70	db2ad
<i>DB2 Connect User's Guide</i>	Provides concepts, programming, and general usage information for the DB2 Connect products.	SC09-2954 db2c0x70	db2c0
<i>DB2 Query Patroller Administration Guide</i>	Provides an operational overview of the DB2 Query Patroller system, specific operational and administrative information, and task information for the administrative graphical user interface utilities.	SC09-2958 db2dwx70	db2dw
<i>DB2 Query Patroller User's Guide</i>	Describes how to use the tools and functions of the DB2 Query Patroller.	SC09-2960 db2wwx70	db2ww
<i>Glossary</i>	Provides definitions for terms used in DB2 and its components. Available in HTML format and in the <i>SQL Reference</i> .	No form number db2t0x70	db2t0
<i>Image, Audio, and Video Extenders Administration and Programming</i>	Provides general information about DB2 extenders, and information on the administration and configuration of the image, audio, and video (IAV) extenders and on programming using the IAV extenders. It includes reference information, diagnostic information (with messages), and samples.	SC26-9929 dmbu7x70	dmbu7
<i>Information Catalog Manager Administration Guide</i>	Provides guidance on managing information catalogs.	SC26-9995 db2dix70	db2di

Table 1. DB2 Information (continued)

Name	Description	Form Number PDF File Name	HTML Directory
<i>Information Catalog Manager Programming Guide and Reference</i>	Provides definitions for the architected interfaces for the Information Catalog Manager.	SC26-9997 db2bix70	db2bi
<i>Information Catalog Manager User's Guide</i>	Provides information on using the Information Catalog Manager user interface.	SC26-9996 db2aix70	db2ai
<i>Installation and Configuration Supplement</i>	Guides you through the planning, installation, and setup of platform-specific DB2 clients. This supplement also contains information on binding, setting up client and server communications, DB2 GUI tools, DRDA AS, distributed installation, the configuration of distributed requests, and accessing heterogeneous data sources.	GC09-2957 db2iyx70	db2iy
<i>Message Reference</i>	Lists messages and codes issued by DB2, the Information Catalog Manager, and the Data Warehouse Center, and describes the actions you should take. You can order both volumes of the Message Reference in the English language in North America with the form number SBOF-8932.	Volume 1 SC09-2978 db2m1x70 Volume 2 SC09-2979 db2m2x70	db2m0
<i>OLAP Integration Server Administration Guide</i>	Explains how to use the Administration Manager component of the OLAP Integration Server.	SC27-0782 db2dpx70	n/a
<i>OLAP Integration Server Metaoutline User's Guide</i>	Explains how to create and populate OLAP metaoutlines using the standard OLAP Metaoutline interface (not by using the Metaoutline Assistant).	SC27-0784 db2upx70	n/a
<i>OLAP Integration Server Model User's Guide</i>	Explains how to create OLAP models using the standard OLAP Model Interface (not by using the Model Assistant).	SC27-0783 db2lpx70	n/a
<i>OLAP Setup and User's Guide</i>	Provides configuration and setup information for the OLAP Starter Kit.	SC27-0702 db2ipx70	db2ip
<i>OLAP Spreadsheet Add-in User's Guide for Excel</i>	Describes how to use the Excel spreadsheet program to analyze OLAP data.	SC27-0786 db2epx70	db2ep

Table 1. DB2 Information (continued)

Name	Description	Form Number PDF File Name	HTML Directory
<i>OLAP Spreadsheet Add-in User's Guide for Lotus 1-2-3</i>	Describes how to use the Lotus 1-2-3 spreadsheet program to analyze OLAP data.	SC27-0785 db2tpx70	db2tp
<i>Replication Guide and Reference</i>	Provides planning, configuration, administration, and usage information for the IBM Replication tools supplied with DB2.	SC26-9920 db2e0x70	db2e0
<i>Spatial Extender User's Guide and Reference</i>	Provides information about installing, configuring, administering, programming, and troubleshooting the Spatial Extender. Also provides significant descriptions of spatial data concepts and provides reference information (messages and SQL) specific to the Spatial Extender.	SC27-0701 db2sbx70	db2sb
<i>SQL Getting Started</i>	Introduces SQL concepts and provides examples for many constructs and tasks.	SC09-2973 db2y0x70	db2y0
<i>SQL Reference, Volume 1 and Volume 2</i>	Describes SQL syntax, semantics, and the rules of the language. This book also includes information about release-to-release incompatibilities, product limits, and catalog views. You can order both volumes of the <i>SQL Reference</i> in the English language in North America with the form number SBOF-8933.	Volume 1 SC09-2974 db2s1x70 Volume 2 SC09-2975 db2s2x70	db2s0
<i>System Monitor Guide and Reference</i>	Describes how to collect different kinds of information about databases and the database manager. This book explains how to use the information to understand database activity, improve performance, and determine the cause of problems.	SC09-2956 db2f0x70	db2f0
<i>Text Extender Administration and Programming</i>	Provides general information about DB2 extenders and information on the administration and configuring of the text extender and on programming using the text extenders. It includes reference information, diagnostic information (with messages) and samples.	SC26-9930 desu9x70	desu9

Table 1. DB2 Information (continued)

Name	Description	Form Number PDF File Name	HTML Directory
<i>Troubleshooting Guide</i>	Helps you determine the source of errors, recover from problems, and use diagnostic tools in consultation with DB2 Customer Service.	GC09-2850 db2p0x70	db2p0
<i>What's New</i>	Describes the new features, functions, and enhancements in DB2 Universal Database, Version 7.	SC09-2976 db2q0x70	db2q0
DB2 Installation and Configuration Information			
<i>DB2 Connect Enterprise Edition for OS/2 and Windows Quick Beginnings</i>	Provides planning, migration, installation, and configuration information for DB2 Connect Enterprise Edition on the OS/2 and Windows 32-bit operating systems. This book also contains installation and setup information for many supported clients.	GC09-2953 db2c6x70	db2c6
<i>DB2 Connect Enterprise Edition for UNIX Quick Beginnings</i>	Provides planning, migration, installation, configuration, and task information for DB2 Connect Enterprise Edition on UNIX-based platforms. This book also contains installation and setup information for many supported clients.	GC09-2952 db2cyx70	db2cy
<i>DB2 Connect Personal Edition Quick Beginnings</i>	Provides planning, migration, installation, configuration, and task information for DB2 Connect Personal Edition on the OS/2 and Windows 32-bit operating systems. This book also contains installation and setup information for all supported clients.	GC09-2967 db2c1x70	db2c1
<i>DB2 Connect Personal Edition Quick Beginnings for Linux</i>	Provides planning, installation, migration, and configuration information for DB2 Connect Personal Edition on all supported Linux distributions.	GC09-2962 db2c4x70	db2c4
<i>DB2 Data Links Manager Quick Beginnings</i>	Provides planning, installation, configuration, and task information for DB2 Data Links Manager for AIX and Windows 32-bit operating systems.	GC09-2966 db2z6x70	db2z6

Table 1. DB2 Information (continued)

Name	Description	Form Number PDF File Name	HTML Directory
<i>DB2 Enterprise - Extended Edition for UNIX Quick Beginnings</i>	Provides planning, installation, and configuration information for DB2 Enterprise - Extended Edition on UNIX-based platforms. This book also contains installation and setup information for many supported clients.	GC09-2964 db2v3x70	db2v3
<i>DB2 Enterprise - Extended Edition for Windows Quick Beginnings</i>	Provides planning, installation, and configuration information for DB2 Enterprise - Extended Edition for Windows 32-bit operating systems. This book also contains installation and setup information for many supported clients.	GC09-2963 db2v6x70	db2v6
<i>DB2 for OS/2 Quick Beginnings</i>	Provides planning, installation, migration, and configuration information for DB2 Universal Database on the OS/2 operating system. This book also contains installation and setup information for many supported clients.	GC09-2968 db2i2x70	db2i2
<i>DB2 for UNIX Quick Beginnings</i>	Provides planning, installation, migration, and configuration information for DB2 Universal Database on UNIX-based platforms. This book also contains installation and setup information for many supported clients.	GC09-2970 db2ixx70	db2ix
<i>DB2 for Windows Quick Beginnings</i>	Provides planning, installation, migration, and configuration information for DB2 Universal Database on Windows 32-bit operating systems. This book also contains installation and setup information for many supported clients.	GC09-2971 db2i6x70	db2i6
<i>DB2 Personal Edition Quick Beginnings</i>	Provides planning, installation, migration, and configuration information for DB2 Universal Database Personal Edition on the OS/2 and Windows 32-bit operating systems.	GC09-2969 db2i1x70	db2i1
<i>DB2 Personal Edition Quick Beginnings for Linux</i>	Provides planning, installation, migration, and configuration information for DB2 Universal Database Personal Edition on all supported Linux distributions.	GC09-2972 db2i4x70	db2i4

Table 1. DB2 Information (continued)

Name	Description	Form Number PDF File Name	HTML Directory
<i>DB2 Query Patroller Installation Guide</i>	Provides installation information about DB2 Query Patroller.	GC09-2959 db2iwx70	db2iw
<i>DB2 Warehouse Manager Installation Guide</i>	Provides installation information for warehouse agents, warehouse transformers, and the Information Catalog Manager.	GC26-9998 db2idx70	db2id
Cross-Platform Sample Programs in HTML			
Sample programs in HTML	Provides the sample programs in HTML format for the programming languages on all platforms supported by DB2. The sample programs are provided for informational purposes only. Not all samples are available in all programming languages. The HTML samples are only available when the DB2 Application Development Client is installed. For more information on the programs, refer to the <i>Application Building Guide</i> .	No form number	db2hs
Release Notes			
<i>DB2 Connect Release Notes</i>	Provides late-breaking information that could not be included in the DB2 Connect books.	See note #2.	db2cr
<i>DB2 Installation Notes</i>	Provides late-breaking installation-specific information that could not be included in the DB2 books.	Available on product CD-ROM only.	
<i>DB2 Release Notes</i>	Provides late-breaking information about all DB2 products and features that could not be included in the DB2 books.	See note #2.	db2ir

Notes:

1. The character *x* in the sixth position of the file name indicates the language version of a book. For example, the file name *db2d0e70* identifies the English version of the *Administration Guide* and the file name *db2d0f70* identifies the French version of the same book. The following letters are used in the sixth position of the file name to indicate the language version:

Language	Identifier
Brazilian Portuguese	b

Bulgarian	u
Czech	x
Danish	d
Dutch	q
English	e
Finnish	y
French	f
German	g
Greek	a
Hungarian	h
Italian	i
Japanese	j
Korean	k
Norwegian	n
Polish	p
Portuguese	v
Russian	r
Simp. Chinese	c
Slovenian	l
Spanish	z
Swedish	s
Trad. Chinese	t
Turkish	m

2. Late breaking information that could not be included in the DB2 books is available in the Release Notes in HTML format and as an ASCII file. The HTML version is available from the Information Center and on the product CD-ROMs. To view the ASCII file:
 - On UNIX-based platforms, see the `Release.Notes` file. This file is located in the `DB2DIR/Readme/%L` directory, where `%L` represents the locale name and `DB2DIR` represents:
 - `/usr/lpp/db2_07_01` on AIX
 - `/opt/IBMDB2/V7.1` on HP-UX, PTX, Solaris, and Silicon Graphics IRIX
 - `/usr/IBMDB2/V7.1` on Linux.
 - On other platforms, see the `RELEASE.TXT` file. This file is located in the directory where the product is installed. On OS/2 platforms, you can also double-click the **IBM DB2** folder and then double-click the **Release Notes** icon.

Printing the PDF Books

If you prefer to have printed copies of the books, you can print the PDF files found on the DB2 publications CD-ROM. Using the Adobe Acrobat Reader, you can print either the entire book or a specific range of pages. For the file name of each book in the library, see Table 1 on page 50.

You can obtain the latest version of the Adobe Acrobat Reader from the Adobe Web site at <http://www.adobe.com>.

The PDF files are included on the DB2 publications CD-ROM with a file extension of PDF. To access the PDF files:

1. Insert the DB2 publications CD-ROM. On UNIX-based platforms, mount the DB2 publications CD-ROM. Refer to your *Quick Beginnings* book for the mounting procedures.
2. Start the Acrobat Reader.
3. Open the desired PDF file from one of the following locations:
 - On OS/2 and Windows platforms:
x:\doc\language directory, where *x* represents the CD-ROM drive and *language* represent the two-character country code that represents your language (for example, EN for English).
 - On UNIX-based platforms:
/cdrom/doc/%L directory on the CD-ROM, where */cdrom* represents the mount point of the CD-ROM and *%L* represents the name of the desired locale.

You can also copy the PDF files from the CD-ROM to a local or network drive and read them from there.

Ordering the Printed Books

You can order the printed DB2 books either individually or as a set (in North America only) by using a sold bill of forms (SBOF) number. To order books, contact your IBM authorized dealer or marketing representative, or phone 1-800-879-2755 in the United States or 1-800-IBM-4Y0U in Canada. You can also order the books from the Publications Web page at <http://www.elink.ibm.com/pbl/pbl>.

Two sets of books are available. SBOF-8935 provides reference and usage information for the DB2 Warehouse Manager. SBOF-8931 provides reference and usage information for all other DB2 Universal Database products and features. The contents of each SBOF are listed in the following table:

Table 2. Ordering the printed books

SBOF Number	Books Included
SBOF-8931	<ul style="list-style-type: none"> • Administration Guide: Planning • Administration Guide: Implementation • Administration Guide: Performance • Administrative API Reference • Application Building Guide • Application Development Guide • CLI Guide and Reference • Command Reference • Data Movement Utilities Guide and Reference • Data Warehouse Center Administration Guide • Data Warehouse Center Application Integration Guide • DB2 Connect User's Guide • Installation and Configuration Supplement • Image, Audio, and Video Extenders Administration and Programming • Message Reference, Volumes 1 and 2 • OLAP Integration Server Administration Guide • OLAP Integration Server Metaoutline User's Guide • OLAP Integration Server Model User's Guide • OLAP Integration Server User's Guide • OLAP Setup and User's Guide • OLAP Spreadsheet Add-in User's Guide for Excel • OLAP Spreadsheet Add-in User's Guide for Lotus 1-2-3 • Replication Guide and Reference • Spatial Extender Administration and Programming Guide • SQL Getting Started • SQL Reference, Volumes 1 and 2 • System Monitor Guide and Reference • Text Extender Administration and Programming • Troubleshooting Guide • What's New
SBOF-8935	<ul style="list-style-type: none"> • Information Catalog Manager Administration Guide • Information Catalog Manager User's Guide • Information Catalog Manager Programming Guide and Reference • Query Patroller Administration Guide • Query Patroller User's Guide

DB2 Online Documentation

Accessing Online Help

Online help is available with all DB2 components. The following table describes the various types of help.

Type of Help	Contents	How to Access...
<i>Command Help</i>	Explains the syntax of commands in the command line processor.	<p>From the command line processor in interactive mode, enter:</p> <p style="padding-left: 40px;"><i>? command</i></p> <p>where <i>command</i> represents a keyword or the entire command.</p> <p>For example, <i>? catalog</i> displays help for all the CATALOG commands, while <i>? catalog database</i> displays help for the CATALOG DATABASE command.</p>
<i>Client Configuration Assistant Help</i>	Explains the tasks you can perform in a window or notebook. The help includes overview and prerequisite information you need to know, and it describes how to use the window or notebook controls.	From a window or notebook, click the Help push button or press the F1 key.
<i>Command Center Help</i>		
<i>Control Center Help</i>		
<i>Data Warehouse Center Help</i>		
<i>Event Analyzer Help</i>		
<i>Information Catalog Manager Help</i>		
<i>Satellite Administration Center Help</i>		
<i>Script Center Help</i>		
<i>Message Help</i>	Describes the cause of a message and any action you should take.	<p>From the command line processor in interactive mode, enter:</p> <p style="padding-left: 40px;"><i>? XXXnnnnn</i></p> <p>where <i>XXXnnnnn</i> represents a valid message identifier.</p> <p>For example, <i>? SQL30081</i> displays help about the SQL30081 message.</p> <p>To view message help one screen at a time, enter:</p> <p style="padding-left: 40px;"><i>? XXXnnnnn more</i></p> <p>To save message help in a file, enter:</p> <p style="padding-left: 40px;"><i>? XXXnnnnn > filename.ext</i></p> <p>where <i>filename.ext</i> represents the file where you want to save the message help.</p>

Type of Help	Contents	How to Access...
SQL Help	Explains the syntax of SQL statements.	<p>From the command line processor in interactive mode, enter:</p> <pre>help <i>statement</i></pre> <p>where <i>statement</i> represents an SQL statement.</p> <p>For example, help SELECT displays help about the SELECT statement.</p> <p>Note: SQL help is not available on UNIX-based platforms.</p>
SQLSTATE Help	Explains SQL states and class codes.	<p>From the command line processor in interactive mode, enter:</p> <pre>? <i>sqlstate</i> or ? <i>class code</i></pre> <p>where <i>sqlstate</i> represents a valid five-digit SQL state and <i>class code</i> represents the first two digits of the SQL state.</p> <p>For example, ? 08003 displays help for the 08003 SQL state, while ? 08 displays help for the 08 class code.</p>

Viewing Information Online

The books included with this product are in Hypertext Markup Language (HTML) softcopy format. Softcopy format enables you to search or browse the information and provides hypertext links to related information. It also makes it easier to share the library across your site.

You can view the online books or sample programs with any browser that conforms to HTML Version 3.2 specifications.

To view online books or sample programs:

- If you are running DB2 administration tools, use the Information Center.
- From a browser, click **File** —> **Open Page**. The page you open contains descriptions of and links to DB2 information:
 - On UNIX-based platforms, open the following page:

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

where %L represents the locale name.

- On other platforms, open the following page:

```
sql1lib\doc\html\index.htm
```

The path is located on the drive where DB2 is installed.

If you have not installed the Information Center, you can open the page by double-clicking the **DB2 Information** icon. Depending on the system you are using, the icon is in the main product folder or the Windows Start menu.

Installing the Netscape Browser

If you do not already have a Web browser installed, you can install Netscape from the Netscape CD-ROM found in the product boxes. For detailed instructions on how to install it, perform the following:

1. Insert the Netscape CD-ROM.
2. On UNIX-based platforms only, mount the CD-ROM. Refer to your *Quick Beginnings* book for the mounting procedures.
3. For installation instructions, refer to the `CDNAVnn.txt` file, where *nn* represents your two character language identifier. The file is located at the root directory of the CD-ROM.

Accessing Information with the Information Center

The Information Center provides quick access to DB2 product information. The Information Center is available on all platforms on which the DB2 administration tools are available.

You can open the Information Center by double-clicking the Information Center icon. Depending on the system you are using, the icon is in the Information folder in the main product folder or the Windows **Start** menu.

You can also access the Information Center by using the toolbar and the **Help** menu on the DB2 Windows platform.

The Information Center provides six types of information. Click the appropriate tab to look at the topics provided for that type.

Tasks	Key tasks you can perform using DB2.
Reference	DB2 reference information, such as keywords, commands, and APIs.
Books	DB2 books.
Troubleshooting	Categories of error messages and their recovery actions.
Sample Programs	Sample programs that come with the DB2 Application Development Client. If you did not install the DB2 Application Development Client, this tab is not displayed.
Web	DB2 information on the World Wide Web. To access this information, you must have a connection to the Web from your system.

When you select an item in one of the lists, the Information Center launches a viewer to display the information. The viewer might be the system help viewer, an editor, or a Web browser, depending on the kind of information you select.

The Information Center provides a find feature, so you can look for a specific topic without browsing the lists.

For a full text search, follow the hypertext link in the Information Center to the **Search DB2 Online Information** search form.

The HTML search server is usually started automatically. If a search in the HTML information does not work, you may have to start the search server using one of the following methods:

On Windows

Click **Start** and select **Programs** → **IBM DB2** → **Information** → **Start HTML Search Server**.

On OS/2

Double-click the **DB2 for OS/2** folder, and then double-click the **Start HTML Search Server** icon.

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

Note: The Search function is not available in the Linux, PTX, and Silicon Graphics IRIX environments.

Using DB2 Wizards

Wizards help you complete specific administration tasks by taking you through each task one step at a time. Wizards are available through the Control Center and the Client Configuration Assistant. The following table lists the wizards and describes their purpose.

Note: The Create Database, Create Index, Configure Multisite Update, and Performance Configuration wizards are available for the partitioned database environment.

Wizard	Helps You to...	How to Access...
<i>Add Database</i>	Catalog a database on a client workstation.	From the Client Configuration Assistant, click Add .
<i>Back up Database</i>	Determine, create, and schedule a backup plan.	From the Control Center, right-click the database you want to back up and select Backup → Database Using Wizard .

Wizard	Helps You to...	How to Access...
<i>Configure Multisite Update</i>	Configure a multisite update, a distributed transaction, or a two-phase commit.	From the Control Center, right-click the Databases folder and select Multisite Update .
<i>Create Database</i>	Create a database, and perform some basic configuration tasks.	From the Control Center, right-click the Databases folder and select Create → Database Using Wizard .
<i>Create Table</i>	Select basic data types, and create a primary key for the table.	From the Control Center, right-click the Tables icon and select Create → Table Using Wizard .
<i>Create Table Space</i>	Create a new table space.	From the Control Center, right-click the Table Spaces icon and select Create → Table Space Using Wizard .
<i>Create Index</i>	Advise which indexes to create and drop for all your queries.	From the Control Center, right-click the Index icon and select Create → Index Using Wizard .
<i>Performance Configuration</i>	Tune the performance of a database by updating configuration parameters to match your business requirements.	From the Control Center, right-click the database you want to tune and select Configure Performance Using Wizard . For the partitioned database environment, from the Database Partitions view, right-click the first database partition you want to tune and select Configure Performance Using Wizard .
<i>Restore Database</i>	Recover a database after a failure. It helps you understand which backup to use, and which logs to replay.	From the Control Center, right-click the database you want to restore and select Restore → Database Using Wizard .

Setting Up a Document Server

By default, the DB2 information is installed on your local system. This means that each person who needs access to the DB2 information must install the same files. To have the DB2 information stored in a single location, perform the following steps:

1. Copy all files and subdirectories from `\sql11ib\doc\html` on your local system to a Web server. Each book has its own subdirectory that contains all the necessary HTML and GIF files that make up the book. Ensure that the directory structure remains the same.

2. Configure the Web server to look for the files in the new location. For information, refer to the NetQuestion Appendix in the *Installation and Configuration Supplement*.
3. If you are using the Java version of the Information Center, you can specify a base URL for all HTML files. You should use the URL for the list of books.
4. When you are able to view the book files, you can bookmark commonly viewed topics. You will probably want to bookmark the following pages:
 - List of books
 - Tables of contents of frequently used books
 - Frequently referenced articles, such as the ALTER TABLE topic
 - The Search form

For information about how you can serve the DB2 Universal Database online documentation files from a central machine, refer to the NetQuestion Appendix in the *Installation and Configuration Supplement*.

Searching Information Online

To find information in the HTML files, use one of the following methods:

- Click **Search** in the top frame. Use the search form to find a specific topic. This function is not available in the Linux, PTX, or Silicon Graphics IRIX environments.
- Click **Index** in the top frame. Use the index to find a specific topic in the book.
- Display the table of contents or index of the help or the HTML book, and then use the find function of the Web browser to find a specific topic in the book.
- Use the bookmark function of the Web browser to quickly return to a specific topic.
- Use the search function of the Information Center to find specific topics. See “Accessing Information with the Information Center” on page 63 for details.

Appendix B. Notices

IBM may not offer the products, services, or features discussed in this document in all countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

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

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:
INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make

improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licenses of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Canada Limited
Office of the Lab Director
1150 Eglinton Ave. East
North York, Ontario
M3C 1H7
CANADA

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information may contain examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information may contain sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_. All rights reserved.

Trademarks

The following terms, which may be denoted by an asterisk(*), are trademarks of International Business Machines Corporation in the United States, other countries, or both.

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

The following terms are trademarks or registered trademarks of other companies:

Microsoft, Windows, and Windows NT are trademarks or registered trademarks of Microsoft Corporation.

Java or all Java-based trademarks and logos, and Solaris are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Tivoli and NetView are trademarks of Tivoli Systems Inc. in the United States, other countries, or both.

UNIX is a registered trademark in the United States, other countries or both and is licensed exclusively through X/Open Company Limited.

Other company, product, or service names, which may be denoted by a double asterisk(**) may be trademarks or service marks of others.

Index

A

- abstract data types 23
- ABSVAL function 37
- add database wizard 64, 65
- address cleansing 16
- Address Windowing Extensions (AWE), database buffer pool 31
- agents, multiple 34
- AIX, 64-Bit Support 32
- ALTER TABLE statement 30
- application development client 6
- apply, replication 26
- ARCHIVE LOG command 30
- archiving active log files 30
- AS/400
 - authID 40
 - schema 40
 - userID 40
- ASNSAT command 26
- Assist wizard
 - OLE DB 11
- Assist wizard, MQSeries 10
- atomic compound statements 24
- authID
 - AS/400 40
 - name 39
 - OS/390 40
- AWE, Address Windowing Extensions 31

B

- back up database wizard 64
- backup
 - closing logs 30
 - cross-platform 43
 - delta 35
 - incremental 35
 - named pipe 35
 - offline 33
 - online 33
 - split image 35
- backup, HP-UX and Solaris Operating Environment 43
- BIND options 41
- books 49, 59
- BPI model 18
- business intelligence
 - introducing 2, 15

- business process intelligence (BPI)
 - model 18

C

- CALL command 11
- calling stored procedures 11
- capture, replication 26
- capture and apply 26
- change target table 16
- class loader, Java 12
- closing logs after backup 30
- columns, identity 30
- columns, index key 30
- COM+ technologies 45
- command
 - ARCHIVE LOG 30
 - ASNSAT 26
 - db2cap 23
 - DB2UPDV7 43
- command prompt tag, importing 20
- COMMIT tag 20
- commit.tag, metadata template 20
- Common Warehouse Metamodel
 - i2 20
 - SAP 20
 - WebSphere Site Analyzer (WSA) 20
- compiled SQL procedures 43
- compound SQL, dynamic 24
- configuration parameter
 - locklist 32
 - TRACKMOD 35
- configure multisite update wizard 46, 64
- control center 15
- control center, OS/390 41
- control flow logic statements 24
- control statements, SQL 24
- converting dynamic SQL 23
- crash recovery 34
- create database wizard 65
- create parallel containers 37
- create table function, wizard 11
- create table space wizard 65
- create table wizard 65
- create target table 16
- creating SQL statements 29
- cross-platform, backup and restore 43

D

- data encryption 33
- data links 26
- data links manager 29
- data management
 - device containers 37
 - enhancements 29
 - introducing 2
- data management storage (DMS)
 - device containers 37
- data sets, managing 41
- data type, DATALINK 26
- data types 23
- data warehouse center 15
- database
 - backup 33
 - crash recovery 34
 - delta backup 35
 - dual file logging 34
 - incremental backup 35
 - named pipe backup 35
 - recovery 33
 - rollforward recovery 34
 - split image 35
 - suspended I/O 33
- database backup
 - HP-UX 43
 - Solaris Operating Environment 43
- database buffer pool
 - 64-Bit Support 32
 - Address Windowing Extensions (AWE) 31
 - size 31
- database management systems
 - AIX 21
 - Microsoft SQL Server 21
 - Oracle 21
 - Solaris Operating Environment 21
 - SQL queries 21
 - Sybase 21
- database middleware system 6
- database object, SEQUENCE 36
- database restore
 - HP-UX 43
 - Solaris Operating Environment 43
- database system monitor 22

- DataGuide, information control manager 20
 - DataJoiner 21, 26
 - DATALINK data type 26
 - datamarts 3
 - DataPropagator 26
 - DB2 CLI Unicode 32
 - DB2 Connect
 - 64-Bit Support 32
 - COM+ technologies 45
 - control center 41
 - data access 21
 - DB2 Relational Connect 4
 - DCL SNA product 45
 - dynamic SQL conversion 23
 - federated system 4
 - heterogeneous distributed query 21
 - Microsoft Transaction Server 45
 - multisite updates 45
 - OLE DB 29
 - single distributed query 21
 - static SQL profiling 23
 - Unicode standard 32
 - web starter kit 46
 - Windows 32-bit operating systems 42
 - DB2 Connect Enterprise Edition on Linux on S/390 47
 - DB2 Data Links Manager
 - Solaris Operating Environment 44
 - Tivoli Space Manager 44
 - DB2 Everyplace 39
 - DB2 family 39
 - DB2 family database 21
 - DB2 for Linux
 - parallel processing 12
 - shared-nothing architecture 12
 - DB2 for OS/390 46
 - DB2 library
 - books 49
 - Information Center 63
 - language identifier for books 57
 - late-breaking information 58
 - online help 60
 - ordering printed books 59
 - printing PDF books 58
 - searching online information 66
 - setting up document server 65
 - structure of 49
 - viewing online information 62
 - wizards 64
 - DB2 Life Sciences Data Connect 6
 - DB2 Net Search Extender
 - text search engine 5
 - DB2_NEWLOGPATH2, registry variable 34
 - DB2 OLAP server 15
 - DB2 OLE DB table function 19
 - DB2 Optimizer 5
 - DB2 Relational Connect 4, 21
 - DB2 Server for VM and VSE 46
 - DB2 table function, MQSeries message queue 19
 - DB2 Text Extender
 - search interface 5
 - DB2 Text Information Extender
 - HTML format 5
 - text retrieval 5
 - XML format 5
 - DB2 Universal Database Enterprise Edition, Linux 12
 - DB2 Universal Database Enterprise Edition for Linux on S/390 12
 - DB2 Universal Database Extended - Enterprise Edition on Linux 12
 - DB2 view 19
 - DB2 view, MQSeries message queue 19
 - DB2 Warehouse Manager Connectors
 - i2 TradeMatrix BPI 18
 - SAP R/3 18
 - Web 18
 - DB2 XML Extender interaction tool 17
 - DB2ARCHIVELOG 30
 - db2cap command 23
 - DB2INIDB utility 33, 35
 - db2look, structured data types 23
 - DB2UPDV7 command 43
 - DCL SNA product 47
 - DDL generation 41
 - DECRYPT_BIN function 33
 - DECRYPT_CHAR function 33
 - decrypting string data 33
 - delta backup 35
 - demographic augmentation, Trillium 16
 - DFS 29
 - DiscoveryLink 6
 - disk full error 31
 - distributed query 21
 - distributed transaction processing 45
 - DMS device containers 37
 - DTS, Data Transaction Services 19
 - dual logging 34
 - dynamic allocation, utility control using 41
 - dynamic compound statements control flow logic statements 24
 - SQL variables 24
 - dynamic SQL conversion 23
 - dynamic SQL statements 24
- ## E
- e-business
 - enhancements 9
 - introducing 2
 - e-Video Central tutorial 17
 - ENCRYPT function 33
 - encrypting string data 33
 - ERwin import tag 20
 - Essbase 5
 - export, schema name 40
 - extended markup language extenders 9
- ## F
- federated system
 - DB2 Relational Connect 4
 - feedback to IBM 2
 - filesystem, virtual 44
 - first steps, installation 1
 - foreignkey.tag, metadata template 20
 - FSM, filesystem 44
 - functions
 - ABS or ABSVAL 37
 - DECRYPT_BIN 33
 - DECRYPT_CHAR 33
 - ENCRYPT 33
 - GETHINT 33
 - MULTIPLY_ALT 37
 - ROUND 37
 - SQL-bodied 24
 - structured data type transform 24
 - structured data types 23
- ## G
- generate DDL 41
 - GETHINT function 33
 - global snapshots, system monitor 22
 - GUI 20
- ## H
- handheld devices 39
 - help wizards 2
 - heterogeneous database nicknames 21
 - SQL procedures 21
 - tables 21

- HP-UX
 - backup database 43
 - JDBC 13
 - restore database 43
- HP-UX, 64-Bit Support 32
- HTML
 - sample programs 57
- HTML, DB2 Text Information Extender 5
- I**
 - i2 load step 18
 - i2 OLAP cube 18
 - i2 source 20
 - i2 TradeMatrix BPI 18
 - i2 TradeMatrix BPI connectors 18
 - IBM ERwin MetaData Extract Program
 - command prompt tag 20
 - GUI 20
 - import tag 20
 - metadata, linking 20
 - schema 20
 - star schema 20
 - target tables 20
 - warehouse center 20
 - IBM WebSphere Application Server 17
 - IBM WebSphere Site Analyzer (WSA) 19
 - identity columns 30
 - import, schema name 40
 - import from
 - DB2 XML Extender Data Access Definition (DAD) file 19
 - MQSeries queue 19
 - incompatibilities, migrating 6
 - incremental backup 35
 - index keys
 - columns 30
 - index type, user-defined 25
 - index wizard 65
 - Information Center 2, 63
 - information control manager 20
 - information feedback 2
 - installation
 - first steps 1
 - launchpad 1
 - installing
 - Netscape browser 63
 - integration tools, Vality and Evolutionary Technologies, Inc. 16
 - isolation level 36
- J**
 - Java, writing stored procedures
 - in 12
 - Java class loader 12
 - JDBC
 - AIX 13
 - HP-UX 13
 - Linux 13
 - Solaris Operating Environment 13
 - JOB STEP 41
- K**
 - Kerberos 42
- L**
 - language identifier
 - books 57
 - late-breaking information 58
 - launchpad
 - installation 1
 - warehouse 1
 - library access 2
 - Linux
 - creating table spaces 13
 - JDBC 13
 - raw I/O 13
 - Linux on S/390 47
 - backup 12
 - JDBC 13
 - recovery 12
 - Linux on S/390, DB2 UDB EEE 12
 - LIST 41
 - LOAD authority 40
 - lock space 45
 - locklist limit 32
 - log
 - closing after backup 30
 - size limit 31
 - log archive 30
 - log file
 - archiving 30
 - log path 31
- M**
 - managing data sets 41
 - match, Trillium 16
 - merge, Trillium 16
 - metadata
 - COMMIT tag 20
 - import 20
 - interim commits 20
 - metadata extract program 20
 - metadata interchange 20
 - metadata templates
 - commit.tag 20
 - metadata templates (*continued*)
 - foreignkey.tag 20
 - primarykey.tag 20
 - primarykeyadditional.tag 20
 - methods, structured data types 23
 - Microsoft Data Transaction Services 18
 - Microsoft OLE DB 18
 - Microsoft OLE DB and Data Transaction Services 19
 - Microsoft OLE DB standard 11
 - Microsoft SQL Server databases 4, 21
 - Microsoft Transaction Server (MTS) 45
 - migrating
 - incompatibilities 6
 - to DB2 Enterprise - Extended Edition 7
 - to Version 7 6
 - visual warehouse 7
 - mirror utility, database backup 35
 - mobile computing
 - handheld devices 39
 - modeler, process 15
 - moving aggregates, SQL 16
 - MQSeries
 - Assist wizard 10
 - integration with WebSphere 17
 - message queue 18, 19
 - messages 9
 - user-defined function 10
 - LIST 41
 - MQSeries Assist wizard 10
 - MQSeries integration 10
 - MTS, Microsoft Transaction Server 45
 - multiple agents
 - crash recovery 34
 - database rollforward recovery 34
 - multiple COM+ components 45
 - MULTIPLY_ALT function 37
 - multisite updates 46
- N**
 - name cleansing 16
 - named pipe
 - backup 35
 - restore 35
 - nested stored procedures 11
 - Net.Data 9
 - net search extender 6
 - NetBackup 29
 - Netscape browser
 - installing 63
 - NEXTVAL expression 36

nicknames, SQL procedures 21

O

object relational 23
OLAP 16
OLAP integration server 15
OLAP starter kit 5, 15, 16
OLE DB 21, 29
OLE DB Assist wizard 11
OLE DB provider 19
online help 60
online information
 searching 66
 viewing 62
Oracle databases 4, 21
OS/390
 authID 40
 control center 41
 schema 40
 userID 40
 utilities 41
overview 10

P

packages, new 3
parallel containers, size 37
parallel processing, Linux 12
partitioning key, update 27
PDF 58
performance
 crash recovery 34
 database rollforward
 recovery 34
 DMS device containers 37
 prefetchers 37
performance configuration
 wizard 65
persistent stored module 11
primarykey.tag, metadata
 template 20
primarykeyadditional.tag, metadata
 template 20
printing PDF books 58
process modeler 15
products, new 3
PTF 45

Q

QMF 3, 22
query
 virtual database 6
query, distributed 21
query compiler 24
query optimizer, Linux 12
query patroller 4, 22
quick tour, DB2 2

R

red borders on required fields 16
referential integrity, typed tables 25
registry variable
 DB2_BLOCK_ON_LOG_DISK_FULL 31
 DB2_INDEX_2BYTEVARLEN 30
 DB2_NEWLOGPATH2 34
relational OLAP 16
release notes 58
remote computing 39
remote files 26
rename table space 31
REORG, structured data types 23
replication
 DataPropagator 26
 schema 40
 Solaris Operating
 Environment 27
 userID 40
required fields
 warehouse 16
restart OS/390 utilities 41
restore, HP-UX and Solaris
 Operating Environment 43
restore, named pipe 35
restore wizard 65
rollbacks using savepoints 31
rollforward recovery, database 34
ROUND function 37

S

sample database 1
sample programs
 cross-platform 57
 HTML 57
samples, stored procedures 11
SAP 20
SAP R/3 18
SAP R/3 connectors 18
satellites
 DB2 Universal Database Personal
 Edition 43
 DB2 Universal Database
 Workgroup Edition 43
savepoints 31
scalability, Linux 12
scalar function
 ABS or ABSVAL 37
 DECRYPT_BIN 33
 DECRYPT_CHAR 33
 ENCRYPT 33
 GETHINT 33
 MULTIPLY_ALT 37
 ROUND 37
schema
 AS/400 40
 name 39
 name for export 40
 name for import 40
 OS/390 40
 replication 40
schema, creating 20
searching
 online information 64, 66
 text 6
sequence object 36
setting up document server 65
shared-nothing architecture,
 Linux 12
Simple Object Access Protocol
 (SOAP) 9, 17
SmartGuides
 wizards 2, 64
snapshot utility, database
 backup 35
snapshots, system monitor 22
SOAP, Simple Object Access
 Protocol 9, 17
software developer's kit 6
Solaris Operating Environment 44
 backup database 43
 replication 27
 restore database 43
Solaris Operating Environment,
 64-Bit Support 32
Solaris Operating Environment, DB2
 Connect 47
 spatial extender 5
 specifying isolation level 36
split image, backup 35
split mirror, database 33
SQL Assist 29
SQL-bodied functions 24
SQL enhancements 16
SQL functions 24
 string data 33
SQL procedure language 11
SQL procedures
 compiled 43
 nicknames 21
SQL queries, DBMS 21
SQL statements 24
 dynamic 24
 MQSeries 10
 SQL Assist 29
SQL trapping with query
 patroller 22
SQL variables 24
standby utility, database backup 35

- star schema builder 15
- statement-level, isolation 36
- stored procedure builder 10
 - MQSeries Assist wizard 10
- stored procedures
 - authID 40
 - building 10
 - calling 11
 - Java 12
 - nested 11
 - samples 11
 - schema 40
 - SQL 11
 - userID 40
 - Visual Basic 11
- string data, encryption 33
- structured data type transforms 24
- structured data types 23
- suspended I/O 33
- Sybase databases 4, 21
- SYSIBM schema
 - functions
 - ABS or ABSVAL 37
 - MULTIPLY_ALT 37
 - ROUND 37
- Sysplex 45
- system monitor 22
- T**
- table function wizard 11
- table space
 - DB2 UDB for Linux 13
 - rename 31
- tables, temporary 23
- target table
 - change 16
 - create 16
 - name 16
 - table space 16
- target tables, creating 20
- TEMPLATE 41
- templates
 - Data Warehouse Center 20
 - metadata 20
- temporary tables 23
- text search engine 5
- Tivoli Space Manager (TSM) 44
- TRACKMOD configuration
 - parameter 35
- transform functions, structured data
 - types 24
- trapping SQL with query
 - patroller 22
- triggers
 - SQL 24
 - typed tables 25
- Trillium
 - address cleansing 16
 - demographic augmentation 16
 - matching 16
 - merging 16
 - name cleansing 16
- TSM, Tivoli Space Manager 44
- tutorial
 - business intelligence 17
 - e-Video Central 17
- typed tables 25
- U**
- Unicode standard,
 - enhancements 32
- updating partitioning key 27
- USE OF TABLESPACE authority 41
- user-defined extended index
 - type 25
- user-defined function, MQSeries 10
- user defined utility IDs 41
- userID
 - AS/400 40
 - name 39
 - OS/390 40
 - replication 40
- utilities, restarting OS/390 41
- utilities, user defined ID 41
- utility control statements 41
- V**
- Vality and Evolutionary
 - Technologies, Inc. 16
- variables, SQL 24
- viewing
 - online information 62
- virtual database 6
- Visual Basic stored procedures 11
- Visual C++ 42, 43
- visual studio 43
- visual warehouse, migrating 7
- W**
- warehouse 3
 - address cleansing 16
 - center 15
 - data marts 15
 - DataGuide 20
 - DB2 OLE DB table function 19
 - DB2 Relational Connect 4
 - DB2 table function 19
 - DB2 view 19
 - DB2 Warehouse Manager
 - Connectors 18
 - DB2 XML Extender Data Access
 - Definition (DAD) file 19
- warehouse 3 (*continued*)
 - i2 load step 18
 - i2 TradeMatrix BPI 18
 - IBM ERwin MetaData Extract
 - Program 20
 - information control manager 20
 - interim commits 20
 - launchpad 1
 - manager 3
 - metadata 20
 - Microsoft OLE DB and Data
 - Transaction Services (DTS) 19
 - MQSeries message queue 19
 - name cleansing 16
 - process modeler 15
 - red borders on required
 - fields 16
 - schema modeler 15
 - target tables 16
 - Trillium 16
 - web clickstream data 18
 - XML documents 19
 - XML objects 20
- warehouse, SAP R/3 18
- warehouse metadata 20
- warehouse source 16
- warehouse target 16
- warehouse target definition 19
- web clickstream data 18
- web connectors, warehouse 19
- web services 9
- web starter kit, DB2 Connect 46
- web traffic data, connectors 19
- Websphere Commerce data,
 - connectors 19
- WebSphere integration with
 - MQSeries 17
- Websphere runtime environment 9
- WebSphere Site Analyzer (WSA) 20
- wildcards, utility control using 41
- Windows
 - userID 42
 - Visual C++ 42
 - Windows 2000 42
- Windows, QMF for 22
- Windows 2000 42
- wizard
 - restore database 65
- wizards
 - add database 64, 65
 - back up database 64
 - completing tasks 64
 - configure multisite update 64
 - create database 65
 - create table 65

- wizards (*continued*)
 - create table space 65
 - DB2 table function 19
 - DB2 view 19
 - index 65
 - MQSeries Assist 10
 - MQSeries message queue 19
 - OLE DB Assist 11
 - performance configuration 65
 - Version 7 design 2
- workgroup edition
 - UNIX 5
- WSA, IBM WebSphere Site Analyzer 19
- WSA, WebSphere Site Analyzer 20

X

- XML, DB2 Text Information Extender 5
- XML documents, MQSeries queue 19
- XML extenders 9
- XML objects, CWM
 - exporting 20
 - importing 20

Contacting IBM

If you have a technical problem, please review and carry out the actions suggested by the *Troubleshooting Guide* before contacting DB2 Customer Support. This guide suggests information that you can gather to help DB2 Customer Support to serve you better.

For information or to order any of the DB2 Universal Database products contact an IBM representative at a local branch office or contact any authorized IBM software remarketer.

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

- 1-800-237-5511 for customer support
- 1-888-426-4343 to learn about available service options

Product Information

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

- 1-800-IBM-CALL (1-800-426-2255) or 1-800-3IBM-OS2 (1-800-342-6672) to order products or get general information.
- 1-800-879-2755 to order publications.

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

The DB2 World Wide Web pages provide current DB2 information about news, product descriptions, education schedules, and more.

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

The DB2 Product and Service Technical Library provides access to frequently asked questions, fixes, books, and up-to-date DB2 technical information.

Note: This information may be in English only.

<http://www.elink.ibm.com/pbl/pbl/>

The International Publications ordering Web site provides information on how to order books.

<http://www.ibm.com/education/certify/>

The Professional Certification Program from the IBM Web site provides certification test information for a variety of IBM products, including DB2.

ftp.software.ibm.com

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

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

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

On CompuServe: GO IBMDB2

Enter this command to access the IBM DB2 Family forums. All DB2 products are supported through these forums.

For information on how to contact IBM outside of the United States, refer to Appendix A of the *IBM Software Support Handbook*. To access this document, go to the following Web page: <http://www.ibm.com/support/>, and then select the IBM Software Support Handbook link near the bottom of the page.

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



Part Number: CT79QNA



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

SC09-2976-01



(1P) P/N: CT79QNA

