IBM® DB2 Universal Database™

What's New

Version 8

SC09-4848-00

IBM® DB2 Universal Database™

What's New

Version 8

SC09-4848-00

Before using this information and the product it supports, be sure to read the general information under *Notices*.

You can order IBM publications online or through your local IBM representative.

- To order publications online, go to the IBM Publications Center at www.ibm.com/shop/publications/order
- To find your local IBM representative, go to the IBM Directory of Worldwide Contacts at www.ibm.com/planetwide

To order DB2 publications from DB2 Marketing and Sales in the United States or Canada, call 1-800-IBM-4YOU (426-4968).

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.

Contents

DB2 Universal Database and DB2 Connect Version 8 provide cutting-edge business solutions. vii

Chapter 1. Introduction to DB2 Universal Database Version 8 1
Product and packaging changes 1
Component replacements and name changes . 1
Terminology changes 2
National language support 2
Discontinued and deprecated functions . . . 3
Migration to Version 8 5
Down level server support 5
Environments that should not be migrated to Version 8.1 8

Chapter 2. Manageability enhancements. . 11
Load enhancements 11
Storage Management tool 12
Flush the package cache 13
Logging enhancements 13
Backup and recovery enhancements 15
New administration notification log 15
Multiple service level install for UNIX . . . 16
Version ID for packages 16
Database Maintenance Mode: QUIESCE. . . 17
ON SCHEMA capability added to REORGCHK 17
RUNSTATS command enhancements. . . . 17
Tools for monitoring the health of your systems 18
Event monitors can now write to SQL tables 20
Improved deadlock event monitoring . . . 20
Snapshot and event monitoring: controlling time stamp collection 21
Snapshot retrieval through SQL 21
DB2 is Tivoli Ready 22
Type-2 indexes 22
Other manageability enhancements 23

Chapter 3. Performance enhancements . . 25
Multidimensional clustering. 25
Prefetching enhancements 25
Page cleaner I/O improvements 26
Catalog and authorization caching on databases with multiple partitions 27
Threading of Java UDFs and stored procedures 27
Connection Concentrator. 27
Materialized query tables 28
Query routing enhancements 28
User-maintained materialized query tables 28
Materialized query tables on nicknames. . 29
Incremental maintenance of materialized query tables using a staging table. . . . 29

Chapter 4. Availability enhancements. . . 31
Online table load 31
Online table reorganization 31
Online index reorganization. 32
Configurable online configuration parameters 32
Online buffer pool creation, deletion, and resizing 33
Incremental maintenance of materialized query tables during load append 34
DMS container operations 35

Chapter 5. Scalability enhancements . . . 37
Compression of nulls and defaults 37
INSERT through UNION ALL views. . . . 37
Logging enhancements for improved scalability 37
Changes to satellite administration 37

Chapter 6. Usability enhancements . . . 39
DB2 Administration Server (DAS) enhancements 39
DB2 Setup wizard enhancements 40
New wizards and GUI tools 40
New Configuration Assistant 42
Control Center 390 enhancements. 42
Structural changes to DB2 HTML documentation 42
Separately installed documentation 43
Regularly updated DB2 documentation . . . 44
Improved documentation search 44
Support for multi-language installations on Windows platforms 45

Chapter 7. Serviceability enhancements . . 47
INSPECT command for checking database architectural integrity 47
DB2 trace facility enhancements 47

Chapter 8. Replication enhancements . . 49
Replication Center 49
New replication control table structures . . . 50
Capture program enhancements 51
Multiple Capture programs 53
Apply program and other enhancements . . 53
Replication monitoring 54
DB2 Data Links Manager replication enhancements 55

Chapter 9. Data Warehouse Center and Warehouse Manager enhancements . . . 57
Warehouse server (AIX) 57
Warehouse agent (Linux). 57
Clean transformer enhancements 57
Common warehouse metamodel (CWM) XML support enhancements 58
Data Warehouse Center column mapping enhancements 58
Data Warehouse Center cascading processes 59
Multiple wait support. 59
SQL Select and Update step. 60
Information Catalog Manager enhancements 60

Chapter 10. Application development enhancements 63
Enhancements to routines (stored procedures, user-defined functions, and methods) . . . 63
Development Center 66
SQL Assist enhancements 67
Other SQL enhancements 68
Dynamic dispatch of methods 68
Windows management instrumentation . . . 68
New Call Level Interface (CLI) LOAD functionality 69
Declared global temporary table enhancements 69
IBM OLE DB Provider for DB2. 69
Web services applications 70
JDBC driver enhancements 71
WebSphere Studio and WebSphere Application Server available with DB2 Universal Developer's Edition 73
DB2 XML support enhancements 73

Chapter 11. Federated Systems 75
Federated Systems 75

Chapter 12. Business intelligence enhancements 77
DB2 OLAP Server and DB2 OLAP Integration Server changes 77
DB2 Spatial Extender enhancements 77

Chapter 13. DB2 family enhancements . . 79
Multiplatform tools for DB2 Universal Database 79
DB2 Data Links Manager enhancements . . 81
DB2 XML Extender enhancements 82
DB2 Net Search Extender 83
Host and iSeries applications can use two-phase commit with TCP/IP 83

Appendix A. DB2 Universal Database technical information 85
Overview of DB2 Universal Database technical information 85
FixPaks for DB2 documentation 85
Categories of DB2 technical information. . 85
Printing DB2 books from PDF files 93
Ordering printed DB2 books 94
Accessing online help 94
Finding topics by accessing the DB2 Information Center from a browser 96
Finding product information by accessing the DB2 Information Center from the administration tools 98
Viewing technical documentation online directly from the DB2 HTML Documentation CD. 100
Updating the HTML documentation installed on your machine 100
Copying files from the DB2 HTML Documentation CD to a Web server. . . . 102
Troubleshooting DB2 documentation search with Netscape 4.x. 103
Searching the DB2 documentation 104
Online DB2 troubleshooting information . . 105
Accessibility 105
Keyboard Input and Navigation 106
Accessible Display 106
Alternative Alert Cues 106
Compatibility with Assistive Technologies 106
Accessible Documentation 106
DB2 tutorials 107

DB2 Information Center accessed from a browser 108

Appendix B. Notices 109
Trademarks 112

Index. 115

Contacting IBM 119
Product information 119

DB2 Universal Database and DB2 Connect Version 8 provide cutting-edge business solutions

Database management software is now the core of enterprise computing. Companies need access to a wide range of information such as XML documents, streaming video, and other rich media types. New ways of working bring new requirements such as digital rights management. The e-business evolution makes continuous availability a necessity and is driving the convergence of transaction, business intelligence, and content management applications as companies integrate their business operations. DB2 Universal Database Version 8 can help your organization meet these challenges.

Whether your business is large or small today, you need a database that can grow with you. DB2 Version 8 is that database. It responds quickly to peaks in transaction demand on your web site, expands to hold growing amounts of information distributed in a number of different databases, and grows with your information infrastructure from one processor to multiple processors to massively parallel clusters. The integration of partitioning and clustering technology into the new DB2 Universal Database Enterprise Server Edition means that it is flexible enough to meet these needs. Now, you don't need to make all of your decisions today about your future growth.

DB2 Universal Database Version 8 also adds *self-managing and resource tuning* (SMART) database technology that lets database administrators choose to configure, tune and manage their databases with enhanced automation. SMART database management means administrators spend less time managing routine tasks and more time focussing on tasks that help enterprises gain and maintain a sustainable competitive advantage.

Industry-leading multidimensional clustering of tables reduces the indexing workload for your DBAs while providing data clustering for fast data query. This means you get better decision-making information sooner, with less system management overhead and improved data warehousing applications.

DB2's innovative query rewrite and optimization technologies and performance configuration capabilities lead the industry. This means you can spend more time analyzing business information and less time on performance improvement and tuning issues.

DB2's built-in planned and unplanned availability capabilities ensure your business applications are available whenever you need them. Whether switching to a standby database server if an unexpected database failure

occurs, or carrying out online maintenance, DB2 makes sure all of your business applications remain available. Online utilities like index rebuild, index create, and table load, as well as configuration parameters that can be changed without stopping the database, mean improved performance and high availability.

DB2 Version 8 delivers the right data management solutions for the new millennium. No other database provides the kinds of advanced performance, availability, scalability, and manageability features that Version 8 provides. We invite you to discover ″what's new″ with DB2 Universal Database Version 8.

Chapter 1. Introduction to DB2 Universal Database Version 8

Product and packaging changes

DB2® UDB Enterprise Edition (EE) and DB2 UDB Enterprise-Extended Edition (EEE) have been merged into a single product, DB2 UDB Enterprise Server Edition (ESE). The ability to create and manage multiple database partitions is part of the ESE product. If you want to create multiple database partitions on a single SMP server, you can do this using the base ESE product. If, however, you want to create multiple DB2 database partitions across more than one physical server (that is, a clustered hardware configuration), you must acquire a separate license.

DB2 UDB Workgroup Edition is now called DB2 UDB Workgroup Server Edition.

The following functions have been removed from the DB2 UDB Runtime Client:

- Client Configuration Assistant
- Command Center

The command line processor is included and can be used to administer the DB2 UDB Runtime Client.

The functionality of DB2 UDB Satellite Edition Version 6 has been merged into DB2 UDB Personal Edition Version 8. Existing DB2 UDB Satellite Edition customers are encouraged to begin the migration to DB2 UDB Personal Edition Version 8.

The DB2 OLAP Starter Kit is not available in DB2 Universal Database Version 8.

Related concepts:

- "Changes to satellite administration" on page 37

Component replacements and name changes

The Client Configuration Assistant has been renamed to the Configuration Assistant and significant enhancements to its function have been made.

The Development Center replaces the Stored Procedure Builder. The Development Center incorporates many more functions and features than the Stored Procedure Builder.

The Performance Configuration wizard has been renamed to the Configuration Advisor, and the Workload Performance wizard has been renamed to the Design Advisor.

Related concepts:

- "Development Center" on page 66
- "New Configuration Assistant" on page 42

Terminology changes

The following terminology changes have been made in Version 8:

- A long tablespace is now called a *large tablespace*.
- A nodegroup is now a *database partition group*.
- Before Version 8, the term *online index reorganization* was used to describe the process of merging index leaf pages while the index was online. This function is now referred to as *online index defragmentation of leaf pages*. A more complete online index reorganization facility has been added for Version 8.
- Country code has been changed to *territory code*.

Version 8 also introduces the term *materialized query table*. A materialized query table is a table whose definition is based on the result of a query, and whose data is in the form of precomputed results that are taken from one or more tables on which the materialized query table definition is based. A materialized query table whose fullselect contains a GROUP BY clause is summarizing data from the tables that are referenced in the fullselect. Such a materialized query table is also known as a *summary table*. A summary table is a specialized type of materialized query table.

Related concepts:

- "Online index reorganization" on page 32

National language support

In Version 8 all client-server data flows will use DRDA. From a code page conversion viewpoint, conversion will now be done on the receiver. Therefore, conversion tables will be installed on the client.

Support has been added for the following code pages:

- Arabic code page 425
- Latin-1 HOST code page 1047
- Unicode V3.1

Many code pages and code page conversion tables have been enhanced to provide support for the euro currency symbol (€). Support for the euro currency symbol is provided by default with these enhanced code pages.

If you want to use the non-euro versions of the code pages, they are available for download from ftp://ftp.software.ibm.com. This FTP server is anonymous. If you are connecting via the command line, log in as user "anonymous" and use your e-mail address as your password. After logging in, change to the conversion tables directory:

```
> cd ps/products/db2/info/vr8/conv
```

Related tasks:

- "Enabling and disabling euro symbol support" in the *Administration Guide: Planning*

Related reference:

- "Conversion table files for euro-enabled code pages" in the *Administration Guide: Planning*

Discontinued and deprecated functions

The following communication protocols are no longer supported:

- IPX/SPX as a DB2® client-server protocol. This means that DB2 Version 8 servers cannot accept IPX/SPX connections, and DB2 Version 8 clients cannot be configured to use IPX/SPX.
- SUNLINK SNA because SUN has announced that protocol's discontinuation.

The following operating system environments are not supported in Version 8:

- OS/2®
- PTX® or NUMA-Q®
- Windows® 95

The Generalize Replication Subscription function of the Satellite Administration Center is not supported.

The asnmobile function and the asnjet function are removed from Replication.

The db2alert.log (alerts) facility has been removed. Users of the db2alerts.log should use the administration notification log as a replacement.

The performance monitor capability of the Control Center has been removed. Users of the performance monitor should examine the functions of the Health Center (which is part of the Control Center) and the DB2 Performance Expert for Multiplatforms, Version 1 (a separate add-on tool) to replace the performance monitoring function they are using.

The type 3 JDBC driver (also called the 'net' driver or applet driver) is deprecated as of Version 8. This means that, although it will be fully supported in this release, you should migrate existing DB2 JDBC applets to the new type 4 driver, because the type 3 driver will be phased out entirely in a future version.

Distributed Computing Environment (DCE) changes

Support for DCE security has been removed in response to the industry move towards Kerberos as the mechanism for secure network authentication and single sign-on. Windows 2000 customers should consider moving to Kerberos as a replacement for DCE Security prior to migrating to DB2 UDB Version 8. Kerberos support for Windows 2000 was made available in Version 7.2. In a future delivery of DB2 UDB Version 8 will extend Kerberos support to UNIX and Linux servers and clients.

Lightweight Directory Access Protocol (LDAP) has become the industry standard for the enterprise class directory implementation. Customers should consider moving to LDAP as a replacement for DCE Directory support prior to migrating to DB2 UDB Version 8. LDAP support is available on all DB2 UDB Version 8 supported platforms with the exception of HP-UX and Linux. LDAP support for HP-UX and Linux is planned for future delivery of DB2 UDB Version 8.

Authentication at the DB2 Connect gateway is no longer possible

Version 8 removes the ability to authenticate a user at the DB2 Connect gateway. Authentication can only be done at the client (using CLIENT authentication) or at the server (using SERVER or SERVER_ENCRYPT). These options must be cataloged at the client in the database directory, or left as NOT_SPEC.

DCS and DCS_ENCRYPT now have exactly the same meaning as SERVER and SERVER_ENCRYPT. Any authentication entries in the database manager configuration file, or entries in the database catalog with authentication DCS or DCS_ENCRYPT will be migrated to SERVER and SERVER_ENCRYPT, respectively. Any attempt to specify DCS or DCS_ENCRYPT will result in these types being mapped to SERVER or SERVER_ENCRYPT.

If there is a DCS entry for a given database catalog entry (signifying that DB2 Connect is being used), and the authentication type is SERVER or SERVER_ENCRYPT, the type will be migrated to CLIENT to preserve existing behavior. If there is no DCS catalog entry, migration does not take place.

There is a case where a value of DCS in the database manager configuration file at the server had a different meaning than a value of SERVER. For federated systems, a value of DCS or DCS_ENCRYPT meant that no authentication would take place at the federated gateway; it was expected that authentication would take place at the final data source. A new database manager configuration parameter was added to accommodate this: FED_NOAUTH. If this parameter is set to YES, and the authentication type is SERVER or SERVER_ENCRYPT, there will be no authentication at the federated gateway, preserving the old behavior. This value is properly set during migration: If the authentication type is DCS or DCS_ENCRYPT in the database manager configuration file, FED_NOAUTH will be set to YES.

Related concepts:

- "Migrating Applications" in the *Application Development Guide: Building and Running Applications*
- "Changes to satellite administration" on page 37
- "New administration notification log" on page 15
- "Multiplatform tools for DB2 Universal Database" on page 79

Related tasks:

- "Alert Center -- Overview: Health Center help" in the *Help: Health Center*

Migration to Version 8

You can find information about migrating to Version 8 by following the links at the end of this section.

Down level server support

As you move your environment from Version 7 to Version 8, if you are in a situation where you migrate your client machines to Version 8 before you migrate all of your servers to Version 8, there are several restrictions and limitations. These restrictions and limitations are not associated with DB2 Connect; nor with zSeries, OS/390, or iSeries database servers.

For Version 8 clients to work with Version 7 servers, you need to configure/enable the use of DRDA Application Server capability on the Version 7 server. For information on how to do this, refer to the Version 7 *Installation and Configuration Supplement*.

To avoid the known restrictions and limitations, you should migrate all of your servers to Version 8 before you migrate any of your client machines to Version 8. If this is not possible, then you should know that when accessing Version 7 servers from Version 8 clients, there is no support available for:

- Some data types:
  - Large object (LOB) data types.
  - User-defined distinct types (UDTs).
  - DATALINK data types.

    The DATALINK data type allows for the management of external data found in non-relational storage. The DATALINK data type, references files that physically reside on file systems external to DB2 Universal Database.
- Some security capabilities:
  - Authentication type SERVER_ENCRYPT.

    SERVER_ENCRYPT is a method to encrypt a password. The encrypted password is used with the user ID to authenticate the user.
  - Changing passwords.

    You are not able to change passwords on the Version 7 server from a Version 8 client.
- Certain connections and communication protocols:
  - Instance requests that require an ATTACH instead of a connection.

    ATTACH is not supported from a Version 8 client to a Version 7 server.
  - The only supported network protocol is TCP/IP.

    Other network protocols like SNA, NetBIOS, IPX/SPX, and others are not supported.
- Some application features and tasks:
  - The DESCRIBE INPUT statement is not supported with one exception for ODBC/JDBC applications.

    In order to support Version 8 clients running ODBC/JDBC applications accessing Version 7 servers, a fix for DESCRIBE INPUT support must be applied to all Version 7 servers where this type of access is required. This fix is associated with APAR IY30655 and will be available before the Version 8 General Availability date. Use the “Contacting IBM” information in any DB2 Universal Database document to find out how to get the fix associated with APAR IY30655.

    The DESCRIBE INPUT statement is a performance and usability enhancement to allow an application requestor to obtain a description of input parameter markers in a prepared statement. For a CALL statement, this includes the parameter markers associated with the IN and INOUT parameters for the stored procedure.
  - Two-phase commit.

The Version 7 server cannot be used as a transaction manager database when using coordinated transactions that involve Version 8 clients. Nor can a Version 7 server participate in a coordinated transaction where a Version 8 server may be the transaction manager database.

– XA-compliant transaction managers.

  An application using a Version 8 client cannot use a Version 7 server as an XA resource. This includes WebSphere, Microsoft COM+/MTS, BEA WebLogic, and others that are part of a transaction management arrangement.

– Monitoring.
– Utilities.

  Those utilities that can be initiated by a client to a server are not supported when the client is at Version 8 and the server is at Version 7.

– SQL statements greater than 32 KB in size.

In addition to these limitations and restrictions for Version 8 clients working with Version 7 servers, there are also similar limitations and restrictions for Version 8 tools working with Version 7 servers.

The following Version 8 tools support only Version 8 servers:

- Control Center
- Task Center
- Journal
- Satellite Administration Center
- Information Catalog Center (including the Web-version of this center)
- Health Center (including the Web-version of this center)
- License Center
- Spatial Extender
- Tools Settings

The following Version 8 tools support Version 7 servers (with some restrictions) and Version 8 servers:

- Configuration Assistant (This tool has different components, of which only the import/export configuration file can be used with Version 7 servers; all of the components work with Version 8)
- Data Warehouse Center
- Replication Center
- Command Center (including the Web-version of this center)
- SQL Assist
- Development Center

- Visual Explain

In general, any Version 8 tool that is only launched from within the navigation tree of the Control Center, or any details view based on these tools, will not be available or accessible to Version 7 and earlier servers. You should consider using the Version 7 tools when working with Version 7 or earlier servers.

Related concepts:

- "Migrating Applications" in the *Application Development Guide: Building and Running Applications*
- "Migration recommendations" in the *Quick Beginnings for DB2 Servers*

Related tasks:

- "Migrating databases" in the *Quick Beginnings for DB2 Servers*
- "Migrating instances (UNIX)" in the *Quick Beginnings for DB2 Servers*
- "Migrating DB2 (Windows)" in the *Quick Beginnings for DB2 Servers*
- "Migrating DB2 (UNIX)" in the *Quick Beginnings for DB2 Servers*
- "Migrating DB2 Personal Edition (Windows)" in the *Quick Beginnings for DB2 Personal Edition*
- "Migrating DB2 Personal Edition (Linux)" in the *Quick Beginnings for DB2 Personal Edition*
- "Migrating databases on DB2 Personal Edition (Windows)" in the *Quick Beginnings for DB2 Personal Edition*
- "Migrating instances and databases on DB2 Personal Edition (Linux)" in the *Quick Beginnings for DB2 Personal Edition*

Related reference:

- "Migration restrictions" in the *Quick Beginnings for DB2 Servers*
- "Version 8 incompatibilities between releases" in the *Administration Guide: Planning*

Environments that should not be migrated to Version 8.1

The following environments should *not* be migrated to DB2® Universal Database Version 8:

DB2 Relational Connect and DB2 Life Sciences Data Connect environments
IBM is restructuring and enhancing its offerings to focus on information integration. These activities will include introducing new functionality that replaces and extends federated functionality previously available with DB2 Relational Connect and DB2 Life Sciences Data Connect. Details will be announced later.

Customers accessing federated data sources using DB2 Relational Connect Version 7 or DB2 Life Sciences Data Connect Version 7 should wait until this new functionality is available before upgrading to DB2 Universal Database Version 8.

DB2 Universal Database Version 8.1 has built-in capability to federate relational data across the IBM family of databases, including DB2 and Informix IDS. Customers who only want to use federated data from DB2 and Informix IDS can begin upgrading to Version 8.1.

DB2 Query Patroller environments

IBM plans to release Version 8 of DB2 Query Patroller, which is intended for use with DB2 Universal Database Version 8 databases. However, DB2 Query Patroller Version 8 is not being made available at this time. Customers using Version 7.2 of the DB2 Query Patroller Version 7.2, which shipped with DB2 Warehouse Manager Version 7.2, should not upgrade to DB2 Universal Database Version 8 until Version 8 of DB2 Query Patroller is available.

DB2 Query Patroller Version 8 will deliver enhanced functionality to better manage and control all aspects of query submission.

Chapter 2. Manageability enhancements

Load enhancements

Several enhancements have been made to the load utility in Version 8. New functionality has been added to simplify the process of loading data into both single partition and multi-partition database environments.

Load operations now take place at the table level. This means that the load utility no longer requires exclusive access to the entire table space, and concurrent access to other table objects in the same table space is possible during a load operation. Further, table spaces that are involved in the load operation are not quiesced. When the COPY NO option is specified for a recoverable database, the table space will be placed in the backup pending table space state when the load operation begins.

Another feature that has been added to the load utility is the ability to query pre-existing data in a table while new data is being loaded. You can do this by specifying the READ ACCESS option of the LOAD command.

The LOCK WITH FORCE option has also been introduced in this release. It allows you to force applications to release the locks they have on a table, allowing the load operation to proceed and to acquire the locks it needs.

You can now load data in partitioned database environments using the same commands (LOAD, db2load) and APIs (db2load, db2LoadQuery) used in single partition database environments. The AutoLoader utility (db2atld) and the AutoLoader control file are no longer needed.

Through the use of the new CURSOR file type, you can now load the results of an SQL query into a database without having to export them to a data file first.

Prior to Version 8, following a load operation the target table remained in check pending state if it contained generated columns. The load utility will now generate column values, and you no longer need to issue the SET INTEGRITY statement after a load operation into a table that contains generated columns and has no other table constraints.

The functionality of the LOAD QUERY command has also been expanded. It now returns the table state of the target table into which data is being loaded as well as the status information it previously included on a load operation in

progress. The LOAD QUERY command can be used to query table states whether or not a load operation is in progress on a particular table.

The Control Center now has a load wizard to assist you in the set up of a load operation.

Related concepts:

- "Loading Data in a Partitioned Database - Overview" in the *Data Movement Utilities Guide and Reference*
- "New wizards and GUI tools" on page 40

Related reference:

- "db2LoadQuery - Load Query" in the *Administrative API Reference*
- "LOAD QUERY Command" in the *Command Reference*
- "LOAD Command" in the *Command Reference*
- "db2Load - Load" in the *Administrative API Reference*

Storage Management tool

A Storage Management tool is now available through the Control Center. From this tool you can access the Storage Management view that displays a snapshot of the storage for a particular database, database partition group, or table space.

Statistical information can be captured periodically and displayed depending on the object chosen:

- For table spaces, information is displayed from the system catalogs and database monitor for tables, indexes and containers defined under the scope of the given table space.
- For databases or database partition groups, information is displayed for all the table spaces defined in the given database or database partition group.
- For databases, information is also collected for all the database partition groups within the database.

You can use the information displayed in this view to monitor various aspects of your storage, such as space usage for table spaces, data skew (database distribution) for database partition groups, and capture cluster ratio of indexes for database partition groups and table spaces.

From the Storage Management view you can also set thresholds for data skew, space usage, and index cluster ratio. A warning or alarm flag will let you know if a target object exceeds a specified threshold.

Related concepts:

- "New wizards and GUI tools" on page 40

Flush the package cache

This function introduces the ability to remove cached dynamic SQL statements from the package cache using a new SQL statement, FLUSH PACKAGE CACHE. This statement allows you to remove cached dynamic SQL statements from the package cache by invalidating them. The invalidation of a cached dynamic SQL statement allows current users of the cached statement to continue, but forces any new requests for the same statement to compile and create a new cached entry.

Although most normal activities that affect the validity of cached dynamic SQL statements are already handled by DB2® by invalidating the affected cached entries, certain activities, such as the new online update of database and database manager configuration parameters, are not. This statement allows you to manually invalidate cached dynamic SQL statements for those scenarios not automatically handled by DB2.

Related concepts:

- "Configuration parameter tuning" in the *Administration Guide: Performance*

Related tasks:

- "Configuring DB2 with configuration parameters" in the *Administration Guide: Performance*

Related reference:

- "Configuration parameters summary" in the *Administration Guide: Performance*
- "FLUSH PACKAGE CACHE statement" in the *SQL Reference, Volume 2*

Logging enhancements

Dual logging was introduced in Version 7.2 (also known as Version 7 FixPak 3). Dual logging was enabled by setting the DB2® registry variable DB2NEWLOGPATH2 to YES. Support for dual logging was limited to UNIX® and you did not have direct control over the path to which the second set of logs were written. The path defaulted to the value in the database configuration parameter *logpath* with a '2' appended.

In Version 8, dual logging is provided on all platforms supported by DB2 UDB. In addition the enablement and configuration of the second log path is

controlled by the database configuration parameter *mirrorlogpath*. When a value is present in *mirrorlogpath* dual logging is enabled, and the second log path is specified in the configuration parameter.

With single path logging, two database configuration parameters were used:

- *logpath* told you where it was currently.
- *newlogpath* was used to specify where the log path would be at the next database activation.

These parameters will continue to work the same way in Version 8. For the second log path in a dual logging configuration, there is only one configuration parameter: *mirrorlogpath*. Use the GET DB CFG SHOW DETAIL command to determine both its current value and the value it will have on the next database activation, which is reported in the delayed value column.

The maximum amount of log space that can be defined has increased from 32 GB to 256 GB. This provides support for more concurrent transactions and transactions that do more work.

Infinite active logging is also new in Version 8. Its allows an active unit of work to span the primary logs and archive logs, effectively allowing a transaction to use an infinite number of log files. Without infinite active log enabled, the log records for a unit of work must fit in the primary log space. Infinite active log is enabled by setting *logsecond* to -1. Infinite active log can be used to support environments with large jobs that require more log space than you would normally allocate to the primary logs.

The block on log disk full function that was introduced in Version 7 is now set using the database configuration parameter *blk_log_dsk_ful* in Version 8. Block on log disk full allows you to specify that DB2 should not fail when running applications on disk full condition from the active log path. When you enable this option, DB2 will retry every five minutes allowing you to resolve the full disk situation and allowing the applications to complete.

Related concepts:

- "Understanding Recovery Logs" in the *Data Recovery and High Availability Guide and Reference*
- "Log Mirroring" in the *Data Recovery and High Availability Guide and Reference*

Related reference:

- "Configuration Parameters for Database Logging" in the *Data Recovery and High Availability Guide and Reference*

Backup and recovery enhancements

DB2® UDB Version 8 includes several backup and recovery enhancements:

XBSA support
Backup can now interface with solutions from storage vendors that have implemented the XBSA industry standard interface.

Restoring to systems with different code pages
You can now restore a database backup to a system with a different code page. For example, a backup taken on an 819 code page system can be restored to an 850 code page system.

Faster table space recovery
During the recovery of a table space, only the log files required to recover the table space are processed; log files that are not required are skipped. If the log files are being retrieved from the archive, the user exit is only asked to retrieve the needed log files.

Point in time rollforward recovery to local time
When using point in time (PIT) rollforward recovery, you can now specify the time as local time. This makes it easier to rollforward to a specific point in time, and eliminates potential errors due to the translation of local to GMT time.

Related reference:

- "BACKUP DATABASE Command" in the *Command Reference*
- "RESTORE DATABASE Command" in the *Command Reference*
- "ROLLFORWARD DATABASE Command" in the *Command Reference*

New administration notification log

DB2® now places log information in two distinct logs, depending on its intended use.

Administration notification log
When significant events occur, DB2 writes information to the administration notification log. The information is intended for use by database and system administrators. Many notification messages provide additional information to supplement the SQLCODE that is provided. The type of event and the level of detail of the information gathered are determined by the NOTIFYLEVEL configuration parameter. However, detailed diagnostic information is not written to this log.

db2diag.log
Diagnostic information about errors is recorded in this text log file.

This information is used for problem determination and is intended for DB2 customer support. The level of detail of the information is determined by the DIAGLEVEL configuration parameter.

Multiple service level install for UNIX

Starting with FixPak 1, DB2® Version 8 supports multiple service (FixPak) levels simultaneously on the same system. This support is only available with DB2 UDB Enterprise Server Edition on UNIX® platforms.

Support for multiple service levels allows:

- The testing of a new service level while an older service level continues to support a production environment. Once testing is complete, the production environment can be switched to the new service level.
- Different teams to share one system with different DB2 code levels. For example, a team that has been developing an application for some time can continue to use the service level of DB2 that they have been using, while a team starting a new project can use the latest DB2 service level.

More information on multiple service levels will be available when Version 8 FixPak 1 is released.

Version ID for packages

You now have the ability to support multiple versions of the same package in a single system by using the new version option for packages.

The intent of this option is to allow multiple packages that share both the schema and package-id to co-exist in the system catalogs. This option will allow you to introduce and test a new version of a package on the system without affecting users of the existing version of the package. The support of the version option will allow ongoing package maintenance to occur without interruption of end-user access to the system.

The PREP, BIND, REBIND and DROP PACKAGE facilities have been enhanced to support package versioning.

Related concepts:

- "Package Versioning" in the *Application Development Guide: Programming Client Applications*

Database Maintenance Mode: QUIESCE

Using the new QUIESCE command, you can force all users off an instance or a database and put it into a quiesced mode for database maintenance activities.

The QUIESCE command provides exclusive access to the instance or database without having to force all users off the instance or database and then trying to stop users from attaching or connecting from outside the database engine (for example, by shutting down all transaction managers).

Only users with the correct authority can attach to the instance or connect to the database. During this quiesced period system administration can be performed on the instance or database. After administration is completed you can then unquiesce the database (using the UNQUIESCE command) and again allow other users to connect to the database, without having to shutdown and perform another database start.

Related reference:

- "QUIESCE Command" in the *Command Reference*
- "UNQUIESCE Command" in the *Command Reference*

ON SCHEMA capability added to REORGCHK

The capability of REORGCHK has been extended to allow you to specify that it should be run for a schema. This augments the current capability to run it for all tables, or for the tables created by the user issuing the command.

Related reference:

- "REORGCHK Command" in the *Command Reference*

RUNSTATS command enhancements

The RUNSTATS command is enhanced to improve the performance of statistics collection, and to provide additional options. The runstats utility now has the ability to:

- Collect additional statistics, such as statistics on column combinations, and prefetching statistics on the table, index, and index-to-table relationship
- Accept a list of index names (previously available only with the API)
- Accept a list of columns on which statistics are to be collected

- Accept distribution statistics limits: NUM_FREQVALUES and NUM_QUANTILES values at the table level (without having to change the configuration parameters, and then disconnect and reconnect all users)
- Accept individual column NUM_FREQVALUES and NUM_QUANTILES values
- Perform a faster (sampled) collection of DETAILED index statistics

Related concepts:

- "Guidelines for collecting and updating statistics" in the *Administration Guide: Performance*
- "Detailed index statistics" in the *Administration Guide: Performance*

Related tasks:

- "Collecting catalog statistics" in the *Administration Guide: Performance*
- "Collecting distribution statistics for specific columns" in the *Administration Guide: Performance*
- "Collecting index statistics" in the *Administration Guide: Performance*
- "Determining when to reorganize tables" in the *Administration Guide: Performance*

Related reference:

- "RUNSTATS Command" in the *Command Reference*
- "db2Runstats - Runstats" in the *Administrative API Reference*

Tools for monitoring the health of your systems

With Version 8, DB2 introduces two new features to help you monitor the health of your DB2 systems: the Health Monitor and the Health Center. These tools add a *management by exception* capability to DB2 Universal Database by alerting you to potential system health issues. This enables you to address health issues before they become real problems that affect your system's performance.

The Health Monitor is a server-side tool that constantly monitors the health of the instance, even without user interaction. If the Health Monitor finds that a defined threshold has been exceeded (for example, the available log space is not sufficient), or if it detects an abnormal state for an object (for example, an instance is down), the Health Monitor will raise an alert.

When an alert is raised two things can occur:

- Alert notifications can be sent by e-mail or to a pager address, allowing you to contact whoever is responsible for a system.

- Preconfigured actions can be taken. For example, a script or a task (implemented from the new Task Center) can be run.

A *health indicator* is a system characteristic that the Health Monitor checks. The Health Monitor comes with a set of predefined thresholds for these health indicators. The Health Monitor checks the state of your system against these health-indicator thresholds when determining whether to issue an alert. Using the Health Center, commands, or APIs, you can customize the threshold settings of the health indicators, and define who should be notified and what script or task should be run if an alert is issued.

The Health Center provides the graphical interface to the Health Monitor. You use it to configure the Health Monitor, and to see the rolled up alert state of your instances and database objects. Using the Health Monitor's drill-down capability, you can access details about current alerts and obtain a list of recommended actions that describe how to resolve the alert.

You can follow one of the recommended actions to address the alert. If the recommended action is to make a database or database manager configuration change, a new value will be recommended and you can implement the recommendation by clicking on a button. In other cases, the recommendation will be to investigate the problem further by launching a tool, such as the CLP or the new Memory Visualizer.

The Health Center and Control Center are integrated through Health Beacons. Health Beacons in the Control Center provide notifications about new alerts in the Health Center. Beacons are implemented on all Control Center windows and notebooks; simply click on a Health Beacon to access the Health Center.

Version 8 also provides a new Web Health Center that can be used to access the Health Monitor information from a Web browser or PDA.

You can also use new DB2 commands and APIs to retrieve health information from the Health Monitor, allowing you to integrate DB2 Health Monitoring with existing system-wide monitoring solutions.

The Health Monitor gathers information about the health of the system using new interfaces that do not impose a performance penalty. It does not turn on any snapshot monitor switches to collect information. The Health Monitor is enabled by default when a instance is created; you can deactivate it using the database manager configuration parameter *health_mon*.

Related tasks:

- "Alert Center -- Overview: Health Center help" in the *Help: Health Center*

Event monitors can now write to SQL tables

Event monitors can now write data to SQL tables instead of files or pipes. This provides the following advantages over PIPE and FILE event monitors:

- Event monitors gather large amounts of data and can take up a lot of disk space. Storing the data in an SQL table means you can define exactly which data elements you need to keep while discarding the rest. Storing the data in an SQL table makes it easy to prune the data using timestamps or other relational criteria.
- Event monitor data has been difficult to interpret, and it has been challenging to correlate event monitor data written to pipes and files with other data. With event monitor data located in tables, users can now query and aggregate data easily using SQL.
- Event monitors that write to tables can capture data on any or all database partitions. By inserting into the same table on each database partition, DB2® achieves true global event monitoring for event monitors.

Related concepts:

- "Event monitor table management" in the *System Monitor Guide and Reference*

Related tasks:

- "Creating a table event monitor" in the *System Monitor Guide and Reference*

Improved deadlock event monitoring

Enhancements to the Version 8 deadlock event monitor help system and database administrators determine why deadlocks occur.

The deadlock event monitor now provides more information than it did in previous releases. For example, the monitor now identifies which statements are involved in the deadlock and it pinpoints which locks each application involved in the deadlock is holding.

To reduce the wasted space associated with unnecessary connection header events, the monitor now only generates connection header events for the deadlock participants rather than for every connection to a database.

Related reference:

- "Locks and deadlocks monitor elements" in the *System Monitor Guide and Reference*

Snapshot and event monitoring: controlling time stamp collection

The collection of time and timestamp data elements is now controlled by the TIMESTAMP monitor switch. The switch is on by default. When you turn the switch off, the database manager skips any timestamp operating system calls when determining time or timestamp-related monitor elements.

When CPU usage approaches 100%, turning the switch off can provide a rapid performance boost. Also, turning the TIMESTAMP switch off can reduce the overall cost of other data under monitor switch control.

Related concepts:

- "System monitor switches" in the *System Monitor Guide and Reference*

Related tasks:

- "Setting monitor switches from a client application" in the *System Monitor Guide and Reference*
- "Setting monitor switches from the CLP" in the *System Monitor Guide and Reference*

Snapshot retrieval through SQL

Snapshots can now be taken using SELECT statements against a set of new user-defined table functions. Taking snapshots through SQL enables you to retrieve system status data from your application without having to implement DB2® API calls. You can now use all of the features of SQL to filter and join the different types of snapshot information.

You use can SQL to retrieve a new snapshot, or to process data from a previous DB2 API call that was written to a file. Use the new WRITE TO FILE option on a previous get snapshot API call to save data to a file.

Related tasks:

- "Capturing a database snapshot using SQL" in the *System Monitor Guide and Reference*

Related reference:

- "Snapshot monitor SQL table functions" in the *System Monitor Guide and Reference*

DB2 is Tivoli Ready

DB2® Version 8 is Tivoli® Ready. When a DB2 Version 8 product is installed, the required signature files are created so that Tivoli Inventory and Discovery can investigate a machine and find DB2.

The Tivoli Manager for DB2 is used for manageability purposes, including the following:

- Start and stop tasks for each server component
- Recovery tasks for each server component
- Process monitors for each server component
- An event adapter for use if the application collects or sends events or alerts
- Software distribution file packages for all desktop components
- Inventory signatures for all components
- Tivoli Global Enterprise Manager (GEM) instrumentation (Level 3) for all business application server components that connect to other applications
- Icons for application in X Pixmap Format (XPM) format

Type-2 indexes

Version 8 adds support for type-2 indexes. The primary advantages of type-2 indexes are:

- They improve concurrency because the use of next-key locking is reduced to a minimum. Most next-key locking is eliminated because a key is marked deleted instead of being physically removed from the index page. For information about key locking, refer to topics that discuss the performance implications of locks.
- An index can be created on columns that have a length greater than 255 bytes.
- A table must have only type-2 indexes before online table reorg and online table load can be used against the table.
- They are required for the new multidimensional clustering facility.

All new indexes are created as type-2 indexes, except when you add an index on a table that already has type-1 indexes. In this case the new index will also be a type-1 index because you cannot mix type-1 and type 2 indexes on a table.

All indexes created before Version 8 were type-1 indexes. To convert type-1 indexes to type-2 indexes, uses the REORG INDEXES command. To find out what type of index exists for a table, use the INSPECT command.

Related concepts:

- "Index performance tips" in the *Administration Guide: Performance*
- "Index cleanup and maintenance" in the *Administration Guide: Performance*

Related reference:

- "REORG INDEXES/TABLE Command" in the *Command Reference*
- "INSPECT Command" in the *Command Reference*

Other manageability enhancements

RENAME index

DB2® now allows you to rename indexes, which offers a time saving. The ability to rename an existing index allows you to first create a new index, then remove an old one, and change the name of the new index so that it is used in place of the old one, without any performance impact to your users.

Null and default compression

You can now save disk space for tables that will have many NULLs and SYSTEM DEFAULT values.

AUTOCONFIGURE command

AUTOCONFIGURE is a new command that recommends and optionally applies new values for buffer pool sizes, database configuration, and database manager configuration. This command provides initial tuning for a database, to which additional tuning can be applied.

AUTOCONFIGURE may also be used with the CREATE DATABASE command to configure databases as soon as they are created.

Related tasks:

- "Defining referential constraints" in the *Administration Guide: Implementation*

Related reference:

- "RENAME statement" in the *SQL Reference, Volume 2*
- "AUTOCONFIGURE Command" in the *Command Reference*

Chapter 3. Performance enhancements

Multidimensional clustering

Multidimensional clustering (MDC) provides an elegant method for flexible, continuous, and automatic clustering of data along multiple dimensions. This results in significant improvement in the performance of queries, as well as significant reduction in the overhead of data maintenance operations, such as reorganization, and index maintenance operations during insert, update, and delete operations. Multidimensional clustering is primarily intended for data warehousing and large database environments, and it can also be used in online transaction processing (OLTP) environments.

MDC enables a table to be physically clustered on more than one key (or dimension) simultaneously. Prior to Version 8, DB2® only supported single-dimensional clustering of data, through clustering indexes. Using a clustering index, DB2 maintains the physical order of data on pages in the key order of the index, as records are inserted and updated in the table. Clustering indexes greatly improve the performance of range queries that have predicates containing one or more keys of the clustering index. With good clustering, only a portion of the table needs to be accessed and, when the pages are sequential, more efficient prefetching can be performed.

With MDC, these benefits are extended to more than one dimension, or clustering key. In terms of query performance, range queries involving any combination of specified dimensions of the table will benefit from clustering. Not only will these queries access only those pages having records with the correct dimension values, these qualifying pages will be grouped by extents. Furthermore, although a table with a clustering index can become unclustered over time as space fills up in the table, an MDC table is able to maintain its clustering over all dimensions automatically and continuously, thus eliminating the need to reorganize the table to restore the physical order of the data.

Related concepts:

- "Multidimensional clustering" in the *Administration Guide: Planning*

Prefetching enhancements

In Version 8 prefetching can be improved by creating block based buffer pools.

When a block based buffer pool is available, the prefetching code recognizes this and will use block I/Os to read multiple pages into the buffer pool in a single I/O significantly improving the performance of prefetching. The BLOCKSIZE parameter of the CREATE and ALTER BUFFERPOOL SQL statement defines the size of the blocks, and hense the number of pages read from disk in a block I/O.

By default, the buffer pools are page-based, which means that contiguous pages on disk are prefetched into non-contiguous pages in memory. Sequential prefetching can be enhanced if contiguous pages can be read from disk into contiguous pages within a buffer pool.

You can create block-based buffer pools for this purpose. A block-based buffer pool consists of both a page area and a block area. The page area is required for non-sequential prefetching workloads. The block area consists of blocks where each block contains a specified number of contiguous pages, which is referred to as the block size.

Related concepts:

- "Prefetching data into the buffer pool" in the *Administration Guide: Performance*

Related reference:

- "ALTER BUFFERPOOL statement" in the *SQL Reference, Volume 2*
- "CREATE BUFFERPOOL statement" in the *SQL Reference, Volume 2*

Page cleaner I/O improvements

Version 8 exploits asynchronous I/O facilities to improve I/O performance. This can significantly improve page cleaning performance.

On AIX, asynchronous I/O is not always enabled; it must be enabled before DB2® Version 8 can be successfully installed.

Related reference:

- "Number of Asynchronous Page Cleaners configuration parameter - num_iocleaners" in the *Administration Guide: Performance*

Catalog and authorization caching on databases with multiple partitions

This feature extends the existing catalog cache to provide a cache at each partition of a partitioned database. The cached information will include SYSTABLE information and authorization information.

These caching enhancements will help to improve the overall performance of:

- Binding packages and compiling SQL statements, including usage of user-defined functions and stored procedures.
- Operations that involve checking database-level privileges.
- Operations that involve checking privileges for user-defined functions and stored procedures.

In particular, the performance of applications which are connected at non-catalog partitions will greatly improve.

Threading of Java UDFs and stored procedures

Routines (stored procedures, UDFs, and methods) are now implemented using a thread-based model that results in a dramatic perfomance increase for database servers running numerous routines. Routines that are defined as thread-safe will run in a single fenced-mode process. There is one process for Java™ routines and another process for non-Java routines to reduce the amount of context switching for users that run large numbers of fenced mode routines. For Java routines, this will also allow resource sharing of the JVM.

It is assumed that existing non-Java routines that are migrated to Version 8 are not thread-safe. Java routines will be migrated with the assumption that they are thread-safe. Users who want to modify pre-existing routines need to drop and re-create them, or use the appropriate alter SQL command. New routines are created with the aforementioned defaults if no thread-safe/non-thread-safe value is specified at creation. Again, non-Java routines will not be thread-safe, and Java routines will be thread-safe.

Connection Concentrator

For Internet applications with many relatively transient connections, or similar kinds of applications, the connection concentrator improves performance by allowing many more client connections to be processed efficiently. It also reduces memory use for each connection and decreases the number of context switches.

Related concepts:

- "Connection-concentrator improvements for client connections" in the *Administration Guide: Performance*

Materialized query tables

A materialized query table (MQT) is a table whose definition is based on the result of a query, and whose data is in the form of precomputed results that are taken from one or more tables on which the materialized query table definition is based.

Prior to Version 8, DB2 UDB supported summary tables, also known as automatic summary tables (ASTs). Summary tables are now considered to be a special type of MQT whose fullselect contains a GROUP BY clause summarizing data from the tables referenced in the fullselect.

The following MQT enhancements can result in improved query performance.

Query routing enhancements

Queries can now be routed to MQTs whose definitions contain a join that is not aggregated. Prior to Version 8, an MQT definition could only reference a join that was aggregated. For example, in Version 8 the following table, which contains a join, can be created to store the customer and account information for bad accounts:

```
CREATE TABLE bad_account AS
  (SELECT customer_name, customer_id, a.balance
    FROM account a, customers c
    WHERE status IN ('delinquent', 'problematic', 'hot')
    AND a.customer_id = c.customer_id)
    DATA INITIALLY DEFERRED REFRESH DEFERRED
```

If a user asks whether an account is delinquent, the DB2 UDB optimizer recognizes that the MQT has cached the requested information, and instead of accessing the base table ACCOUNT, DB2 accesses BAD_ACCOUNT, which provides a better response time and can be used to return customer information.

User-maintained materialized query tables

Many custom applications maintain and load tables that are really precomputed data representing the result of a query. By identifying a table as a user-maintained materialized query table, dynamic query performance can be improved. Such MQTs are maintained by users, rather than by the system. Update, insert, and delete operations are permitted against user-maintained MQTs.

Setting appropriate special registers allows the query optimizer to take advantage of the precomputed query result that is already contained in the user-maintained MQT.

Materialized query tables on nicknames

This feature allows you to cache remote data locally on your DB2 Universal Database instance. Remote data resides in databases that are supported by relational DBMS instances such as Oracle or Sybase, or even other instances of DB2 UDB.

MQTs can reference a combination of nicknames and local tables. Such materialized query tables can be created with the refresh deferred option only. Queries against nicknames or tables are rewritten and optimized in relation to these MQTs.

Routing a query to the MQT when all criteria for matching and routing are satisfied, yields better performance than getting results from the remote table.

It is possible to query a nickname even if the remote table for which the nickname was created becomes unavailable. If this nickname has a materialized query table defined on it, and all routing criteria match, the query will only need to select data from the MQT.

Maintenance is performed locally by means of the REFRESH TABLE command. (There is no way to keep track of updates to tables in a remote database.) Maintenance is always deferred; refresh immediate materialized query tables (defined on nicknames) are not supported.

Incremental maintenance of materialized query tables using a staging table

You can now incrementally refresh an MQT defined with the REFRESH DEFERRED option. If a refresh deferred MQT is to be incrementally maintained, it must have a staging table associated with it. The staging table associated with an MQT is created with the CREATE TABLE SQL statement.

When insert/delete/update statements modify the underlying tables of an MQT, the changes resulting from these modifications are propagated and immediately appended to a staging table as part of the same statement. The propagation of these changes to the staging table is similar to the propagation of changes that occurs during the incremental refresh of immediate MQTs.

A REFRESH TABLE statement is used to incrementally refresh the MQT. If a staging table is associated with the MQT, the system may be able to use the staging table that supports the MQT to incrementally refresh it. The staging table is pruned when the refresh is complete. Prior to Version 8, a refresh deferred MQT was regenerated from scratch when performing a refresh table

operation. MQTs can now be incrementally maintained, providing significant performance improvement. For information about the situations under which a staging table will not be used to incrementally refresh an MQT, see the "SQL Reference".

You can also use this new facility to eliminate the high lock contention caused by the immediate maintenance of refresh immediate MQTs. If the data in the MQT does not need to be current to the second, changes can be captured in a staging table and applied on any schedule.

Chapter 4. Availability enhancements

Online table load

When loading data into a table in Version 8, the table space in which the table resides will no longer be locked. Users have full read and write access to all the tables in the table space, except for the table being loaded. For the table being loaded, the existing data in the table will be available for read access if the load is appending data to the table.

These new load features significantly improve the availability of the data and help customers deal with the maintenance of large data volumes and shrinking maintenance windows.

Related concepts:

• "Load enhancements" on page 11

Online table reorganization

DB2® now provides two methods of reorganizing tables:

Online

Online table reorganization allows applications to access the table during the reorganization. In addition, online table reorganization can be paused and resumed later by anyone with the appropriate authority by using the schema and table name.

Online table reorganization is allowed only on tables with type-2 indexes and without extended indexes.

Offline

The offline method provides faster table reorganization, especially if you do not need to reorganize LOB or LONG data. LOBS and LONG data are no longer reorganized unless specifically requested. In addition, indexes are rebuilt in order after the table is reorganized. Read-only applications can access the original copy of the table except during the last phases of the reorganization, in which the shadow copy replaces the original copy and the indexes are rebuilt.

Both online and offline reorganizations have been enhanced to improve support for multi-partition databases. You can reorganize a single partition, a set of partitions, or all partitions.

Related concepts:

- "Table reorganization" in the *Administration Guide: Performance*

Related reference:

- "REORG INDEXES/TABLE Command" in the *Command Reference*

Online index reorganization

New for DB2® Version 8 is the ability to read and update a table and its existing indexes during an index reorganization using the new REORG INDEXES command.

During online index reorganization, the entire index object (that is, all indexes on the table) is rebuilt. A "shadow copy" of the index object is made, leaving the original indexes and the table available for read and write access. Any concurrent transactions that update the table are logged. Once the logged table changes have been forward-fitted and the new index (the shadow copy) is ready, the new index is made available. While the new index is being made available all access to the table is prohibited.

The default behavior of the REORG INDEXES command is ALLOW NO ACCESS, which places an exclusive lock on the table during the reorganization process, but you can also specify ALLOW READ ACCESS or ALLOW WRITE ACCESS to permit other transactions to read from or update the table.

Indexes can now be created in large table spaces (formerly long table spaces). In situations where the existing indexes consume more than 32 GB, this will allow you to allocate sufficient space to accommodate the two sets of indexes that will exist during the online index reorganization process.

Related reference:

- "REORG INDEXES/TABLE Command" in the *Command Reference*

Configurable online configuration parameters

Over 50 configuration parameters can now be set online. Changes to these configurable online configuration parameters take immediate effect without the need to stop and start the instance, or deactivate and activate the database. You no longer have to disconnect users when you fine tune your system, giving you more flexibility for deciding when to change the configuration.

Key database and database manager configuration parameters can be set online. For example, memory heaps such as catalogcache_sz, pckcache_sz, stmtheap, sortheap, util_heap_sz are dynamic allowing you to adjust memory usage as workloads vary over time. Other parameters such as locklist size, maxlocks, and dlchktime (dead lock check time) will allow you to adjust the locking characteristics of your database system which can improve performance.

You can choose to defer a change to a configurable online configuration parameter so that the configuration change will be made at the next instance start or database activation. A SHOW DETAILS option has been added to the GET DATABASE and GET DATABASE MANAGER CONFIGURATION commands that will list both the current value and the value that will be used at the next instance start or database activation.

In a few cases you can set the parameter you are configuring to automatic, and DB2® will then adjust its value automatically as workload on the system changes. For example, setting maxappls to automatic says there is no limit to the maximum number of applications, except when memory is exhausted. The GET DATABASE and GET DATABASE MANAGER CONFIGURATION commands have been changed to indicate the configuration values set to automatic and their current values.

Related tasks:

- "Configuring DB2 with configuration parameters" in the *Administration Guide: Performance*

Related reference:

- "Configuration parameters summary" in the *Administration Guide: Performance*

Online buffer pool creation, deletion, and resizing

In Version 8, you can change how DB2® uses memory without stopping database activity. You can change your buffer pool allocations and alter database and database manager configuration parameters that affect memory use while DB2 is running.

You can add a new buffer pool, alter the size of an existing buffer pool, or drop a buffer pool without stopping database activity. New options have been added to the CREATE and ALTER BUFFERPOOL SQL statements:

- IMMEDIATE, the default, will attempt to make the changes while DB2 is running.

- DEFERRED will allow you to defer the changes to the next database activation, and provides compatibility with the behavior of previous versions.

DROP BUFFERPOOL has immediate effect as it did previous versions. However, in Version 8 the memory is made available immediately to the database shared memory, and can be reused for other memory allocations.

With the ability to change your buffer pool allocations and update configuration parameters online you can now customize memory usage to the task. For example, if you have a prime shift memory allocation that is optimized for query performance (a large buffer pool), you could use a script to optimize memory usage for a load operation. The script would:

- Reduce the buffer pool size.
- Increase the utility heap and sort heap.
- Run loads with the configuration optimized for load.
- Return parameters to prime shift values when the loads are completed.

Related reference:

- "ALTER BUFFERPOOL statement" in the *SQL Reference, Volume 2*
- "CREATE BUFFERPOOL statement" in the *SQL Reference, Volume 2*

Incremental maintenance of materialized query tables during load append

A materialized query table is a dependent of an underlying table. Before Version 8, if data was appended to the underlying table during a load operation, the materialized query table was unavailable until the load completed and the materialized query table was maintained. Furthermore the materialized query table was completely rebuilt, which was often a lengthy operation.

With Version 8, the materialized query table can remain available during the load append operation on the underlying table. When the load of the appended rows in the underlying table is complete, the materialized query table may be refreshed incrementally using only the appended data, which significantly reduces the time to update it.

For example, if the materialized query table is an aggregate (automatic summary table), for rows appended to the underlying table that correspond to new groups in the aggregate, new summary rows will be inserted. For appended rows that correspond to existing groups in the aggregate, existing rows will be updated. Although the aggregate table remains unavailable

during this maintenance phase, when a small number of rows are appended to the underlying table (when compared to the size of the table), the time the aggregate is unavailable is reduced.

The ability to incrementally maintain a materialized query table is not restricted to aggregates. Many materialized query tables can be incrementally maintained. These changes significantly improve the availability of materialized query tables to your users.

Related concepts:

- "Refreshing Dependent Immediate Materialized Query Tables" in the *Data Movement Utilities Guide and Reference*

DMS container operations

If you use database-managed table spaces, Version 8 will allow you to drop a container from a table space, reduce the size of existing containers, and add new containers to a table space such that a rebalance does not occur.

- If a DMS table space was initially over-allocated, this can now be corrected.
- If the amount of data that resides in a table space has decreased significantly and the resulting ″waste″ in space is permanent, this extra space can now be reduced.
- You can add containers to a DMS table space using the new BEGIN STRIP SET option of the ALTER command, and a rebalance operation will not occur. Space added in this way is immediately available for use.

These operations can be accomplished online with full access to the table space, so there is no need to disconnect users.

Related concepts:

- "How containers are added and extended in DMS table spaces" in the *Administration Guide: Planning*
- "How containers are dropped and reduced in DMS table spaces" in the *Administration Guide: Planning*

Chapter 5. Scalability enhancements

Compression of nulls and defaults

Version 8 provides the ability to store nulls and system default values very efficiently. If you specify the VALUE COMPRESSION clause when creating a table, a new data row format will be used, which provides more efficient storage of NULL and 0 length values. It also enables system default values to be compressed when the COMPRESS SYSTEM DEFAULT column clause is used.

These changes can significantly reduce the disk storage required for large tables that contain nulls and system default values.

INSERT through UNION ALL views

You can now INSERT into UNION ALL views. This augments the UPDATE and DELETE capabilities provided in previous releases.

This can be very useful in HSM environments where the view ranges over a partitioned table hierarchy of which the parts are stored on tertiary storage. It also supports impved database scalability by providing SELECT, INSERT, UPDATE and DELETE support on a single view over multiple tables.

Logging enhancements for improved scalability

Version 8 logging now supports two features that improve scalability:

- An infinite active log which allows a single transaction to span both the primary logs and the archive logs.
- A maximum log space of up to 256 GB. This will provide support for larger workloads: more concurrent transactions and transactions that do more work.

Related concepts:

- "Logging enhancements" on page 13

Changes to satellite administration

The DB2CTLSRV instance and the SATCTLDB database can be created in a partitioned database environment. This enhancement allows you to exploit the scalability and additional processing power that is available in a partitioned database environment.

For conceptual and administrative information about the satellite environment, see the DB2 Version 6 manual, "Administering Satellites Guide and Reference", GC09-2821-00. You should ignore the information in that manual describing replication in a satellite environment. Updated information about the satellite environment is scheduled to be available after DB2 Version 8 is released.

Chapter 6. Usability enhancements

DB2 Administration Server (DAS) enhancements

For DB2® Version 8, the DB2 Administration Server (DAS) is no longer an instance. It is now a separate server process that supports TCP/IP communications.

The DAS is used to assist with tasks on DB2 servers. You must have a DAS running if you want to use tools like the Configuration Assistant, the Control Center, the Replication Center, or the Development Center.

The DAS assists the Control Center, Development Center, Replication Center, and Configuration Assistant when working on the following administration tasks:

- Enabling remote administration of DB2 servers.
- Providing the facility for job management, including the ability to schedule the running of both DB2 and operating system command scripts. These command scripts are user-defined.
- Defining the scheduling of jobs, viewing the results of completed jobs, and performing other administrative tasks against jobs located either remotely or locally to the DAS using the Task Center.
- On Windows and UNIX platforms, providing a means for discovering information about the configuration of DB2 instances, databases, and other DB2 administration servers in conjunction with the DB2 Discovery utility. This information is used by the Configuration Assistant and the Control Center to simplify and automate the configuration of client connections to DB2 databases.

In Version 8, the DB2 Administration Server will be available on zSeries™ (OS/390 and z/OS™ only) and iSeries™ platforms, as well as all supported Windows® and UNIX® platforms. The DAS on zSeries and iSeries can support the Control Center, Development Center, and Replication Center in administrative tasks on DB2 for zSeries and iSeries.

The Version 8 DAS on Windows and UNIX includes a scheduler to run tasks, such as DB2 and operating system command scripts, that are defined using the Task Center. Task information such as the commands to be run; schedule, notification, and completion actions associated with the task; and run results are stored in a set of tables called the Tools Catalog.

Related concepts:

- "DB2 Administration Server" in the *Administration Guide: Implementation*

DB2 Setup wizard enhancements

Installing DB2® on Windows® and UNIX® platforms has never been easier.

New for Version 8 on UNIX platforms is the DB2 Setup wizard, which was previously available only for DB2 on Windows platforms. The DB2 Setup wizard is a graphical user interface (GUI) that allows you to specify installation settings and initial configuration values. The DB2 Setup wizard replaces the db2setup text-based install utility on all UNIX platforms.

Also new on UNIX platforms for Version 8:

DB2 Launchpad
: Use the DB2 Launchpad GUI to access the Release Notes, installation prerequisites, and the Quick Tour, as well as the DB2 Setup wizard. The DB2 Launchpad was previously only available on Windows platforms.

DB2 Instance Setup wizard
: A new DB2 Instance Setup wizard is now available on UNIX platforms only to help you manage your instances after DB2 has been installed. You can create new instances or modify existing instances using this GUI tool.

On both Windows and UNIX platforms, you can now use the DB2 Setup wizard to generate a response file. A response file enables you to install DB2 on large numbers of computers simultaneously across your network. Response files contain the same values and parameters as those selected during the initial installation of DB2 on your local system, so you can be certain that each response file-based installation on your remote systems is the same.

New wizards and GUI tools

Several new wizards have been added to provide step-by-step guidance when you are creating objects, manipulating data, or configuring your environment.

- The Memory Visualizer helps you to uncover and fix memory-related problems on a DB2 instance. It uses visual displays and plotted graphs to help you understand memory components and their relationships to one another. You can invoke it from a Health Center recommendation or use it on its own as a monitoring tool.

- Use the Redistribute Data wizard to create an effective redistribution plan for your database partition group. You can use this wizard to redistribute data onto newly added database partitions, to redistribute data away from database partitions that you want to drop, or to create a customized redistribution plan.
- The Backup and Restore wizards have been rewritten to simplify their use. They also support the features added to backup and restore in Version 8.
- Use the Configure Database Logging wizard to configure data logging options for a database. You specify whether you want to use circular logging or archive logging, and how you want to handle your log files (archiving and media), and the wizard makes recommendations including the size of your primary and secondary log files, and the number and size of buffers.
- Use the Add Partitions wizard to create a partition and add it to one or more database partition groups. Use the wizard to add a new partition to your instance, assign the partition to one or more database partition groups, set the containers for the partition, and schedule the add operation.
- Use the Alter Database Partition Group wizard to modify an existing database partition group. You can specify whether you want to add an available partition to a database partition group, or drop an existing database partition from a database partition group.
- Use the Storage Management view to monitor the storage state of a partitioned database. In the Storage Management view, you can take storage snapshots for a database, a database partition group, or a table space.
- Use the Load wizard to load data into a selected table. The Load wizard guides you through load configuration and the selection of options. You can save the load configuration you define. The Load wizard also lets you copy an existing load task and use the setting values of the existing load task for your new load task.
- Use the new Design Advisor to optimize workload performance by creating a set of indexes recommended by DB2ADVIS. The Design Advisor replaces the Create Index wizard that was available in previous releases. The Design Advisor recommends which indexes to create in your database to improve performance for a given set of SQL statements in a workload.
- You can now determine the service level of DB2 installed on a system or instance using the About function on these objects. This function provides the same information as the **db2level** command, and it can also be run from a client against a remote server.

New Configuration Assistant

Previously known as the Client Configuration Assistant, the new Configuration Assistant has been tightly integrated with the Control Center and enhanced with many new features such as:

- The ability to invoke the Control Center from the Configuration Assistant.
- The option to configure both local and remote servers, including DB2® Connect servers.
- The ability to create configuration templates without affecting the local configuration.
- Import and export capabilities for exchanging configuration templates with other systems.
- Improved response time for discovery requests along with the option to refresh the list of discovered objects at any time.
- The ability to view and update applicable database manager configuration parameters and DB2 registry variables.

Control Center 390 enhancements

In Version 8, Control Center 390 support for DB2® for z/OS™ and OS/390® Version 6 and 7 servers has been significantly enhanced to support specific requirements of that platform. For example, Control Center 390 now supports:

- Copying an entire DB2 subsystem. The new Subsystem cloning wizard that generates the Job Control Language (JCL) required to homogeneously copy (clone) an entire subsystem to a target subsystem.
- Support for the DB2 unload utility has also been added to the Control Center for OS/390, which lets you unload data from one source object (i.e., a table space) to a sequential data set in external formats.
- You can select multiple table objects within the same table space.
- Utility execution is now optimized by supporting parallel execution.
- DB2 for OS/390 batch processing functions have been integrated into the Control Center: Build JCL, Create JCL and Custom JCL.

Structural changes to DB2 HTML documentation

The HTML topics are written as stand-alone Web pages that contain just the information you need to understand a concept or complete a task. Accessing documentation in HTML format gives you the following advantages:

Unified table of contents

The entire set of HTML topics appear within a single table of contents, so you no longer have to search through several different books to find the topic you need. Instead, drill down through the table of contents by the type of information that you need. To understand a concept, start at Concepts. To complete a task, start at Tasks. To view command syntax or get help for an error message, start at Reference. To launch your application development efforts quickly, view some sample code under Samples.

To enable the expanding and collapsing behavior of the unified table of contents, you must use one of the following browsers:

- Netscape 6.1 or greater
- Microsoft® Internet Explorer 5.0 or greater

Linking

Linking between topics has improved greatly. You will find links at the bottom of almost every topic, classified by the type of information to which it links. Topics can link to other topics, to tools help, or to external Web sites to provide you with the information you need to complete your tasks.

Master index

As an alternative means of accessing the topics, a master index is available in HTML format.

Separately installed documentation

The installation of DB2's product documentation is now independent of the product itself. You can install all of the documentation or a subset of it.

The HTML CD has its own installation program, so you can choose to install the separate HTML documentation during installation of DB2, or at any time after DB2® is installed.

One of the benefits of a separate install is that you can now install the documentation and Information Center on a machine that does not have DB2 installed. For example, you could install it on your company's internal web server rather than on your users' computer so that you save space on the individual machines and ensure that everyone has access to current information.

Regularly updated DB2 documentation

The HTML version of the DB2® information library is now updated with every DB2 FixPak. You can upgrade your documentation to the most recent version, regardless of the level of your product. Only documentation that was installed using the HTML CD documentation installation program can be updated using the documentation FixPak. The documentation FixPak will not update HTML that was copied directly from the HTML CD.

When you install DB2 tools such as the DB2 Control Center, help for these tools is installed at the same time. Updates to the tools help, when they are available, will be applied when you install a new DB2 FixPak.

If you install the DB2 HTML documentation, you can install any available updated topics from the DB2 Information Center by selecting **Information Center → Update Local Documentation** from the menu.

Whenever you update your information, you will be updating the English documentation as well as any other national language documentation on your computer. Information in all languages is kept at the same level.

A topic listing all the new and updated documentation for any given FixPak is available from the Information Center navigation tree.

PDF versions of the documentation are updated with every point release of DB2, and are available from the DB2 support site (`http://www.ibm.com/software/data/db2/udb/winos2unix/support`). A subset of this information will be refreshed with selected FixPaks. The PDF version of the Release Notes will contain a record of all updates since the last point release.

Improved documentation search

Searching DB2® online HTML documentation is now easy and convenient. All you need is a Java-enabled browser with Java™ support turned on. You must use Netscape 6.1 or higher, or Microsoft® Internet Explorer 5.0 or higher, if you want to use the HTML search function.

Search indexes for all supported languages are shipped with DB2, so you don't have to wait for indexes to be created before you perform a search.

Support for multi-language installations on Windows platforms

DB2® now supports the installation of multiple language locales with a single installation of DB2 on Windows® platforms.

When you install DB2 in a language other than English, you should be aware that English documentation and help is also installed automatically.

Switching the default system language also switches the language in which DB2 is displayed, assuming that the language was chosen during DB2 installation. Otherwise, English is displayed.

This enhancement does not require the Multilingual User Interface (MUI) Pack for Windows.

Chapter 7. Serviceability enhancements

INSPECT command for checking database architectural integrity

New with the release of DB2® Version 8 is the INSPECT command.

INSPECT allows you to inspect table spaces and tables for their architectural integrity while your database remains online.

Related reference:

- "db2Inspect - Inspect database" in the *Administrative API Reference*
- "INSPECT Command" in the *Command Reference*

DB2 trace facility enhancements

Enhancements to the DB2® trace facility mean improved efficiency. In most cases, the level of degradation is significantly reduced. As well, the size of the trace buffer is now larger so that information can be gathered more easily.

Related reference:

- "db2trc - Trace Command" in the *Command Reference*

Chapter 8. Replication enhancements

Replication Center

The DB2® Replication Center is a new user interface tool that you can use to set up and administer your replication environment. The Replication Center supports administration for DB2-to-DB2 replication environments, and administration for replication between DB2 and non-DB2 relational databases. The DB2 Replication Center is part of the DB2 Control Center set of tools. The Replication Center includes all of the replication functionality previously available from the DB2 Control Center and the DB2 DataJoiner® Replication Administration (DJRA) tool.

You can use the Replication Center to:

- Define defaults in profiles for creating control tables, source objects, and target objects
- Create replication control tables
- Register replication sources
- Create subscription sets and add subscription-set members to subscription sets
- Operate the Capture program
- Operate the Apply program
- Monitor the replication process

You can also use the Replication Center to perform many other replication administration tasks.

The Replication Center has a launchpad that allows you to perform the basic functions needed to set up a DB2 replication environment. The launchpad also shows you graphically how the different steps are related to one another.

You cannot use the Version 8 Control Center to administer a replication environment; use the Replication Center. As of Version 8, the DB2 DataJoiner Replication Administration (DJRA) tool is not available with DB2.

Related concepts:

- "Using the DB2 Replication Center" in the *Replication Guide and Reference*

Related tasks:

- "Getting started with the Replication Center: Replication Center help" in the *Help: Replication Center*

New replication control table structures

DB2® stores all definitional and operational information for replication in the replication control tables. Substantial changes have been made to the control table structures in Version 8 to support new function and to improve usability. New tables have been added, some existing tables have been changed, and a few tables have been made obsolete.

The control tables used by replication now support 128-byte table names and 30-byte column names for databases that support long names.

The following new tables have been added:

- IBMSNAP_APPENQ ensures that only one Apply program is running for a single Apply qualifier.
- IBMSNAP_CAPENQ ensures that only one Capture program is running for a single Capture schema.
- IBMSNAP_CAPMON contains operational statistics for monitoring the progress of the Capture program.
- IBMSNAP_CAPSCHEMAS contains the names of all Capture schemas.
- IBMSNAP_PRUNE_SET coordinates the pruning of CD tables.
- IBMSNAP_RESTART enables the Capture program to resume capturing from the correct point in the log or journal.
- IBMSNAP_SIGNAL contains signals used to control the Capture program.

The following new tables have been added for the Replication Alert Monitor:

- IBMSNAP_ALERTS contains a history of all alerts issued by the Replication Alert Monitor.
- IBMSNAP_CONDITIONS contains alert conditions for each monitored server.
- IBMSNAP_CONTACTGRP maps contacts with groups.
- IBMSNAP_CONTACTS contains contact names and addresses.
- IBMSNAP_GROUPS contains contact groups.
- IBMSNAP_MONENQ ensures that only one Monitor process is running for a single Monitor qualifier.
- IBMSNAP_MONSERVERS contains the most recent time that the Replication Alert Monitor monitored a Capture or Apply control server.
- IBMSNAP_MONTRACE traces Replication Alert Monitor activity.
- IBMSNAP_MONTRAIL contains a history of Monitor activity for every Monitor cycle.

Most of the previously existing replication control tables have been updated to support changes for the Version 8 replication architecture.

The following tables from previous versions of DB2 replication are now obsolete:

- IBMSNAP_CRITSEC is replaced by IBMSNAP_SIGNAL.
- IBMSNAP_WARMSTART is replaced by IBMSNAP_RESTART.

The migration utility (ASNMIG8) ensures that any data from control tables in your existing DB2 replication environment are migrated to your Version 8 replication environment.

DB2 replication does not support the new control tables for z/VM, VSE, or Microsoft® Jet systems. You can use the DB2 Control Center or DJRA to administer your Version 7 control tables for z/VM, VSE, and Microsoft Jet.

Related reference:

- "List of tables used at the Apply control server" in the *Replication Guide and Reference*
- "List of tables used at the Capture control server" in the *Replication Guide and Reference*
- "Tables at a glance" in the *Replication Guide and Reference*
- "List of tables used at the Monitor control server" in the *Replication Guide and Reference*

Capture program enhancements

You can start the Capture and Apply programs in any order. The new replication control tables ensure that the Capture and Apply programs communicate effectively with each other, and with the Replication Alert Monitor.

The default startmode for the Capture program has been changed. Using this new default startmode, the Capture program will perform a cold start when you start it for the first time, but will only perform warm starts on subsequent restarts, and will not automatically switch to a cold start.

In previous releases, the Capture program could either capture or prune data, but could not do both tasks at once. In Version 8, it can do both. A separate pruning thread is initialized when the Capture program is started, and it remains active while the Capture program is running.

Note: These improvements already exist in previous versions of DB2® DataPropagator™ for AS/400.

Normal and retention-limit pruning are performed by these separate threads so that performance throughput does not suffer at the expense of keeping data storage minimized.

You can supply operational parameters for the Capture program when you start it. If the Capture program is already active, you can dynamically change the operational parameters.

The Capture program is controlled by signals that you or the Apply program store in the Signal (IBMSNAP_SIGNAL) table. The signal table provides a way to communicate with the Capture program through log records. Not only does the signal table let the Apply program tell the Capture program when to start capturing data, it also allows for precise termination of log record reading and for user-defined signals through log records. The Capture program uses these signals to determine when to start or stop capturing changes for a particular table, and whether it must perform update-anywhere replication. It can also provide the log sequence number to set a precise end point for Apply events.

When you register a table for replication, you can specify the row-capture rule for that object. For each registered table, you can specify whether you want the Capture program to capture changes for a row whenever *any* column of the table changes or only when a *registered* column changes. For example, if you want to register a table with 100 columns for replication, but you want to capture changes for only two of the columns and only when these two columns are updated, you would specify the row-capture rule when you register those two tables. In previous versions of DB2, you could not specify the row-capture rule for each registration. You had to specify the row-capture rule when you started the Capture program; therefore, the row-capture rule was applied to all registered tables even if you wanted the extra processing only on a small subset of tables. The start-up parameter is not available in Version 8.

You can add columns to your replication source without reinitializing the Capture program or stopping and starting it. The extra column values are not added to the change data (CD) table. On UNIX, Windows, and z/OS™ systems, you can also alter the CD table while the Capture program is running.

In update-anywhere replication, a change made to a replica table is captured and replicated to all other tables. If you have a multi-tier scenario, or if you have replica tables that are logical partitions of the source table, you might not want changes from one replica table forwarded to all the others. In Version 8, you can specify if you want changes recaptured and forwarded to other replicas, or if you want the Capture program to ignore changes that have already been captured by another replica table.

In Version 7, you could create only one Windows® service to operate all of your Capture and Apply programs. In Version 8, you can create separate services for each Capture and Apply program and for the Replication Alert Monitor, and start or stop them individually.

Related tasks:

- "Operating the Capture program" in the *Replication Guide and Reference*

Multiple Capture programs

You can run multiple instances of the Capture program on a given server by creating multiple sets of Capture control tables. All relevant Capture control tables have a unique table schema, which is called the *Capture schema*. The default Capture schema is ASN.

You can start each of the Capture programs with different operational parameters, and each program runs independently of the other. You might want to use multiple Capture programs on a given server for the following reasons:

- To support multiple non-DB2 relational sources on a single federated server.
- For z/OS™ operating systems, to support a mix of ASCII, EBCDIC, and UNICODE source tables.
- To achieve higher throughput and lower latency.

Related tasks:

- "Operating the Capture program" in the *Replication Guide and Reference*

Apply program and other enhancements

In Version 8, joins between replication control tables have been eliminated in some situations, resulting in substantial improvements in performance. For example:

- User copy target tables can be populated without a join on the fetch.
- The column IBMSNAP_UOWID is replaced by IBMSNAP_COMMITSEQ in change data (CD) tables. The IBMSNAP_COMMITSEQ column allows the Apply program to process certain target table types without having to join the CD table with the UOW table.
- CD tables that are not involved in replica scenarios with conflict detection can be pruned without a join. In cases where a join between the CD table and the UOW table is required, the join is made using the IBMSNAP_COMMITSEQ column.

You can choose to have the Apply program commit changes transactionally, in commit-sequence order for all subscription-set members at the same time, and you can specify how many transactions the Apply program should process before committing its work. This feature is useful in a data warehouse scenario or if there are referential integrity constraints on the target table.

In Version 7, you could create only one service to operate all of your Capture and Apply programs. In Version 8, you can create separate services for each Capture and Apply program and for the Replication Alert Monitor, and start or stop them individually.

In Version 7, the password file used by the Apply program and the Replication Analyzer contained plain text, not encrypted information. In Version 8, the passwords in the password file are encrypted and no passwords are stored in plain text. A new command (**asnpwd**) enables you to create and maintain the password file.

Related tasks:

- "Operating the Apply program" in the *Replication Guide and Reference*

Replication monitoring

You can monitor your Version 8 replication environment in the following ways:

Set up monitoring so that you're notified when something occurs

The Replication Alert Monitor monitors the Capture and Apply programs for you, based on the criteria you provide when you set it up. In the Replication Center, you define alert conditions for criteria that you want to monitor, and specify people who should be contacted automatically when those conditions are met. For example, you might want to set a threshold for the amount of memory used by the Capture program. If memory usage exceeds the set value, notification is sent by the Replication Alert Monitor to the contact that you specified. For the Apply program, you might want to define an alert condition so that notification is sent if a transaction is rejected in an update-anywhere configuration.. You can also specify that you want to be notified when error conditions occur.

You can start the Replication Alert Monitor using either the Replication Center or the **asnmon** command. You can use a Monitor qualifier to start more than one instance of the Monitor program. Starting multiple Monitor programs allows you to distribute the workload among Monitors or to ensure that a mission-critical application has a dedicated Monitor process.

Query current replication processes and statistics

You can spot-check how the replication programs are doing:

- You can query the status of each thread of the Capture and Apply programs. Use this type of monitoring, for example, to determine whether the Capture program is capturing or is idle. Status is available in the Replication Center for each Capture or Apply control server, or by using the **asnccmd** and **asnacmd** commands.
- You can show details of statistics stored in certain replication control tables. Use the Replication Center to display statistics that help you troubleshoot problems for any Capture schema or Apply qualifier. You can refresh the query to update the statistics in the display.

Analyze trends and historical data in replication statistics

You can use the Replication Center to query historical data to review how your replication environment performed during a particular time interval. You can generate reports, then analyze the following information:

- Messages: What error and informational messages have the Capture and Apply programs issued?
- Capture throughput: On average, how many rows were processed in the CD table for a given period of time?
- Capture latency: How recently did the Capture program commit data into the CD table?
- Apply throughput: On average, how many rows were processed in the target table for a given period of time?
- End-to-end latency: On average, how much time elapsed between the time when the Capture program read the changes from the DB2® log and the Apply program replicated the changes to the target tables?

Related tasks:

- "Monitoring replication" in the *Replication Guide and Reference*
- "Replication Alert Monitor : Replication Center help" in the *Help: Replication Center*

DB2 Data Links Manager replication enhancements

If you have a DATALINK value pointing to an external file, you can retrieve consistent versions the file if the column is defined with RECOVERY YES. In past releases, DB2® would replicate the latest copy of the file and could not guarantee that the file being replicated was consistent with the replicated database data values.

You can maintain the same target file across multiple changes in the source database.

For the AIX® and Windows® operating systems, and the Solaris operating environment, you can connect to the DB2 Data Links Manager replication daemon (DLFM_ASNCOPYD), which is part of the DB2 Data Links Manager product, to retrieve and store linked files for replication. You do not need to start and maintain a separate ASNDLCOPYD daemon as in previous releases. On OS/400, you still have to start and maintain a separate ASNDLCOPYD daemon.

Related concepts:

- "DB2 Data Links Manager enhancements" on page 81

Chapter 9. Data Warehouse Center and Warehouse Manager enhancements

Warehouse server (AIX)

For version 8, DB2® Warehouse Manager capabilities have been extended to include support for the warehouse server and associated logger, initialization, external trigger, and mail notification programs on AIX. This support removes the Windows® dependency for the warehouse server.

Warehouse agent (Linux)

For version 8, DB2® Warehouse Manager capabilities have been extended to include support for the warehouse agent, programs, and transformers on Linux (for 32-bit Intel processors). Linux distributions with the following levels are supported: kernel 2.4.7, glibc 2.2.4.

Clean transformer enhancements

The clean transformer has been rewritten in Version 8. Improvements have been made to enhance the clean transformer's performance and function. The following cleaning rules are provided:

- Carry over
- Find and replace
- Discretize into ranges
- Clip
- Convert case
- Encode invalid values

New enhancements include:

- Better null support
- More data types supported
- Generate all, first, and last matches
- Error on multiple or no matches
- Differentiator and ordering columns for rules
- Use of automatic summary tables

The Version 7.2 clean transformer will remain in the Data Warehouse Center as a deprecated program.

Some of these new functions and performance enhancements are not available for the zSeries™ and iSeries™ platforms.

Related concepts:

- "Clean Data transformer" in the *Data Warehouse Center Administration Guide*

Related tasks:

- "Cleaning data : Data Warehouse Center help" in the *Help: Data Warehouse Center*

Common warehouse metamodel (CWM) XML support enhancements

The common warehouse metamodel (CWM) is a model-based architecture that contains features that are effective in, and mappable to, a broad range of representative warehouse configurations. CWM 1.0 was published in February 2001 and is adopted as an OMG standard. Metadata that is either exported or imported in this format can be interchanged with other warehousing tools that interpret metadata that conforms to the CWM standard.

You can now use the Common Warehouse Metamodel (CWM) utility to import or export CWM objects from an XML file and to convert these objects to Data Warehouse Center objects.

Related tasks:

- "Importing MQSeries messages and XML metadata" in the *Data Warehouse Center Administration Guide*
- "Exporting a tag language file: Data Warehouse Center help" in the *Help: Data Warehouse Center*
- "Importing objects : Information Catalog Center help" in the *Help: Data Warehouse Center*

Related reference:

- "Metadata mappings between the Data Warehouse Center and CWM XML objects and properties" in the *Data Warehouse Center Administration Guide*

Data Warehouse Center column mapping enhancements

The column mapping feature in the Data Warehouse Center provides you with even more control over how columns are mapped between sources and targets. The enhancements in this release include:

- The Source Columns and target columns windows can scroll either separately or together.
- You can map all columns by column name and data type, by column name only, or by position.
- You can remove all column mappings at one time.
- You can now edit the names and types of new target columns in the column mapping page.
- You can search for a column in either the source or target lists using the Find Columns feature.

Related concepts:

- "Column mapping" in the *Data Warehouse Center Administration Guide*

Related tasks:

- "Creating a column mapping: Data Warehouse Center help" in the *Help: Data Warehouse Center*

Data Warehouse Center cascading processes

In DB2® Version 7.2 you could use task flow features to run a step based on the success, failure, or completion of another step. In DB2 Version 8, you can trigger a step based on the success, failure, or completion of multiple steps by defining and enabling a schedule and task flow for the processes that contain the steps. In addition to having more control over how steps run, you can also monitor the progress of cascaded processes in the Work In Progress window.

Related concepts:

- "Process task flow" in the *Data Warehouse Center Administration Guide*

Related tasks:

- "Creating a cascade relationship based on success or failure of a publication : Data Warehouse Center help" in the *Help: Data Warehouse Center*
- " Creating cascade relationships : Data Warehouse Center help" in the *Help: Data Warehouse Center*

Multiple wait support

In Version 8, you can specify through the user interface that a step should wait on multiple steps to complete before running the dependent step. The mechanism for specifying the multiple wait is a process cascade tree. The user can specify the cascaded process will be scheduled upon success, failure, or completion of another process.

Related concepts:

- "Process task flow" in the *Data Warehouse Center Administration Guide*

SQL Select and Update step

With this new step, you can update a target table in the data warehouse without needing to replace the entire table or write additional code. The SQL Select and Update step allows the rows to be updated in the target table based on the key column mapping. The UPDATE statement is generated automatically. After it is generated, the UPDATE statement can be edited before processing, and the update can be processed in place and scheduled to run on demand.

Related concepts:

- "SQL steps" in the *Data Warehouse Center Administration Guide*

Related tasks:

- "Defining an SQL step" in the *Data Warehouse Center Administration Guide*
- "Defining an SQL step: Data Warehouse Center help" in the *Help: Data Warehouse Center*

Information Catalog Manager enhancements

In today's business environment, having access to important company information is a necessity. Gathering pertinent information is often a time-consuming process, and involves searching for data in various locations. In Version 8, the Information Catalog Manager has been enhanced to add a new user interface, the Information Catalog Center. In Version 8, the Information Catalog Manager has been enhanced to add a new user interface, the Information Catalog Center. It manages the business metadata. The Information Catalog Center streamlines the search for business and technical metadata by providing the capability to organize, navigate, and search through metadata objects.

The Information Catalog Manager still contains a web interface for end users, but this interface, the Information Catalog Center for the Web, is a Websphere application that supports the new Information Catalog Manager. The Information Catalog Manager contains a Java™ API that can be used by third party applications to manipulate metadata in the information catalog.

The Information Catalog Center can be invoked from Windows® and UNIX platforms either directly or from one of the other DB2® centers. The Information Catalog Center implements object-level security so that metadata

objects can be viewed by only authorized users. It also contains the ability for administrators to not only create object types, but also relationship types to better organize your metadata for end users. The Information Catalog Manager interface tools are included in the DB2 Administration Client. The DB2 Warehouse Manager also contains these tools plus the ability to initialize a new information catalog.

Additional enhancements for administrators and users of the information catalog include an improved search capability, additional data types for property values, constraints on permitted values for properties, and usability enhancements to tailor the user interface.

Related concepts:

- "Information Catalog Center" in the *Information Catalog Center Administration Guide*

Related tasks:

- "Getting started with the Information Catalog Center" in the *Information Catalog Center Administration Guide*

Chapter 10. Application development enhancements

Enhancements to routines (stored procedures, user-defined functions, and methods)

The term *routine* is used to encompass stored procedures, UDFs, and methods. This reflects the fact that as of DB2 Version 8, parameter styles, data type mappings, and system catalogs are the same for all three routine types.

The following are the main enhancements for routines:

Catalog views

The catalog views for functions, methods, and stored procedures have been merged.

- SYSCAT.ROUTINES describes all the routines that are defined in the database.
- SYSCAT.ROUTINEPARMS describes the parameters to those routines (functions and methods), as well as their return information.

The previous catalog views are still supported to allow existing applications that depend on those views to run.

New privilege for invoking routines

The routine EXECUTE privilege has been defined to explicitly control who can invoke routines (stored procedures, UDFs, and methods). When the routine is used in an SQL statement, the routine definer must have the EXECUTE privilege on any packages used by the routine.

New authorities for registering external routines

New authorities have been defined to explicitly control who can register external routines (stored procedures, UDFs, and methods).

The CREATE_EXTERNAL_ROUTINE authority is required to register external routines. FENCED THREADSAFE routines run as threads inside a shared process. Each of these routines are able to read the memory used by other routine threads in the same process. Therefore, it is possible for one threaded routine to collect sensitive data from other routines in the threaded process. Another risk inherent in the sharing of a single process is that one routine thread with flawed memory management can corrupt other routine threads, or cause the entire threaded process to crash. When granting the CREATE_EXTERNAL_ROUTINE authority, be aware that the recipient

can potentially monitor or corrupt the memory of other FENCED THREADSAFE routines. For more information, see ″Threading of Java UDFs and stored procedures″.

The CREATE_NOT_FENCED_ROUTINE authority is required to register NOT FENCED routines. NOT FENCED routines run in the same process as the database manager, and as a result, can corrupt the database manager's shared memory, or damage the database control structures. Either form of damage will cause the database manager to fail. NOT FENCED routines can also corrupt databases and their tables. When granting the CREATE_NOT_FENCED_ROUTINE authority, be aware that the recipient can potentially gain unrestricted access to the database manager and all its resources. If you have CREATE_NOT_FENCED_ROUTINE authority, you also have CREATE_EXTERNAL_ROUTINE authority.

ALTER statement for external routines

Each routine type now has a an ALTER statement that can be used to change the routine's EXTERNAL NAME to reference a new routine body. For example, instead of dropping and re-registering a Java stored procedure when you need to update its Java method, use the ALTER PROCEDURE statement to make the routine use a new Java method instead of the old method.

Enhancements to NOT FENCED routines

- NOT FENCED routines support nesting and recursion. There are no restrictions on the types of routines that can be nested. For example, FENCED routines can invoke NOT FENCED routines, and vice-versa.
- NOT FENCED stored procedures can return result sets.
- SQL procedures are automatically registered as NOT FENCED.

Enhancements for Java routines

Java routines now support recursion. There are no restrictions on the types of routines that can be nested. For example, Java routines can invoke routines written in other languages, and vice-versa.

The introduction of a thread-based model for implementing routines results in improved performance and scalability for Java routines. For more information, see ″Threading of Java UDFs and stored procedures″.

Library management enhancements

The DB2 library manager dynamically adjusts its library caching according to your workload. For optimal performance consider the following:

- Keep the number of routines in your libraries as small as possible. If you are including multiple routines in the same library, ensure

that you group them based on if they are invoked in the same time frame. With a library caching scheme, it is better to have numerous smaller libraries than few large libraries.

- The load cost for a library in the THREADSAFE C process is paid only once for libraries that are consistently in use by THREADSAFE C routines. After the routine's first invocation, all subsequent invocations, from any thread in the process, do not need to load the routine's library.

SQL in external user-defined functions and methods

External user-defined functions and methods can now contain read-only SQL statements. Both static and dynamic SQL can be used. Other routines can be invoked from a function or method. Such nesting can include recursive invocations of the routine and routines written in different languages. The limit of 90 parameters for PROGRAM TYPE MAIN stored procedures has been removed.

CALL is now a fully compiled statement

The CALL statement is now a fully compiled statement. This means that the CALL statement can now be dynamically prepared in CLI, ODBC, embedded SQL, JDBC, and SQLj. Input arguments to a stored procedure call can be expressions. These arguments are subject to data type checking and promotion.

When invoking procedures from a client application, do not use a host variable as the procedure name in the CALL statement. This allows for package cache reuse, thus enabling subsequent stored procedure invocations to avoid system catalog lookups and compilation.

Related concepts:

- "SQL in External Routines" in the *Application Development Guide: Programming Server Applications*
- "Performance Considerations for Developing Routines" in the *Application Development Guide: Programming Server Applications*
- "Threading of Java UDFs and stored procedures" on page 27
- "Dynamic SQL in SQL-Bodied Routines" in the *Application Development Guide: Programming Server Applications*
- "Security Considerations for Routines" in the *Application Development Guide: Programming Server Applications*
- "Restrictions for Routines" in the *Application Development Guide: Programming Server Applications*
- "Library and Class Management Considerations for Developing Routines" in the *Application Development Guide: Programming Server Applications*

Related reference:

- "CALL statement" in the *SQL Reference, Volume 2*
- "CREATE FUNCTION statement" in the *SQL Reference, Volume 2*
- "CREATE PROCEDURE statement" in the *SQL Reference, Volume 2*
- "CREATE TYPE (Structured) statement" in the *SQL Reference, Volume 2*
- "GRANT (Database Authorities) statement" in the *SQL Reference, Volume 2*
- "REVOKE (Database Authorities) statement" in the *SQL Reference, Volume 2*
- "SYSCAT.ROUTINEPARMS catalog view" in the *SQL Reference, Volume 1*
- "SYSCAT.ROUTINES catalog view" in the *SQL Reference, Volume 1*
- "CREATE METHOD statement" in the *SQL Reference, Volume 2*
- "GRANT (Routine Privileges) statement" in the *SQL Reference, Volume 2*
- "REVOKE (Routine Privileges) statement" in the *SQL Reference, Volume 2*
- "ALTER FUNCTION statement" in the *SQL Reference, Volume 2*
- "ALTER METHOD statement" in the *SQL Reference, Volume 2*
- "ALTER PROCEDURE statement" in the *SQL Reference, Volume 2*

Development Center

In DB2® Version 8, the Development Center replaces the Stored Procedure Builder. The Development Center has many more functions and features than the Stored Procedure Builder. The Development Center provides an easy-to-use interface for developing routines such as stored procedures and user-defined functions (UDFs). A set of wizards makes it easy to perform your development tasks. The Development Center provides a single development environment that supports the entire DB2 family ranging from the workstation to z/OS. You can launch the Development Center as a stand-alone application from the IBM® DB2 Universal Database™ program group or from a DB2 Universal Database center, such as the Control Center, the Command Center, or the Task Center.

With the Development Center, you can:

- Create, build, and deploy Java™ and SQL stored procedures
- Create, build, and deploy user-defined functions:
  - SQL table and scalar UDFs
  - UDFs that read MQSeries® messages
  - UDFs that access OLE DB data sources
  - UDFs that extract data from XML documents
- Debug SQL stored procedures using the integrated debugger

- See the contents of the server for each database connection that is in your project or that you have explicitly added to the Server View. You can also view and work with other database objects such as tables, triggers, and views.
- Export and import routines and project information.

The Development Center also provides a DB2 development add-in for each of the following development environments:

- Microsoft® Visual C++
- Microsoft Visual Basic
- Microsoft Visual InterDev

With the add-ins, you can easily access the features of the Development Center and other DB2 centers from your Microsoft development environment, making it easy for you to develop and incorporate stored procedures and UDFs into your DB2 application development.

Related tasks:

- "About the Development Center: Development Center help" in the *Help: Development Center*

SQL Assist enhancements

With SQL Assist and some knowledge of SQL, you can create SELECT, INSERT, UPDATE, and DELETE statements. SQL Assist is a tool that uses an outline and details panels to help you organize the information that you need to create an SQL statement.

Version 8 enhancements include:

- A re-designed user interface for easy inspection and modification of SQL statement elements
- Assistance for creating table joins
- SQL syntax checking
- The option to copy and paste an existing SQL statement into SQL Assist, then use the SQL Assist interface to make modifications

SQL Assist version 7.2 remains in the Data Warehouse Center as a deprecated product.

Related tasks:

- "Overview of IBM SQL Assist: SQL Assist help" in the *Help: SQL Assist*

Other SQL enhancements

INSTEAD OF triggers are supported

INSTEAD OF triggers are used to perform updates, deletes, inserts, and selects transparently against views even when the view may be too complex to support update operations natively. The use of this type of trigger can simplify application interfaces.

New built-in functions

The following trigonometric functions have been added to the SYSIBM schema: ATANH, COSH, SINH, TANH, ACOS, ASIN, ATAN, ATAN2, COS, COT, SIN, and TAN.

ATANH, COSH, SINH, and TANH are new. The other functions were available under the SYSFUN schema; having them in SYSIBM offers a performance advantage.

TRUNCATE, CEILING, and FLOOR have also been added to the SYSIBM schema. These built-in functions now support decimal values.

Informational constraints

Version 8 introduces a new type of constraint called *informational constraints*. Informational constraints are rules that can be used in query rewrite to improve performance but are not enforced by the database manager.

Often, constraints are enforced by the logic in business applications and it is not desirable to use system enforced constraints since re-verification of the constraints on insert, update and delete operations can be costly. In this case, informational constraints are a better alternative.

Dynamic dispatch of methods

You can now override methods–that is, you can reimplement methods in subtypes to provide more specific functionality. DB2® executes the most specific method at run time, depending on the dynamic type of the subject. This is called dynamic dispatch.

Windows management instrumentation

The DB2® WMI provider allows WMI applications to monitor DB2 server services, enumerate and create databases, configure operational settings and perform database backup, restore, and roll-forward operations.

Related concepts:

- "Windows Management Instrumentation (WMI)" in the *Application Development Guide: Building and Running Applications*

Related reference:

- "Windows Management Instrumentation Samples" in the *Application Development Guide: Building and Running Applications*

New Call Level Interface (CLI) LOAD functionality

In DB2® Version 8, you can import large amounts of data in DB2 CLI using the CLI LOAD functionality.

CLI LOAD provides an interface to the IBM® DB2 LOAD utility from CLI. This functionality allows you to insert data in CLI using LOAD instead of array insert. This option can yield significant performance benefits when large amounts of data need to be inserted.

Related tasks:

- "Importing Data with the CLI LOAD Utility in CLI Applications" in the *CLI Guide and Reference, Volume 1*

Declared global temporary table enhancements

Enhancements to declared global temporary tables (DGTTs) include:

- Index support: the ability to create indexes using the CREATE INDEX statement.
- Undo logging, to support the rollback of data changes to DGTTs.
- Statistics support: improved performance as a result of using the RUNSTATS command to update statistics about the physical characteristics of a temporary table and its associated indexes.

Related concepts:

- "RUNSTATS command enhancements" on page 17

IBM OLE DB Provider for DB2

The IBM® OLE DB Provider for DB2® allows DB2 to act as a resource manager for the OLE DB provider. This support gives OLE DB-based applications the ability to extract or query DB2 data using the OLE interface.

The IBM OLE DB Provider for DB2 offers the following features:

- Support level 0 of the OLE DB provider specification, including some additional level 1 interfaces.
- A free threaded provider implementation, which enables the application to create components in one thread and use those components in any other thread.
- An error lookup service that returns DB2 error messages.

If DB2 Connect™ is installed, OLE DB users can also access data on host database management systems such as DB2 for MVS, DB2 for VSE & VM, and SQL/400.

Related concepts:

- "Purpose of the IBM OLE DB Provider for DB2" in the *Application Development Guide: Programming Client Applications*
- "Application Types Supported by the IBM OLE DB Provider for DB2" in the *Application Development Guide: Programming Client Applications*

Related tasks:

- "Building ADO Applications with Visual Basic" in the *Application Development Guide: Building and Running Applications*
- "Building ADO Applications with Visual C++" in the *Application Development Guide: Building and Running Applications*

Related reference:

- "IBM OLE DB Provider Support for OLE DB Components and Interfaces" in the *Application Development Guide: Programming Client Applications*
- "IBM OLE DB Provider Support for ADO Methods and Properties" in the *Application Development Guide: Programming Client Applications*

Web services applications

Web services are a powerful new programming paradigm that enable the development and deployment of loosely coupled applications within a company or across industries. Web services are based on emerging technologies including Simple Object Access Protocol (SOAP), Web Services Description Language (WSDL), and Universal Description, Discovery, and Integration (UDDI).

A Web service is created by wrapping an application so that it can be accessed using XML messages which, in turn, are wrapped to mask the underlying transport protocol. The service is publicized in a standard-format registry, allowing people and applications to find and use the service over the web.

DB2® can be accessed as a Web service provider, and it is usually teamed with IBM® WebSphere® family products to provide a complete Web services framework.

An easy way to develop data-intensive Web services applications is through DB2's document access definition extension (DADx) programming model. DADx files are simply XML documents that contain stored procedure references, DB2 XML Extender references, or traditional SQL. Best of all, DADx files can be created quickly without any Java™ programming.

For a demonstration of DB2 as a Web service provider, see the IBM Video Central for e-business tutorial at

`http://www.ibm.com/software/data/developer/samples/video/`

Related concepts:

- "Web Services" in the *Application Development Guide: Programming Client Applications*
- "Web Services Architecture" in the *Application Development Guide: Programming Client Applications*
- "Tools for Building Web Applications" in the *Application Development Guide: Programming Client Applications*

JDBC driver enhancements

Many new features and enhancements have been made to the JDBC drivers in DB2® Version 8. Among these changes, the biggest change is the architectural improvement that shortens the code path between the JDBC driver and DB2 servers, which results in better performance and stability for JDBC and SQLj, and any applications built with the DB2 JDBC drivers.

The JDBC drivers in previous releases were built on top of DB2 CLI, a DB2-native C call level interface, which was built on top of several other layers. Version 8 adds a new client layer that communicates with the DB2 server via the DRDA® (Distributed Relational Database Architecture) protocol, and that replaces CLI and a number of layers beneath it. The new JDBC type 4 driver is built on top of this Java™ client layer.

The following is new for JDBC for Version 8:

- Type 4 JDBC driver

  The new type 4 JDBC driver is a two-tier pure Java JDBC driver, which allows a Java client to communicate directly with DB2 servers via DRDA protocol. This driver is designed to replace the type 3 driver. You should migrate applets that use the type 3 JDBC driver to the type 4 driver, in preparation for the end of type 3 driver support.

DB2 also provides a new profile customizer (that is, a new version of `db2profc`) that supports new JDBC 2.0 features, including LOBs and Java expressions as host variables. Designed to support the new type 4 driver, the new customizer also works with the type 2 and type 3 drivers.

- Type 2 and type 3 driver enhancements:
  - The type 2 JDBC driver is Java 2 Enterprise Edition (J2EE) certified, which means it conforms to the Java 2 platform Enterprise Edition specification, and is certified for use with J2EE-compatible products.
  - The JDBC drivers are available on 64-bit platforms.
  - Unicode support is enhanced. All unnecessary code page conversion between a DB2 client and a DB2 server that occurs when connecting to a Unicode database is eliminated. No data loss occurs.
  - OUTPUT LOB parameter support for `CallableStatement` is added.
  - Better memory management, which in turn improves the stability and performance of the drivers.
  - `FetchSize` for `ResultSets` is supported by the type 3 JDBC driver.
  - Support for Datalink as a data type, and Date, Time, and Timestamp values that can be set or retrieved for a particular time zone using a Calendar object.
  - Improved JDBC tracing.
  - Improved performance for `PreparedStatements`.
- The JDBC 1.2 drivers are deprecated.
- The new SQLj translator supports the `java.sql.Blob` and `java.sql.Clob` types of JDBC 2.0, as well as host variable expressions.

Related concepts:

- "Migrating Applications" in the *Application Development Guide: Building and Running Applications*
- "Application and Applet Support in Java with the Type 4 Driver" in the *Application Development Guide: Programming Client Applications*
- "JDBC 2.1 Core API Restrictions by the DB2 JDBC Type 4 Driver" in the *Application Development Guide: Programming Client Applications*
- "JDBC 2.1 Optional Package API Support by the DB2 JDBC Type 4 Driver" in the *Application Development Guide: Programming Client Applications*

WebSphere Studio and WebSphere Application Server available with DB2 Universal Developer's Edition

WebSphere® Studio and WebSphere Application Server are both available with DB2® Universal Developer's Edition. WebSphere Studio is a suite of tools that brings all aspects of Web site developent into a common interface. WebSphere Studio makes it easier than ever to cooperatively create, assemble, publish, and maintain dynamic interactive Web applications.

You can write enterprise Java™ applications that use the WebSphere Application Server to access data in DB2 databases. When your data source is created with the WebSphere Application Server, your application can use the connection pooling that is provided by the WebSphere Application Server. The connection pooling, which conforms to the JDBC core and Optional Package API specifications, allows administrators to tune the pool for optimal performance, and allows programmers to write applications without knowledge of common vendor-specific SQLExceptions.

WebSphere also provides support to encrypt and decrypt data, thus helping to ensure security for your applications.

Related concepts:

- "WebSphere Studio" in the *Application Development Guide: Programming Client Applications*
- "DB2 Developer's Edition Products" in the *Application Development Guide: Programming Client Applications*

DB2 XML support enhancements

DB2® support for XML has been enhanced with the following:

- XML schema validation user-defined function (UDF)

  You can use the `schemavalidate` UDF to validate XML documents.
- REC2XML and COLLATTVAL functions

  You can use the REC2XML function to return a string formatted with XML tags and containing column names and column data. The COLLATTVAL function can be used to return a string with the column name as an attribute value.

Chapter 11. Federated Systems

Federated Systems

A DB2® *federated system* is a special type of distributed database management system (DBMS). A federated system consists of a DB2 instance that operates as a server, a database that serves as the federated database, one or more data sources, and clients (users and applications) who access the database and data sources. With a federated system you can send distributed requests to multiple data sources within a single SQL statement. The power of a DB2 federated system is in its ability to:

- Join data from local tables and remote data sources, as if all the data are local.
- Take advantage of the data source processing strengths, by sending distributed requests to the data sources for processing.
- Compensate for SQL limitations at the data source by processing parts of a distributed request at the federated server.

Federated database systems provide the middleware functionality for outstanding information integration. Built into DB2 Enterprise Server Edition is the ability to federate relational data across IBM's family of databases, including DB2 and Informix™ IDS.

Here is a brief summary of the enhancements in Version 8 to Federated Systems:

- Support on additional operating systems — DB2 for Linux, DB2 for HP-UX, and Windows® 2000
- Write capability to perform INSERT, UPDATE, and DELETE actions on the data sources
- Ability to create remote tables on relational data sources.

Related concepts:

- "Federated systems" in the *Federated Systems Guide*
- "The federated database" in the *Federated Systems Guide*
- "Overview of the tasks to set up a federated system" in the *Federated Systems Guide*

Chapter 12. Business intelligence enhancements

DB2 OLAP Server and DB2 OLAP Integration Server changes

You can use DB2® OLAP Server and its add-on features, such as DB2 OLAP Integration Server and DB2 OLAP Server™ Analyzer, to build online analytical processing (OLAP) applications that are production-ready and Web-ready. These products are sold separately.

Note: The DB2 OLAP Starter Kit is not available in DB2 Universal Database™ Version 8.

DB2 OLAP Server for Version 8 includes the following major improvements:

- DB2 OLAP Server Miner is a no-cost add-on feature of DB2 OLAP Server that automatically mines large volumes of OLAP data. It identifies and reports on specific elements that represent the most outstanding, and unexpected, pieces of information. You can expose problems and opportunities that were previously hidden and act on them quickly.
- Hybrid analysis is a new function of DB2 OLAP Integration Server that you can use to access more data without enlarging your OLAP database. It builds a virtual extension from an OLAP database to the relational database that contains the lowest members of your OLAP hierarchies. The data you access with hybrid analysis is not part of the OLAP database; hybrid analysis maps the relational data to the appropriate hierarchies in your OLAP database.

The Web site for DB2 OLAP Server describes the other improvements for Version 8:

`http://www.ibm.com/software/data/db2/db2olap/`

DB2 Spatial Extender enhancements

With DB2® Spatial Extender, you can include spatial attributes, such as how far your customers live from your offices, 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.

New enhancements in DB2 Spatial Extender Version 8 include:

- An Index Advisor to help you design efficient spatial indexes more quickly.

- More spatial functions that let you derive a broader scope of information from your spatial data.
- A greater ability than with previous versions to integrate user- and vendor-supplied geocoders into your spatial environment.
- More sample code that you can cut and paste into your applications.
- A greater ability to export spatial data to geobrowsers than in previous versions. Geobrowsers let you see results of queries in visual form.
- The ability to export SDE transfer data to files that can be loaded into external data sources.

Related concepts:

- "The purpose of DB2 Spatial Extender" in the *Spatial Extender User's Guide and Reference*

Chapter 13. DB2 family enhancements

Multiplatform tools for DB2 Universal Database

In September of 2001, IBM® delivered application tools designed to enhance DB2® Universal Database across the Microsoft® Windows, HP-UX, Solaris Operating Environment, AIX, and Linux platforms with the introduction of DB2 Web Query Tool for Multiplatforms, Version 1.2 and DB2 Table Editor for Multiplatforms, Version 4.2. Enhanced releases of these products were recently announced.

Now, IBM adds three new significant tools to its multiplatform portfolio. The new and enhanced capabilities extending the DB2 solution include:

DB2 Recovery Expert for Multiplatforms, Version 1

DB2 Recovery Expert for Multiplatforms provides targeted, flexible, and automated recovery of database assets, even as systems remain on-line. The tool provides an easy-to-use environment that even less experienced DBAs can successfully use to complete highly sophisticated and efficient recovery techniques in minimal time. Built-in *self-managing and resource tuning (SMART)* features provide intelligent analysis of altered, corrupted, incorrect, or missing database assets – including tables, indexes, or data – and automate the process of rebuilding those assets to a correct point in time, all without disruption to normal database or business operations. With DB2 Recovery Expert, users can

- Recover database objects without resorting to resource-intensive disaster recovery
- Accurately roll back unwanted data changes throughout a database
- Obtain intelligent assistance when determining the most efficient technique for a recovery situation
- Recreate databases (or selected objects) to new environments
- Support multiple databases on multiple operating systems and hardware platforms

The DB2 Recovery Expert for Multiplatforms supports DB2 Universal Database Version 7 and later.

DB2 Performance Expert for Multiplatforms, Version 1

DB2 Performance Expert for Multiplatforms provides a comprehensive view that consolidates, reports, analyzes, and recommends changes on

DB2 performance-related information. DB2 Performance Expert can selectively employ and integrate the view from all trace, snapshot, event, and DB2 Version 8 Health Monitor outputs. Plus, it provides on-line snapshot reports, a buffer pool analyzer and reporting facility, and can selectively store performance data in its own performance data warehouse which you can study at a later time, both at detailed and rolled up levels of the data. And with its starter set of SMART features, DB2 Performance Expert also provides recommendations for system tuning to gain optimum throughput.

DB2 Performance Expert provides a state-of-the-art workstation-based user interface. Because the DB2 Performance Expert provides a common interface across DB2 server platforms, it both simplifies DBA tasks, and reduces the requirements for cross-platform training. The DB2 Performance Expert can be started stand-alone or launched via a plug-in from the DB2 Control Center.

The DB2 Performance Expert for Multiplatforms supports DB2 Universal Database™ Version 8.

DB2 High Performance Unload for Multiplatforms, Version 2.1

DB2 High Performance Unload gives customers a fast and efficient tool for unloading and extracting data for movement across enterprise systems, or for reorganization in-place. The product delivers high levels of parallelism when either unloading or extracting in partitioned database environments, both for DB2 Enterprise Server Edition Version 8, and DB2 Enterprise - Extended Edition (EEE) Version 7. It can process multiple select statements at the same time so that multiple requests of the DB2 UDB will be processed in a single scan of the table. DB2 High Performance Unload also has the capability to:

- Rapidly unload table spaces of all types supported by DB2
- Make the output file-set a tape, disk or named pipe, formatted suitably for optional reorganization or reload
- Simultaneously execute several unloads accessing the same table space
- Do unload against a full backup to eliminate interference with DB2 production databases (the image copy can be the last or any full copy)
- Unload selected rows and columns via a SELECT
- Unload every *n* rows (for example, every fifth row)
- Generate load control statements for subsequent reload
- Optionally unload data and move information from a partitioned table to a single file-set for use in another database environment

- Unload DB2 Universal Database table data from multiple database partitions in a single request

The DB2 High Performance Unload for Multiplatforms supports DB2 Universal Database Version 7 and later.

DB2 Table Editor for Multiplatforms, Version 4.3

DB2 Table Editor makes it possible to offer direct DB2 database access to anyone for creating, reviewing, or updating data. Task-specific forms that are restricted to specific data and actions can be quickly built and rolled out to novice end users, and an expert interface can be used for browsing existing databases and ad hoc actions. A Java-enabled browser is all that is needed on the client machine for using DB2 Table Editor solutions.

The DB2 Table Editor for Multiplatforms supports DB2 Universal Database Version 7 and later.

DB2 Web Query Tool for Multiplatforms, Version 1.3

When any part of your business operations require faster, more creative, or more complex data operations, DB2 Web Query Tool is the ideal authoring environment. DB2 Web Query Tool's advanced SQL functionality extracts the full performance out of DB2 databases while making it easy for users at all levels to access enterprise-wide data. Wherever data access and data integration can bring value, DB2 Web Query Tool can deliver that value.

The DB2 Web Query Tool for Multiplatforms supports DB2 Universal Database Version 7 and later.

DB2 Data Links Manager enhancements

The DB2® Data Links Manager is a separately orderable DB2 feature that enables your applications to manipulate data that resides in both unstructured files and in the relational database management system (RDBMS). The files stored outside of the database reside in file systems accessible over a network, and are managed as if they were stored in the RDBMS. DB2 Data Links Manager provides the integration between the relational database management system (RDBMS) and the external file systems through extensions to DB2 Universal Database.

Version 8 includes the following enhancements:

- DB2 Data Links Manager is now available on the Windows® 2000 operating system, in addition to Windows NT, AIX, and the Solaris Operating Environment.

- DB2 Data Links Manager now supports both the Version 8 and Version 7 Solaris Operating Environments.
- You can now update files while they remain under the control of DB2 Data Links Manager. Files that DB2 Data Links Manager controls are referred to as *linked files*. Before DB2 Version 8, you had to unlink a file, make changes to the file, then relink the file.
- Performance improvements for archive, data recovery, and replication operations on linked files. These improvements reduce the time for beginning backups involving linked files, reduce the time for running the reconcile utility, and improve the throughput on replication of linked files using DB2 replication.
- Improved security features for restricting the linking of files to authorized users.
- A new manual, the *DB2 Data Links Manager Administration Guide and Reference*, is now available. This new manual describes how to administer, use, write applications for, and troubleshoot the DB2 Data Links Manager.

Related concepts:

- "DB2 Data Links Manager" in the *Quick Beginnings for Data Links Manager*
- "DB2 Data Links Manager replication enhancements" on page 55
- "Introduction to Data Links Manager security" in the *DB2 Data Links Manager Administration Guide and Reference*

Related reference:

- "DB2 Data Links Manager system setup and backup recommendations" in the *DB2 Data Links Manager Administration Guide and Reference*

DB2 XML Extender enhancements

The DB2® XML Extender is an end-to-end solution for storing and retrieving XML documents.

The DB2 XML Extender makes it easy to work with XML documents by enabling you to store these structured documents in DB2 databases. The XML Extender also provides you with user-defined functions that assist you in working with XML documents. Entire XML documents can be stored in DB2 databases as character data or stored as external files that are still managed by DB2. Retrieval functions allow you to retrieve either the entire XML document or individual elements or attributes.

For e-commerce, you can use the XML Extender to leverage your critical business information in DB2 databases to engage in business-to-business solutions using XML-based interchange formats.

New features have been added to enhance the XML extender:

- The XML Extender now supports Web services with the Web services Object Runtime Framework (WORF), a set of tools for implementing Web services with DB2. Web services are XML-based application functions that can be started from the Internet.
- MQSeries® applications are also supported by the XML Extender. This feature allows users to send XML documents to, or retrieve documents from, MQSeries message queues.

DB2 Net Search Extender

Textual data is a large and valuable source of unstructured information in any enterprise. Adding smart text search capability to your database allows you to make even more informed business decisions according to your specific needs.

With Version 8 of DB2 Net Search Extender, IBM has merged the in-memory search capabilities of the Version 7.2 DB2 Net Search Extender and the powerful text search capabilities of the Version 7.2 DB2 Text Information Extender into a single, easy-to-use DB2 option. Choose from among word, phrase, boolean, fuzzy, wildcard, proximity, and free-text searches, or use the thesaurus support to search within sections of structured documents, including XML and HTML documents. In-memory search is indispensable for e-commerce, and for any application with high performance and scalability needs.

The merging of these two extenders into a single product represents an integrated solution that follows industry standards by basing the text search interface on the SQL Multimedia (SQL/MM) standard. Integrating text search with DB2's optimizer yields high performance full text search working within SQL queries.

The integrated DB2 Net Search Extender solution can be managed right from the DB2 Universal Database Control Center, significantly simplifying basic tasks like creating a text index. You can even use DB2 Net Search Extender with DB2's built-in federated support to index and search your text data stored in other DB2 and Informix Dynamic Server (IDS) databases.

Host and iSeries applications can use two-phase commit with TCP/IP

With Version 8, host and iSeries™ applications can use two-phase commit when connecting to DB2® Universal Database over TCP/IP.

Appendix A. DB2 Universal Database technical information

Overview of DB2 Universal Database technical information

DB2 Universal Database technical information can be obtained in the following formats:

- Books (PDF and hard-copy formats)
- A topic tree (HTML format)
- Help for DB2 tools (HTML format)
- Sample programs (HTML format)
- Command line help
- Tutorials

This section is an overview of the technical information that is provided and how you can access it.

FixPaks for DB2 documentation

IBM may periodically make documentation FixPaks available. Documentation FixPaks allow you to update the information that you installed from the *DB2 HTML Documentation CD* as new information becomes available.

Note: If you do install documentation FixPaks, your HTML documentation will contain more recent information than either the DB2 printed or online PDF manuals.

Categories of DB2 technical information

The DB2 technical information is categorized by the following headings:

- Core DB2 information
- Administration information
- Application development information
- Business intelligence information
- DB2 Connect information
- Getting started information
- Tutorial information
- Optional component information
- Release notes

The following tables describe, for each book in the DB2 library, the information needed to order the hard copy, print or view the PDF, or locate the HTML directory for that book. A full description of each of the books in

the DB2 library is available from the IBM Publications Center at www.ibm.com/shop/publications/order

The installation directory for the HTML documentation CD differs for each category of information:

htmlcdpath/doc/htmlcd/*%L*/*category*

where:

- *htmlcdpath* is the directory where the HTML CD is installed.
- *%L* is the language identifier. For example, en_US.
- *category* is the category identifier. For example, core for the core DB2 information.

In the PDF file name column in the following tables, the character in the sixth position of the file name indicates the language version of a book. For example, the file name db2d1e80 identifies the English version of the *Administration Guide: Planning* and the file name db2d1g80 identifies the German 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
Arabic	w
Brazilian Portuguese	b
Bulgarian	u
Croatian	9
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
Romanian	8
Russian	r
Simp. Chinese	c
Slovakian	7
Slovenian	l

Spanish	z
Swedish	s
Trad. Chinese	t
Turkish	m

No form number indicates that the book is only available online and does not have a printed version.

Core DB2 information

The information in this category covers DB2 topics that are fundamental to all DB2 users. You will find the information in this category useful whether you are a programmer, a database administrator, or you work with DB2 Connect, DB2 Warehouse Manager, or other DB2 products.

The installation directory for this category is doc/htmlcd/%L/core.

Table 1. Core DB2 information

Name	Form Number	PDF File Name
IBM DB2 Universal Database Command Reference	SC09-4828	db2n0x80
IBM DB2 Universal Database Glossary	No form number	db2t0x80
IBM DB2 Universal Database Master Index	SC09-4839	db2w0x80
IBM DB2 Universal Database Message Reference, Volume 1	GC09-4840	db2m1x80
IBM DB2 Universal Database Message Reference, Volume 2	GC09-4841	db2m2x80
IBM DB2 Universal Database What's New	SC09-4848	db2q0x80

Administration information

The information in this category covers those topics required to effectively design, implement, and maintain DB2 databases, data warehouses, and federated systems.

The installation directory for this category is doc/htmlcd/%L/admin.

Table 2. Administration information

Name	Form number	PDF file name
IBM DB2 Universal Database Administration Guide: Planning	SC09-4822	db2d1x80
IBM DB2 Universal Database Administration Guide: Implementation	SC09-4820	db2d2x80
IBM DB2 Universal Database Administration Guide: Performance	SC09-4821	db2d3x80
IBM DB2 Universal Database Administrative API Reference	SC09-4824	db2b0x80
IBM DB2 Universal Database Data Movement Utilities Guide and Reference	SC09-4830	db2dmx80
IBM DB2 Universal Database Data Recovery and High Availability Guide and Reference	SC09-4831	db2hax80
IBM DB2 Universal Database Data Warehouse Center Administration Guide	SC27-1123	db2ddx80
IBM DB2 Universal Database Federated Systems Guide	GC27-1224	db2fpx80
IBM DB2 Universal Database Guide to GUI Tools for Administration and Development	SC09-4851	db2atx80
IBM DB2 Universal Database Replication Guide and Reference	SC27-1121	db2e0x80
IBM DB2 Installing and Administering a Satellite Environment	GC09-4823	db2dsx80
IBM DB2 Universal Database SQL Reference, Volume 1	SC09-4844	db2s1x80
IBM DB2 Universal Database SQL Reference, Volume 2	SC09-4845	db2s2x80
IBM DB2 Universal Database System Monitor Guide and Reference	SC09-4847	db2f0x80

Application development information

The information in this category is of special interest to application developers or programmers working with DB2. You will find information about supported languages and compilers, as well as the documentation required to access DB2 using the various supported programming interfaces, such as embedded SQL, ODBC, JDBC, SQLj, and CLI. If you view this information online in HTML you can also access a set of DB2 sample programs in HTML.

The installation directory for this category is doc/htmlcd/%L/ad.

Table 3. Application development information

Name	Form number	PDF file name
IBM DB2 Universal Database Application Development Guide: Building and Running Applications	SC09-4825	db2axx80
IBM DB2 Universal Database Application Development Guide: Programming Client Applications	SC09-4826	db2a1x80
IBM DB2 Universal Database Application Development Guide: Programming Server Applications	SC09-4827	db2a2x80
IBM DB2 Universal Database Call Level Interface Guide and Reference, Volume 1	SC09-4849	db2l1x80
IBM DB2 Universal Database Call Level Interface Guide and Reference, Volume 2	SC09-4850	db2l2x80
IBM DB2 Universal Database Data Warehouse Center Application Integration Guide	SC27-1124	db2adx80
IBM DB2 XML Extender Administration and Programming	SC27-1234	db2sxx80

Business intelligence information

The information in this category describes how to use components that enhance the data warehousing and analytical capabilities of DB2 Universal Database.

The installation directory for this category is doc/htmlcd/%L/wareh.

Table 4. Business intelligence information

Name	Form number	PDF file name
IBM DB2 Warehouse Manager Information Catalog Center Administration Guide	SC27-1125	db2dix80
IBM DB2 Warehouse Manager Installation Guide	GC27-1122	db2idx80

DB2 Connect information

The information in this category describes how to access host or iSeries data using DB2 Connect Enterprise Edition or DB2 Connect Personal Edition.

The installation directory for this category is doc/htmlcd/%L/conn.

Table 5. DB2 Connect information

Name	Form number	PDF file name
APPC, CPI-C, and SNA Sense Codes	No form number	db2apx80
IBM Connectivity Supplement	No form number	db2h1x80
IBM DB2 Connect Quick Beginnings for DB2 Connect Enterprise Edition	GC09-4833	db2c6x80
IBM DB2 Connect Quick Beginnings for DB2 Connect Personal Edition	GC09-4834	db2c1x80
IBM DB2 Connect User's Guide	SC09-4835	db2c0x80

Getting started information

The information in this category is useful when you are installing and configuring servers, clients, and other DB2 products.

The installation directory for this category is doc/htmlcd/%L/start.

Table 6. Getting started information

Name	Form number	PDF file name
IBM DB2 Universal Database Quick Beginnings for DB2 Clients	GC09-4832	db2itx80

Table 6. Getting started information (continued)

Name	Form number	PDF file name
IBM DB2 Universal Database Quick Beginnings for DB2 Servers	GC09-4836	db2isx80
IBM DB2 Universal Database Quick Beginnings for DB2 Personal Edition	GC09-4838	db2i1x80
IBM DB2 Universal Database Installation and Configuration Supplement	GC09-4837	db2iyx80
IBM DB2 Universal Database Quick Beginnings for DB2 Data Links Manager	GC09-4829	db2z6x80

Tutorial information

Tutorial information introduces DB2 features and teaches how to perform various tasks.

The installation directory for this category is doc/htmlcd/%L/tutr.

Table 7. Tutorial information

Name	Form number	PDF file name
Business Intelligence Tutorial: Introduction to the Data Warehouse	No form number	db2tux80
Business Intelligence Tutorial: Extended Lessons in Data Warehousing	No form number	db2tax80
Development Center Tutorial for Video Online using Microsoft Visual Basic	No form number	db2tdx80
Information Catalog Center Tutorial	No form number	db2aix80
Video Central for e-business Tutorial	No form number	db2twx80
Visual Explain Tutorial	No form number	db2tvx80

Optional component information

The information in this category describes how to work with optional DB2 components.

The installation directory for this category is doc/htmlcd/%L/opt.

Table 8. Optional component information

Name	Form number	PDF file name
IBM DB2 Life Sciences Data Connect Planning, Installation, and Configuration Guide	GC27-1235	db2lsx80
IBM DB2 Spatial Extender User's Guide and Reference	SC27-1226	db2sbx80
IBM DB2 Universal Database Data Links Manager Administration Guide and Reference	SC27-1221	db2z0x80
IBM DB2 Universal Database Net Search Extender Administration and Programming Guide **Note:** HTML for this document is not installed from the HTML documentation CD.	SH12-6740	N/A

Release notes

The release notes provide additional information specific to your product's release and FixPak level. They also provides summaries of the documentation updates incorporated in each release and FixPak.

Table 9. Release notes

Name	Form number	PDF file name
DB2 Release Notes	See note.	See note.
DB2 Installation Notes	Available on product CD-ROM only.	Available on product CD-ROM only.

Note: The HTML version of the release notes 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/opt/db2_08_01 on AIX
- /opt/IBM/db2/V8.1 on all other UNIX operating systems

Related tasks:

- "Printing DB2 books from PDF files" on page 93

- "Ordering printed DB2 books" on page 94
- "Accessing online help" on page 94
- "Finding product information by accessing the DB2 Information Center from the administration tools" on page 98
- "Viewing technical documentation online directly from the DB2 HTML Documentation CD" on page 100

Printing DB2 books from PDF files

You can print DB2 books from the PDF files on the *DB2 PDF Documentation* CD. Using Adobe Acrobat Reader, you can print either the entire book or a specific range of pages.

Prerequisites:

Ensure that you have Adobe Acrobat Reader. It is available from the Adobe Web site at www.adobe.com

Procedure:

To print a DB2 book from a PDF file:

1. Insert the *DB2 PDF Documentation* CD. On UNIX operating systems, mount the DB2 PDF Documentation CD. Refer to your *Quick Beginnings* book for details on how to mount a CD on UNIX operating systems.
2. Start Adobe Acrobat Reader.
3. Open the PDF file from one of the following locations:
   - On Windows operating systems:

     x:\doc*language* directory, where *x* represents the CD-ROM drive letter and *language* represents the two-character territory code that represents your language (for example, EN for English).
   - On UNIX operating systems:

     /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.

Related tasks:

- "Ordering printed DB2 books" on page 94
- "Finding product information by accessing the DB2 Information Center from the administration tools" on page 98
- "Viewing technical documentation online directly from the DB2 HTML Documentation CD" on page 100

Related reference:

- "Overview of DB2 Universal Database technical information" on page 85

Ordering printed DB2 books

Procedure:

To order printed books:

- Contact your IBM authorized dealer or marketing representative. To find a local IBM representative, check the IBM Worldwide Directory of Contacts at www.ibm.com/planetwide
- Phone `1-800-879-2755` in the United States or `1-800-IBM-4YOU` in Canada.
- Visit the IBM Publications Center at www.ibm.com/shop/publications/order

You can also obtain printed DB2 manuals by ordering Doc Packs for your DB2 product from your IBM Reseller. The Doc Packs are subsets of the manuals in the DB2 library selected to help you to get started using the DB2 product that you purchased. The manuals in the Doc Packs are the same as those that are available in PDF format on the *DB2 PDF Documentation CD* and contain the same content as the documentation that is available on the *DB2 HTML Documentation CD*.

Related tasks:

- "Printing DB2 books from PDF files" on page 93
- "Finding topics by accessing the DB2 Information Center from a browser" on page 96
- "Viewing technical documentation online directly from the DB2 HTML Documentation CD" on page 100

Related reference:

- "Overview of DB2 Universal Database technical information" on page 85

Accessing online help

The online help that comes with all DB2 components is available in three types:

- Window and notebook help
- Command line help
- SQL statement help

Window and notebook help explain the tasks that you can perform in a window or notebook and describe the controls. This help has two types:

- Help accessible from the **Help** button
- Infopops

The **Help** button gives you access to overview and prerequisite information. The infopops describe the controls in the window or notebook. Window and notebook help are available from DB2 centers and components that have user interfaces.

Command line help includes Command help and Message help. Command help explains the syntax of commands in the command line processor. Message help describes the cause of an error message and describes any action you should take in response to the error.

SQL statement help includes SQL help and SQLSTATE help. DB2 returns an SQLSTATE value for conditions that could be the result of an SQL statement. SQLSTATE help explains the syntax of SQL statements (SQL states and class codes).

Note: SQL help is not available for UNIX operating systems.

Procedure:

To access online help:

- For window and notebook help, click **Help** or click that control, then click **F1**. If the **Automatically display infopops** check box on the **General** page of the **Tool Settings** notebook is selected, you can also see the infopop for a particular control by holding the mouse cursor over the control.
- For command line help, open the command line processor and enter:
  - For Command help:

    ```
    ? command
    ```

    where *command* represents a keyword or the entire command.

    For example, `? catalog` displays help for all the CATALOG commands, while `? catalog database` displays help for the CATALOG DATABASE command.
- For Message help:

  ```
  ? XXXnnnnn
  ```

  where *XXXnnnnn* represents a valid message identifier.

  For example, `? SQL30081` displays help about the SQL30081 message.
- For SQL statement help, open the command line processor and enter:

  ```
  ? sqlstate or ? class code
  ```

where *sqlstate* represents a valid five-digit SQL state and *class code* represents the first two digits of the SQL state.

For example, `? 08003` displays help for the 08003 SQL state, while `? 08` displays help for the 08 class code.

Related tasks:

- "Finding topics by accessing the DB2 Information Center from a browser" on page 96
- "Viewing technical documentation online directly from the DB2 HTML Documentation CD" on page 100

Finding topics by accessing the DB2 Information Center from a browser

The DB2 Information Center accessed from a browser enables you to access the information you need to take full advantage of DB2 Universal Database and DB2 Connect. The DB2 Information Center also documents major DB2 features and components including replication, data warehousing, metadata, and DB2 extenders.

The DB2 Information Center accessed from a browser is composed of the following major elements:

Navigation tree
The navigation tree is located in the left frame of the browser window. The tree expands and collapses to show and hide topics, the glossary, and the master index in the DB2 Information Center.

Navigation toolbar
The navigation toolbar is located in the top right frame of the browser window. The navigation toolbar contains buttons that enable you to search the DB2 Information Center, hide the navigation tree, and find the currently displayed topic in the navigation tree.

Content frame
The content frame is located in the bottom right frame of the browser window. The content frame displays topics from the DB2 Information Center when you click on a link in the navigation tree, click on a search result, or follow a link from another topic or from the master index.

Prerequisites:

To access the DB2 Information Center from a browser, you must use one of the following browsers:

- Microsoft Explorer, version 5 or later

- Netscape Navigator, version 6.1 or later

Restrictions:

The DB2 Information Center contains only those sets of topics that you chose to install from the *DB2 HTML Documentation CD*. If your Web browser returns a `File not found` error when you try to follow a link to a topic, you must install one or more additional sets of topics from the *DB2 HTML Documentation CD*.

Procedure:

To find a topic by searching with keywords:

1. In the navigation toolbar, click **Search**.
2. In the top text entry field of the Search window, enter one or more terms related to your area of interest and click **Search**. A list of topics ranked by accuracy displays in the **Results** field. The numerical ranking beside the hit provides an indication of the strength of the match (bigger numbers indicate stronger matches).

   Entering more terms increases the precision of your query while reducing the number of topics returned from your query.
3. In the **Results** field, click the title of the topic you want to read. The topic displays in the content frame.

To find a topic in the navigation tree:

1. In the navigation tree, click the book icon of the category of topics related to your area of interest. A list of subcategories displays underneath the icon.
2. Continue to click the book icons until you find the category containing the topics in which you are interested. Categories that link to topics display the category title as an underscored link when you move the cursor over the category title. The navigation tree identifies topics with a page icon.
3. Click the topic link. The topic displays in the content frame.

To find a topic or term in the master index:

1. In the navigation tree, click the "Index" category. The category expands to display a list of links arranged in alphabetical order in the navigation tree.
2. In the navigation tree, click the link corresponding to the first character of the term relating to the topic in which you are interested. A list of terms with that initial character displays in the content frame. Terms that have multiple index entries are identified by a book icon.

3. Click the book icon corresponding to the term in which you are interested. A list of subterms and topics displays below the term you clicked. Topics are identified by page icons with an underscored title.
4. Click on the title of the topic that meets your needs. The topic displays in the content frame.

Related concepts:

- "Accessibility" on page 105
- "DB2 Information Center accessed from a browser" on page 108

Related tasks:

- "Finding product information by accessing the DB2 Information Center from the administration tools" on page 98
- "Updating the HTML documentation installed on your machine" on page 100
- "Troubleshooting DB2 documentation search with Netscape 4.x" on page 103
- "Searching the DB2 documentation" on page 104

Related reference:

- "Overview of DB2 Universal Database technical information" on page 85

Finding product information by accessing the DB2 Information Center from the administration tools

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

The DB2 Information Center accessed from the tools provides six types of information.

Tasks Key tasks you can perform using DB2.

Concepts
Key concepts for DB2.

Reference
DB2 reference information, such as keywords, commands, and APIs.

Troubleshooting
Error messages and information to help you with common DB2 problems.

Samples
Links to HTML listings of the sample programs provided with DB2.

Tutorials

Instructional aid designed to help you learn a DB2 feature.

Prerequisites:

Some links in the DB2 Information Center point to Web sites on the Internet. To display the content for these links, you will first have to connect to the Internet.

Procedure:

To find product information by accessing the DB2 Information Center from the tools:

1. Start the DB2 Information Center in one of the following ways:
   - From the graphical administration tools, click on the **Information Center** icon in the toolbar. You can also select it from the **Help** menu.
   - At the command line, enter **db2ic**.
2. Click the tab of the information type related to the information you are attempting to find.
3. Navigate through the tree and click on the topic in which you are interested. The Information Center will then launch a Web browser to display the information.
4. To find information without browsing the lists, click the **Search** icon to the right of the list.

   Once the Information Center has launched a browser to display the information, you can perform a full-text search by clicking the **Search** icon in the navigation toolbar.

Related concepts:

- "Accessibility" on page 105
- "DB2 Information Center accessed from a browser" on page 108

Related tasks:

- "Finding topics by accessing the DB2 Information Center from a browser" on page 96
- "Searching the DB2 documentation" on page 104

Viewing technical documentation online directly from the DB2 HTML Documentation CD

All of the HTML topics that you can install from the *DB2 HTML Documentation CD* can also be read directly from the CD. Therefore, you can view the documentation without having to install it.

Restrictions:

As the Tools help is installed from the DB2 product CD and not from the *DB2 HTML Documentation CD*, you must install the DB2 product to view the help.

Procedure:

1. Insert the *DB2 HTML Documentation* CD. On UNIX operating systems, mount the *DB2 HTML Documentation CD*. Refer to your *Quick Beginnings* book for details on how to mount a CD on UNIX operating systems.
2. Start your HTML browser and open the appropriate file:
   - For Windows operating systems:

     ```
     e:\program files\IBM\SQLLIB\doc\htmlcd\%L\index.htm
     ```

     where *e* represents the CD-ROM drive, and %L is the locale of the documentation that you wish to use, for example, **en_US** for English.
   - For UNIX operating systems:

     ```
     /cdrom/program files/IBM/SQLLIB/doc/htmlcd/%L/index.htm
     ```

     where */cdrom/* represents where the CD is mounted, and %L is the locale of the documentation that you wish to use, for example, **en_US** for English.

Related tasks:

- "Finding topics by accessing the DB2 Information Center from a browser" on page 96
- "Copying files from the DB2 HTML Documentation CD to a Web server" on page 102

Related reference:

- "Overview of DB2 Universal Database technical information" on page 85

Updating the HTML documentation installed on your machine

It is now possible to update the HTML installed from the *DB2 HTML Documentation CD* when updates are made available from IBM. This can be done in one of two ways:

- Using the Information Center (if you have the DB2 administration GUI tools installed).
- By downloading and applying a DB2 HTML documentation FixPak .

Note: This will NOT update the DB2 code; it will only update the HTML documentation installed from the *DB2 HTML Documentation CD*.

Procedure:

To use the Information Center to update your local documentation:

1. Start the DB2 Information Center in one of the following ways:
   - From the graphical administration tools, click on the **Information Center** icon in the toolbar. You can also select it from the **Help** menu.
   - At the command line, enter **db2ic**.
2. Ensure your machine has access to the external Internet; the updater will download the latest documentation FixPak from the IBM server if required.
3. Select **Information Center —> Update Local Documentation** from the menu to start the update.
4. Supply your proxy information (if required) to connect to the external Internet.

The Information Center will download and apply the latest documentation FixPak, if one is available.

To manually download and apply the documentation FixPak :

1. Ensure your machine is connected to the Internet.
2. Open the DB2 support page in your Web browser at: www.ibm.com/software/data/db2/udb/winos2unix/support.
3. Follow the link for Version 8 and look for the ″Documentation FixPaks″ link.
4. Determine if the version of your local documentation is out of date by comparing the documentation FixPak level to the documentation level you have installed. This current documentation on your machine is at the following level: **DB2 v8.1 GA**.
5. If there is a more recent version of the documentation available then download the FixPak applicable to your operating system. There is one FixPak for all Windows platforms, and one FixPak for all UNIX platforms.
6. Apply the FixPak:
   - For Windows operating systems: The documentation FixPak is a self extracting zip file. Place the downloaded documentation FixPak in an empty directory, and run it. It will create a **setup** command which you can run to install the documentation FixPak.

- For UNIX operating systems: The documentation FixPak is a compressed tar.Z file. Uncompress and untar the file. It will create a directory named `delta_install` with a script called **installdocfix**. Run this script to install the documentation FixPak.

Related tasks:

- "Copying files from the DB2 HTML Documentation CD to a Web server" on page 102

Related reference:

- "Overview of DB2 Universal Database technical information" on page 85

Copying files from the DB2 HTML Documentation CD to a Web server

The entire DB2 information library is delivered to you on the *DB2 HTML Documentation CD* and may be installed on a Web server for easier access. Simply copy to your Web server the documentation for the languages that you want.

Note: You might encounter slow performance if you access the HTML documentation from a Web server through a low-speed connection.

Procedure:

To copy files from the *DB2 HTML Documentation CD* to a Web server, use the appropriate source path:

- For Windows operating systems:

  `E:\program files\IBM\SQLLIB\doc\htmlcd\%L\*.*`

  where *E* represents the CD-ROM drive and *%L* represents the language identifier.
- For UNIX operating systems:

  `/cdrom/program files/IBM/SQLLIB/doc/htmlcd/%L/*.*`

  where *cdrom* represents the mount point for the CD-ROM drive and *%L* represents the language identifier.

Related tasks:

- "Searching the DB2 documentation" on page 104

Related reference:

- "Supported DB2 interface languages, locales, and code pages" in the *Quick Beginnings for DB2 Servers*
- "Overview of DB2 Universal Database technical information" on page 85

Troubleshooting DB2 documentation search with Netscape 4.x

Most search problems are related to the Java support provided by web browsers. This task describes possible workarounds.

Procedure:

A common problem with Netscape 4.x involves a missing or misplaced security class. Try the following workaround, especially if you see the following line in the browser Java console:

```
Cannot find class  java/security/InvalidParameterException
```

- On Windows operating systems:

  From the *DB2 HTML Documentation CD*, copy the supplied *x*:program files\IBM\SQLLIB\doc\htmlcd*locale*\InvalidParameterException.class file to the java\classes\java\security\ directory relative to your Netscape browser installation, where *x* represents the CD-ROM drive letter and *locale* represents the name of the desired locale.

  Note: You may have to create the java\security\ subdirectory structure.

- On UNIX operating systems:

  From the *DB2 HTML Documentation CD*, copy the supplied */cdrom*/program files/IBM/SQLLIB/doc/htmlcd/*locale*/InvalidParameterException.class file to the java/classes/java/security/ directory relative to your Netscape browser installation, where *cdrom* represents the mount point of the CD-ROM and *locale* represents the name of the desired locale.

  Note: You may have to create the java/security/ subdirectory structure.

If your Netscape browser still fails to display the search input window, try the following:

- Stop all instances of Netscape browsers to ensure that there is no Netscape code running on the machine. Then open a new instance of the Netscape browser and try to start the search again.
- Purge the browser's cache.
- Try a different version of Netscape, or a different browser.

Related tasks:

- "Searching the DB2 documentation" on page 104

Searching the DB2 documentation

You can search the library of DB2 documentation to locate information that you need. A pop-up search window opens when you click the search icon in the navigation toolbar of the DB2 Information Center (accessed from a browser). The search can take a minute to load, depending on the speed of your computer and network.

Prerequisites:

You need Netscape 6.1 or higher, or Microsoft's Internet Explorer 5 or higher. Ensure that your browser's Java support is enabled.

Restrictions:

The following restrictions apply when you use the documentation search:

- Search is not case sensitive.
- Boolean searches are not supported.
- Wildcard and partial searches are not supported. A search on *java** (or *java*) will only look for the literal string *java** (or *java*) and would not, for example, find *javadoc*.

Procedure:

To search the DB2 documentation:

1. In the navigation toolbar, click the **Search** icon.
2. In the top text entry field of the Search window, enter one or more terms (separated by a space) related to your area of interest and click **Search**. A list of topics ranked by accuracy displays in the **Results** field. The numerical ranking beside the hit provides an indication of the strength of the match (bigger numbers indicate stronger matches).

   Entering more terms increases the precision of your query while reducing the number of topics returned from your query.
3. In the **Results** list, click the title of the topic you want to read. The topic displays in the content frame of the DB2 Information Center.

Note: When you perform a search, the first (highest-ranking) result is automatically loaded into your browser frame. To view the contents of other search results, click on the result in the results list.

Related tasks:

- "Troubleshooting DB2 documentation search with Netscape 4.x" on page 103

Online DB2 troubleshooting information

With the release of DB2® UDB Version 8, there will no longer be a *Troubleshooting Guide*. The troubleshooting information once contained in this guide has been integrated into the DB2 publications. By doing this, we are able to deliver the most up-to-date information possible. To find information on the troubleshooting utilities and functions of DB2, access the DB2 Information Center from any of the tools.

Refer to the DB2 Online Support site if you are experiencing problems and want help finding possible causes and solutions. The support site contains a large, constantly updated database of DB2 publications, TechNotes, APAR (product problem) records, FixPaks, and other resources. You can use the support site to search through this knowledge base and find possible solutions to your problems.

Access the Online Support site at www.ibm.com/software/data/db2/udb/winos2unix/support, or by clicking the **Online Support** button in the DB2 Information Center. Frequently changing information, such as the listing of internal DB2 error codes, is now also available from this site.

Related concepts:

- "DB2 Information Center accessed from a browser" on page 108

Related tasks:

- "Finding product information by accessing the DB2 Information Center from the administration tools" on page 98

Accessibility

Accessibility features help users with physical disabilities, such as restricted mobility or limited vision, to use software products successfully. These are the major accessibility features in DB2® Universal Database Version 8:

- DB2 allows you to operate all features using the keyboard instead of the mouse. See "Keyboard Input and Navigation" on page 106.
- DB2 enables you customize the size and color of your fonts. See "Accessible Display" on page 106.
- DB2 allows you to receive either visual or audio alert cues. See "Alternative Alert Cues" on page 106.
- DB2 supports accessibility applications that use the Java™ Accessibility API. See "Compatibility with Assistive Technologies" on page 106.

- DB2 comes with documentation that is provided in an accessible format. See “Accessible Documentation”.

Keyboard Input and Navigation

Keyboard Input

You can operate the DB2 Tools using only the keyboard. You can use keys or key combinations to perform most operations that can also be done using a mouse.

Keyboard Focus

In UNIX-based systems, the position of the keyboard focus is highlighted, indicating which area of the window is active and where your keystrokes will have an effect.

Accessible Display

The DB2 Tools have features that enhance the user interface and improve accessibility for users with low vision. These accessibility enhancements include support for customizable font properties.

Font Settings

The DB2 Tools allow you to select the color, size, and font for the text in menus and dialog windows, using the Tools Settings notebook.

Non-dependence on Color

You do not need to distinguish between colors in order to use any of the functions in this product.

Alternative Alert Cues

You can specify whether you want to receive alerts through audio or visual cues, using the Tools Settings notebook.

Compatibility with Assistive Technologies

The DB2 Tools interface supports the Java Accessibility API enabling use by screen readers and other assistive technologies used by people with disabilities.

Accessible Documentation

Documentation for the DB2 family of products is available in HTML format. This allows you to view documentation according to the display preferences set in your browser. It also allows you to use screen readers and other assistive technologies.

DB2 tutorials

The DB2® tutorials help you learn about various aspects of DB2 Universal Database. The tutorials provide lessons with step-by-step instructions in the areas of developing applications, tuning SQL query performance, working with data warehouses, managing metadata, and developing Web services using DB2.

Before you begin:

Before you can access these tutorials using the links below, you must install the tutorials from the *DB2 HTML Documentation* CD.

If you do not want to install the tutorials, you can view the HTML versions of the tutorials directly from the *DB2 HTML Documentation CD*. PDF versions of these tutorials are also available on the *DB2 PDF Documentation CD*.

Some tutorial lessons use sample data or code. See each individual tutorial for a description of any prerequisites for its specific tasks.

DB2 Universal Database tutorials:

If you installed the tutorials from the *DB2 HTML Documentation* CD, you can click on a tutorial title in the following list to view that tutorial.

Business Intelligence Tutorial: Introduction to the Data Warehouse Center
Perform introductory data warehousing tasks using the Data Warehouse Center.

Business Intelligence Tutorial: Extended Lessons in Data Warehousing
Perform advanced data warehousing tasks using the Data Warehouse Center.

Development Center Tutorial for Video Online using Microsoft® Visual Basic
Build various components of an application using the Development Center Add-in for Microsoft Visual Basic.

Information Catalog Center Tutorial
Create and manage an information catalog to locate and use metadata using the Information Catalog Center.

Video Central for e-business Tutorial
Develop and deploy an advanced DB2 Web Services application using WebSphere® products.

Visual Explain Tutorial
Analyze, optimize, and tune SQL statements for better performance using Visual Explain.

DB2 Information Center accessed from a browser

The DB2® Information Center gives you access to all of the information you need to take full advantage of DB2 Universal Database™ and DB2 Connect™ in your business. The DB2 Information Center also documents major DB2 features and components including replication, data warehousing, the Information Catalog Center, Life Sciences Data Connect, and DB2 extenders.

The DB2 Information Center accessed from a browser has the following features if you view it in Netscape Navigator 6.1 or later or Microsoft Internet Explorer 5 or later. Some features require you to enable support for Java or JavaScript:

Regularly updated documentation
Keep your topics up-to-date by downloading updated HTML.

Search
Search all of the topics installed on your workstation by clicking **Search** in the navigation toolbar.

Integrated navigation tree
Locate any topic in the DB2 library from a single navigation tree. The navigation tree is organized by information type as follows:

- Tasks provide step-by-step instructions on how to complete a goal.
- Concepts provide an overview of a subject.
- Reference topics provide detailed information about a subject, including statement and command syntax, message help, requirements.

Master index
Access the information installed from the *DB2 HTML Documentation CD* from the master index. The index is organized in alphabetical order by index term.

Master glossary
The master glossary defines terms used in the DB2 Information Center. The glossary is organized in alphabetical order by glossary term.

Related tasks:

- "Finding topics by accessing the DB2 Information Center from a browser" on page 96
- "Finding product information by accessing the DB2 Information Center from the administration tools" on page 98
- "Updating the HTML documentation installed on your machine" on page 100

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/region 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/region 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.

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

IBM Canada Limited
Office of the Lab Director
8200 Warden Avenue
Markham, Ontario
L6G 1C7
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 document 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 illustrate 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 are trademarks of International Business Machines Corporation in the United States, other countries, or both, and have been used in at least one of the documents in the DB2 UDB documentation library.

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

The following terms are trademarks or registered trademarks of other companies and have been used in at least one of the documents in the DB2 UDB documentation library:

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

Intel and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

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

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Index

A

accessibility
  features 105
Add Partitions wizard 40
administration notification log 15
  replaces db2alert.log 3
ALTER BUFFERPOOL statement 33
Alter Database Partition Group wizard 40
applications
  two-phase commit for host and iSeries 83
  Web
    connection concentrator 27
    DB2 OLAP Server 77
    WebSphere Studio 73
Apply program
  committing changes by transaction 53
  monitoring 54
  one service per instance of program 53
  passwords encrypted in password file 53
asynchronous I/O, page cleaner performance 26
AUTOCONFIGURE command 23

B

backup utility
  XBSA support 15
block on log disk full (blk_log_dsk_ful) configuration parameter 13
buffer pools
  AUTOCONFIGURE command 23
  block-based, prefetching performance 25
built-in functions
  trigonometric 68

C

caching
  partitioned databases 27
CALL statements
  dynamically preparing 63
Capture program
  enhancements 51
Capture program *(continued)*
  monitoring 54
  multiple instances of 53
  one service per instance of program 53
CC/390 (Control Center 390) enhancements 42
centers
  Health Center 18
  Replication Center 49
clean transformer enhancements 57
Client Configuration Assistant
  renamed to Configuration Assistant 1
code pages
  euro currency symbol 23
  newly supported 2
columns
  mapping enhancements, Data Warehouse Center 58
Common warehouse metamodel (CWM) utility
  XML support 58
configurable online configuration parameters 32
Configuration Assistant (CA)
  enhancements 42
configuration parameters
  configurable online 32
  mirrorlogpath 13
connection concentrator 27
connections
  connection concentrator 27
constraints
  informational 68
containers
  DMS, online management 35
Control Center
  performance monitor removed 3
Control Center for OS/390 and z/OS 40, 42
country codes (renamed to territory code) 2
CREATE BUFFERPOOL statement 33

D

DAS (DB2 Administration Server)
  enhancements 39
data
  checking for required reorganization by schema 17
  Load wizard 40
  null and default values compressed 23
Data Links Manager
  performance 81
  platform availability 81
  replicating DATALINK columns 55
Data Warehouse Center
  clean transformer performance 57
  column mapping enhancements 58
  converting common warehouse metamodel (CWM) objects 58
  SQL Select and Update step 60
  updating target table columns 60
database configuration
  AUTOCONFIGURE command 23
Database Logging wizard 40
database management
  Tivoli Manager for DB2 22
database manager configuration
  AUTOCONFIGURE command 23
database partition groups
  Alter Database Partition Group wizard 40
  Redistribute Data wizard 40
databases
  quiescing for maintenance 17
DATALINK data type
  replication 55
DB2 Administration Server (DAS)
  enhancements 39
DB2 CLI
  LOAD functionality 69
DB2 control server
  partitioned database environments 37
DB2 Data Links Manager
  enhancements 81
  files, updating 81

DB2 documentation search
using Netscape 4.x 103
DB2 Enterprise Edition
merged with Extended Enterprise Edition 1
DB2 High Performance Unload for Multiplatforms tool 79
DB2 HTML documentation
installing 43
links between topics 42
master index 42
searching 44
supported browsers 42
unified table of contents 42
updates 44
DB2 Information Center 108
DB2 OLAP Integration Server
enhancements 77
DB2 OLAP Server
enhancements 77
DB2 Performance Expert for Multiplatforms tool 79
DB2 Recovery Expert for Multiplatforms tool 79
DB2 Relational Connect
enhancements 75
DB2 Setup wizard
enhancements 40
DB2 Table Editor for Multiplatforms tool 79
DB2 tutorials 107
DB2 Universal Database
deprecated function 3
Enterprise Server Edition 1
multilanguage installations on Windows platforms 45
multiplatform tools 79
product and packaging changes 1
Tivoli-ready 22
Workgroup Edition renamed 1
DB2 Warehouse Manager
AIX server 57
Information Catalog Manager renamed 60
Linux support 57
DB2 Web Query Tool for Multiplatforms tool 79
db2alert.log 3
db2Inspect() API 47
db2Load API
enhancements 11
deadlocks
event monitoring of 20
declared temporary tables
indexing 69
statistics 69
undo logging 69
Development Center
description 66
disability 105
DMS (database managed space)
online container management 35
dual logging
mirrorlogpath configuration parameter 13
dynamic dispatch, methods 68

E

euro currency symbol
support 23
support for 2
event monitors
control over time and timestamp data elements 21
for deadlocks 20
writing to SQL tables 20

F

federated systems
enhancements 75
flushing package cache 13

H

Health Center
overview 18

I

indexes
online reorganization of 32
renaming 23
type-2 22
online reorganization of tables 31
infinite active logging 13
Information Catalog Center 60
Information Catalog Manager
renamed to Information Catalog Center 60
informational constraints 68
INSPECT command 47
installing
DB2 HTML documentation 43
DB2 Universal Database
multiple languages on Windows platforms 45
INSTEAD OF triggers 68
IPX/SPX
not supported 3

J

JDBC
driver enhancements 71
JDBC 1.2 driver
deprecated 71

L

load utility
callable from CLI 69
enhancements 11
incremental maintenance of materialized query tables during append 34
Load wizard 40
logs
blocking logging on log disk full condition 13
Database Logging wizard 40
dual logging
mirrorlogpath configuration parameter 13
infinite active logging 13

M

materialized query tables (MQT)
Design Advisor 40
incremental maintenance during load append 34
incremental maintenance using a staging table 28
on nicknames 28
user-maintained 28
with non-aggregated joins 28
methods
dynamic dispatch of 68
overriding 68
monitoring
replication 54
MQT (materialized query tables)
Design Advisor 40
incremental maintenance during load append 34
incremental maintenance using a staging table 28
on nicknames 28
user-maintained 28
with non-aggregated joins 28
multidimensional clustering (MDC) 25
multiple wait
Data Warehouse Center 59

N

national language support (NLS)
enhancements 2

national language support (NLS) *(continued)*
euro currency symbol 2
new code pages, support 2
nodegroups
renamed to database partition group 2
NUMA-Q 3

O

OLE DB provider for DB2
overview 69
online
help, accessing 94
table load 31
online index reorganization
renamed to online index defragmentation of leaf pages 2
ordering DB2 books 94
OS/2 (not supported) 3
overriding
methods 68

P

packages
cache, flushing 13
multiple versions 16
page cleaners
asynchronous I/O, performance 26
partitioned database environments
Add Partitions wizard 40
Alter Database Partition Group wizard 40
Redistribute Data wizard 40
partitioned databases
authorization caching 27
catalog caching 27
performance
Data Links Manager 81
Design Advisor 40
multidimensional clustering tables 25
online reorganization of tables 31
type-2 indexes 22
prefetching
block-based buffer pools 25
printed books, ordering 94
processes, cascading, Data Warehouse Center 59
PTX not supported 3

Q

quiescing database for maintenance 17

R

Redistribute Data wizard 40
renaming
indexes 23
REORG INDEXES command 32
REORG TABLE command 31
REORGCHK command 17
replication
Capture program
enhancements 51
multiple instances of 53
control tables 50
DATALINK columns 55
monitoring 54
performance enhancements 53
Replication Center 49
Replication Analyzer
passwords encrypted in password file 53
Replication Center
monitoring Capture and Apply programs 54
overview 49
response files
creating
DB2 Setup wizard 40
restore utility
restoring database to different code page 15
rollforward utility
rolling forward to local time 15
routines
catalog views 63
creating with Development Center 66
implemented using thread model 27
SQL statements 63
run-time client, functions removed 1
RUNSTATS command
enhancements 17
global temporary tables 69

S

satellite administration
DB2CTLSV instance
partitioned database environment support 37
no support for generalizing replication subscriptions 3, 37
satellite administration *(continued)*
satellite function merged into DB2 Personal Edition 1
searching
DB2 HTML documentation 44
SELECT statement
retrieving
snapshots 21
snapshots
control over time and timestamp data elements 21
retrieving through SQL 21
SQL (Structured Query Language)
authorization
external routines 63
SQL Assist, enhancements 67
SQL Assist
enhancements 67
statistics
RUNSTATS command enhancements 17
Storage Management tool 12
Stored Procedure Builder, replaced by Development Center 1, 66
stored procedures
altering 63
Subsystem Cloning wizard 40
SUNLINK SNA, not supported 3

T

table spaces
checking
integrity, with INSPECT 47
not locked during load 31
tables
changed, replication 50
checking integrity with INSPECT 47
Load wizard 40
materialized query 28
incremental maintenance during load append 34
multidimensional clustering 25
new, replication 50
online reorganization 31
table space not locked during load 31
target, updating with Data Warehouse Center 60
writing event monitors 20
tasks
Task Center, DB2 for OS/390 batch processing 40
territory codes
new name for country code 2

threads
routines implemented using thread-based model 27
Tivoli Storage Manager (TSM)
using 22
tools
multiplatform 79
traces
enhancements 47
triggers
INSTEAD OF activation 68
trigonometric built-in functions 68
troubleshooting
administration notification log 15
DB2 documentation search 103
Health Center 18
online information 105
tracing 47
tutorials 107
two-phase commit
host and iSeries applications 83
type 2 indexes
described 22
online table reorganization 31
type 2 JDBC driver
enhancements 71
type 3 JDBC driver 71
type 4 JDBC driver
DB2 support 71

U

user-defined functions (UDFs)
table functions for snapshots 21

V

version levels
migration to Version 8 5

W

warehouse agents
Linux 57
warehouse server
AIX 57
warehousing
user-maintained materialized query tables 28
Web applications
connection concentrator 27
Web services
DB2 support 70
DB2 XML Extender 82
WebSphere Application Server 73
WebSphere Studio 73
windows
management instrumentation 68
Windows
management instrumentation 68
Windows 95, not supported 3
workloads
DB2 advisor 40
Design Advisor 40

X

XML
common warehouse metamodel (CWM) support 58
DB2 XML Extender 82
REC2XML and COLLATTVAL functions 73
schema validation UDF 73
XML Extender
enhancements 82
MQSeries applications 82

Contacting IBM

In the United States, call one of the following numbers to contact IBM:

- 1-800-237-5511 for customer service
- 1-888-426-4343 to learn about available service options
- 1-800-IBM-4YOU (426-4968) for DB2 marketing and sales

In Canada, call one of the following numbers to contact IBM:

- 1-800-IBM-SERV (1-800-426-7378) for customer service
- 1-800-465-9600 to learn about available service options
- 1-800-IBM-4YOU (1-800-426-4968) for DB2 marketing and sales

To locate an IBM office in your country or region, check IBM's Directory of Worldwide Contacts on the web at www.ibm.com/planetwide

Product information

Information regarding DB2 Universal Database products is available by telephone or by the World Wide Web at www.ibm.com/software/data/db2/udb

This site contains the latest information on the technical library, ordering books, client downloads, newsgroups, FixPaks, news, and links to web resources.

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) to order products or to obtain general information.
- 1-800-879-2755 to order publications.

For information on how to contact IBM outside of the United States, go to the IBM Worldwide page at www.ibm.com/planetwide

Part Number: CT17ZNA

Printed in U.S.A.

(1P) P/N: CT17ZNA

SC09-4848-00

Spine information:

IBM® DB2 Universal Database™ What's New Version 8