

IBM® DB2 Universal Database™



What's New

"Stinger"

IBM® DB2 Universal Database™



What's New

"Stinger"

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

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

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.

© Copyright International Business Machines Corporation 1997 - 2004. All rights reserved.

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

Contents

Chapter 1. Highlights of DB2 UDB "Stinger"	1
---	----------

Chapter 2. Disaster recovery technology enables around-the-clock availability	3
--	----------

Chapter 3. Autonomic computing technology make the DBA's job easier!	5
---	----------

Chapter 4. Enhanced integration with Microsoft .NET, J2EE/WebSphere Studio and Web services improves developer productivity	7
--	----------

Chapter 5. Protecting your valuable information resources.	9
---	----------

Chapter 6. Extend your solutions to include mobile data with DB2 mobility on demand	11
--	-----------

Chapter 7. Find the DB2 documentation you need using the DB2 Information Center	13
--	-----------

Chapter 8. Manageability enhancements	15
--	-----------

Manageability enhancements summary	15
Activity Monitor.	15
Index extensions in partitioned environments	16
System security enhancements for Windows improve usability	16
SQL statement size limit increased to 2 MB.	17
Log file inclusion in backup images	17
Self-tuning backup and restore operations	17
Automating database maintenance	18
Health Center enhancements	18
DB2 Universal Database is compliant with the FIPS 140-2 (Use of certified Cryptographic Module) government standard	19
New command for simplified database recovery	19
Simplified memory configuration and reduced memory errors	19
Advanced alter table functionality from the GUI	20
Support for Windows Local System account	20
Column default values can be altered.	21
Control Center for z/OS and OS/390 enhancements	21
Kerberos authentication support for the UNIX platforms	22
Pattern matching when using the db2look command	22
DB2 embedded application server enhancements	22

Automatic maintenance available when creating a database	23
View historical storage information for all objects in the storage management view	23

Chapter 9. Performance enhancements	25
--	-----------

Performance enhancements summary.	25
Improved RUNSTATS performance through sampling	25
Design advisor enhancements for recommending indexes, MQTs, MDC tables and partitions	26
Specifying a lock wait mode strategy.	27
Improved query execution plans through better cardinality estimation	27
Reduction of the time and space required to optimize complex queries.	28
Automatic statistics profiling	28
Automating database maintenance	28
Throttling of RUNSTATS utility.	29
Large page support in buffer pools on DB2 Universal Database for Linux	30
Multipage file allocation on SMS table spaces enabled by default	30
Automatic setting of prefetch size	30
XA timeout sample code	30

Chapter 10. Availability enhancements	33
--	-----------

Availability enhancements summary	33
High availability disaster recovery (HADR).	33
XML Extender platform support for HP-UX 32-bit and 64-bit PA-RISC	35
Automatic client reroute	35
Utility for demigrating database directories from "Stinger" to version 8.1	35

Chapter 11. Usability enhancements	37
---	-----------

Usability enhancements summary	37
Command Editor	37
Control Center enhancements	37
Utility progress monitoring	38
Changing generated column properties without having to recreate the table	39

Chapter 12. SQL enhancements	41
---	-----------

SQL enhancements summary	41
Column default values can be altered.	41
Changing generated column properties without having to recreate the table	41
SQL statement size limit increased to 2 MB.	42
SQL procedures no longer require a C or C++ compiler	42
Bind option REOPT enables reoptimization of SQL statements.	43
A new SQL/XML publishing function: XMLSERIALIZE.	43

Function XMLNAMESPACES new for "Stinger"	44
Unicode tables in a non-Unicode database supported	44
Specifying a lock wait mode strategy	44
Improved query execution plans through better cardinality estimation	45
Procedure invocation in a trigger body in DB2 UDB Enterprise Edition	46

Chapter 13. Serviceability enhancements 47

Serviceability enhancements summary	47
New utility to retrieve statistics from a running DB2 instance or database (similar to 'onstat' for Informix) db2diag.log format message enhancements	48
Diagnostic log analysis tool for filtering and formatting db2diag.log	48
Trace facility improved	49

Chapter 14. Data Warehouse Center and Warehouse Manager enhancements 51

Data Warehouse Center and Warehouse Manager ETL enhancements summary	51
Warehouse Manager ETL (Extract, Transform, and Load) supports Unicode warehouse control databases	51

Chapter 15. Application development enhancements 53

Application development enhancements summary	53
JDK 1.4 supported by DB2 Universal Database	53
DB2 Universal JDBC Driver improvements	54
OLE DB and ODBC .NET data providers	57
DB2 .NET Data Provider enhancements	57
.NET common language runtime procedures and UDFs	58
IBM DB2 Development Add-In for Microsoft Visual Studio .NET enhancements	58
DB2 Run-Time Client Lite	59
A new SQL/XML publishing function: XMLSERIALIZE	59
Procedure invocation in a trigger body in DB2 UDB Enterprise Edition	60
Unicode tables in a non-Unicode database supported	60
Nested savepoints allow for greater control of rollbacks	61
Bind option REOPT enables reoptimization of SQL statements	61
Simplified association of packages with an application using CURRENT PACKAGE PATH special register	62
SQL procedures no longer require a C or C++ compiler	62
Development Center stored procedure and debugging enhancements	62
Development Center no longer requires C compiler for building SQL stored procedures	63

Development Center supports SQL stored procedure debugging for UNIX 64-bit platforms	63
Development Center supports Java stored procedures on iSeries	63
DB2 WebSphere MQ user-defined function 64-bit support	63
Advanced alter table functionality from the GUI	64
XA timeout sample code	64

Chapter 16. Business intelligence enhancements 65

Business intelligence enhancements summary	65
4 Online import	65
Improved RUNSTATS performance through sampling	65
Design advisor enhancements for recommending indexes, MQTs, MDC tables and partitions	66
DB2 Data Warehouse Edition enhancements including integrated installation	67

Chapter 17. DB2 family enhancements 69

DB2 family enhancements summary	69
DB2 Run-Time Client Lite	69
DB2 Data Warehouse Edition enhancements including integrated installation	70
Find the DB2 documentation you need using the DB2 Information Center	70
Asynchronous I/O support on Linux (2.6 kernel exploitation)	71
DB2 Query Patroller 64-bit support for UNIX, Linux, and Windows	71
Licensing updates	72
Function XMLNAMESPACES new for "Stinger"	73
DB2 Express support on Linux PowerPC 32-bit	73
DB2 for Linux 64-bit support on the zSeries platform	73
DB2 for LinuxPPC 64-bit support	73
DB2 for Linux 64-bit support on iSeries and pSeries platforms	74
DB2 Connect Linux zSeries 64-bit support	74
DB2 extender enhancements	74
DB2 Geodetic Extender	74
Spatial Extender enhancements	74
XML Extender platform support for HP-UX 32-bit and 64-bit PA-RISC	75

Chapter 18. National language enhancements 77

National language enhancements summary	77
Thai and Unicode collation algorithm differences	77
New collators for Unicode databases	77

Chapter 19. Security enhancements 79

Security enhancements summary	79
Common Criteria certification of DB2 Universal Database products	79
System security enhancements for Windows improve usability	79
Support for Windows Local System account	80

Prevention of access to DB2 system files on Windows-based platforms	80
Access token on Windows provides users' group information	80
Custom security plugins for authentication and group management.	81
Kerberos authentication support for the UNIX platforms	81
Encryption of user data using two new authentication types	81

Appendix A. DB2 Universal Database technical information 83

DB2 documentation and help	83
DB2 documentation updates.	83
DB2 Information Center	84
DB2 Information Center installation scenarios	85
Installing the DB2 Information Center using the DB2 Setup wizard (UNIX)	88
Installing the DB2 Information Center using the DB2 Setup wizard (Windows)	90
Invoking the DB2 Information Center.	92
Updating the DB2 Information Center installed on your computer or intranet server	93
Displaying topics in your preferred language in the DB2 Information Center	94
DB2 PDF and printed documentation.	95
Core DB2 information	95
Administration information	95
Application development information	96
Business intelligence information	97
DB2 Connect information.	97

Getting started information	98
Tutorial information	98
Optional component information	98
Release notes	99
Printing DB2 books from PDF files	100
Ordering printed DB2 books	100
Invoking contextual help from a DB2 tool	101
Invoking message help from the command line processor	102
Invoking command help from the command line processor	103
Invoking SQL state help from the command line processor	103
DB2 tutorials	103
DB2 troubleshooting information	104
Accessibility	105
Keyboard input and navigation	105
Accessible display	105
Compatibility with assistive technologies	106
Accessible documentation	106
Dotted decimal syntax diagrams	106
Common Criteria certification of DB2 Universal Database products.	108

Appendix B. Notices 109

Trademarks	111
----------------------	-----

Index 113

Contacting IBM 115

Product information	115
-------------------------------	-----

Chapter 1. Highlights of DB2 UDB "Stinger"

DB2[®] Universal Database (DB2 UDB) "Stinger" delivers new features to address the ever-increasing demands and requirements on the most important data: our customer's. DB2 UDB "Stinger" provides new functionality to assist and resolve challenges that you, the customer, experience each day.

Database Administrators can reap immediate benefits from the broadened autonomic (also known as self-managing) computing solutions offered in DB2 UDB "Stinger". These solutions automate and simplify potentially time consuming and complex database tasks. For example, this release simplifies performance optimization for specific workloads by introducing the Design Advisor, a self-configuring tool that automates database design.

DBAs in Business Intelligence environments can take advantage of several DB2 UDB "Stinger" features including runstats enhancements and online import. RUNSTATS now exploits page and row-level sampling for faster statistics gathering. Faster statistics processing means more frequent updates are feasible, resulting in better optimizer plans and faster BI queries. This feature is used by DB2 in recent TPC-H benchmark results. The Online Import Utility exploits fine-grained locking for greater concurrency in online table loads. It can be applied to trickle-feeding a real-time warehouse.

For application developers, DB2 UDB "Stinger" delivers a significant amount of new capabilities as well as further integration of DB2 tooling into the Microsoft[®] .NET and WebSphere Studio/Java environments. This simplifies the development and deployment of DB2 applications allowing application developers to take advantage of the openness, performance and scalability of DB2, without regard to the back-end database or the chosen application architecture.

IT managers and DBAs will benefit from the integration of industry proven high availability disaster recovery (HADR) technology available in DB2 UDB "Stinger". Line of business managers and the enterprise itself, benefit the most as critical applications face less risk of downtime. HADR technology will help to ensure that the enterprise is business-enabled, around-the-clock without requiring time-consuming application changes.

DBAs and security administrators can provide greater security for their information assets by utilizing the pluggable security architecture. Custom security plug-ins allow DB2 to be plugged into any security scheme that utilizes userids and passwords for authentication, and userids and group membership to implement authorities.

Integrating DB2 into environments utilizing a Windows[®] security schema has been made easier by improving support for Active Directory.

Regardless of your role or the size of your business, DB2 UDB "Stinger" has exciting new features that can assist you with your daily challenges. Read on to learn more about these new features.

Related concepts:

- Chapter 7, "Find the DB2 documentation you need using the DB2 Information Center," on page 13

- Chapter 2, "Disaster recovery technology enables around-the-clock availability," on page 3
- Chapter 3, "Autonomic computing technology make the DBA's job easier!," on page 5
- Chapter 4, "Enhanced integration with Microsoft .NET, J2EE/WebSphere Studio and Web services improves developer productivity," on page 7
- Chapter 5, "Protecting your valuable information resources," on page 9
- Chapter 6, "Extend your solutions to include mobile data with DB2 mobility on demand," on page 11

Related reference:

- "Manageability enhancements summary" on page 15
- "Performance enhancements summary" on page 25
- "Availability enhancements summary" on page 33
- "Usability enhancements summary" on page 37
- "SQL enhancements summary" on page 41
- "Serviceability enhancements summary" on page 47
- "Data Warehouse Center and Warehouse Manager ETL enhancements summary" on page 51
- "Application development enhancements summary" on page 53
- "Business intelligence enhancements summary" on page 65
- "DB2 family enhancements summary" on page 69
- "National language enhancements summary" on page 77
- "Security enhancements summary" on page 79

Chapter 2. Disaster recovery technology enables around-the-clock availability

To stay ahead of the competition and drive profitability, today's e-business enterprises need to be open for business 24 hours a day, seven days a week. To ensure continuous availability for e-business, DB2[®] Universal Database offers an industry-proven high availability disaster recovery (HADR) feature, modeled on the high availability feature of the IBM[®] Informix[®] Dynamic Server. HADR is a database replication feature that provides a high availability solution for both partial and complete site failures. HADR protects against data loss by replicating data changes from a source database, called the primary, to a target database, called the standby. In the event of a partial or complete site failure, the standby database can quickly take over for the primary database.

The automatic client reroute feature can be used with HADR to enable client applications to recover from a loss of communication with the server and to continue working with minimal interruption. You can use automatic client reroute with HADR to have client applications connect to the new primary database after a takeover operation.

Related concepts:

- "High availability disaster recovery (HADR)" on page 33
- "Automatic client reroute" on page 35

Chapter 3. Autonomic computing technology make the DBA's job easier!

If you want to improve the productivity and effectiveness of your database administrators, check out the new manageability features in DB2® Universal Database, "Stinger". An exciting example of these features is the Design advisor, which assists DBAs in making optimal and comprehensive database design decisions. Many DBAs would agree that the decisions made with respect to the design of a database are some of the most challenging, time consuming, and critical to make. This self-configuring tool greatly simplifies the design process by using workload, database, and hardware information to recommend specific performance acceleration options for routine design tasks. Specifically, the new Design advisor assists with the following design tasks:

- Index selection
- Materialized Query Tables selection
- Multidimensional Clustering selection
- Partitioning selection

Growing out of work at the IBM® Almaden Research Lab, this advisor makes it easier and faster to design partitioned data warehouses for high performance.

Another example of a new time saving feature for a critical DBA task is the ability to automate database maintenance activities such as backup, table defragmentation (reorg), and table statistics gathering (runstats). DB2 Universal Database™ "Stinger" provides the ability to easily schedule these tasks to run automatically based on factors determined by the DBA such as available maintenance windows, percentage of log files used, and so on.

Other manageability enhancements in "Stinger" include self-healing features such as the Health Center Recommendation advisor and automated log file management; and self-tuning features for backup and restore operations as well as the ability to throttle backup operations and statistics collection. Also available is enhanced sampling functionality, both row and page level, for faster, and potentially more frequent collection of statistics. This improves query optimization in challenging business intelligence environments.

Related concepts:

- "Log file inclusion in backup images" on page 17
- "Design advisor enhancements for recommending indexes, MQTs, MDC tables and partitions" on page 26
- "Automatic statistics profiling" on page 28
- "Self-tuning backup and restore operations" on page 17
- "Automating database maintenance" on page 18
- "Health Center enhancements" on page 18
- "New command for simplified database recovery" on page 19
- "Throttling of RUNSTATS utility" on page 29
- "Automatic setting of prefetch size" on page 30

Chapter 4. Enhanced integration with Microsoft .NET, J2EE/WebSphere Studio and Web services improves developer productivity

DB2® UDB "Stinger" is a significant release for application developers with significant new SQL features, better support for open standard APIs, web services and seamless integration with the key development environments for both Java™ and Microsoft® .NET. It builds on the strong momentum we've gained since the release of DB2 UDB Version 8 for .NET support in both driver enablement and tooling. The DB2 .NET Data Provider has much improved performance and now supports DB2 UDB server on iSeries™. DB2 UDB servers also support stored procedures written in Microsoft CLR (Common Language Runtime) compliant languages such as Visual Basic .NET and C#. The DB2 tools add-ins to Microsoft Visual Studio .NET are further augmented with schema operation capabilities, wizards for developing and deploying CLR stored procedures as well as WORF (Web services Object Runtime Framework) web services.

For Java developers, we now support JDK 1.4 as a runtime as well as a development environment for all DB2 platforms. We also put in XA support into the DB2 Universal JDBC driver so that it can be used in mission critical enterprise applications.

Web services continues to be our focus as the new application development paradigm. "Stinger" has enhanced web service support as a provider. SOAP engine transparency is now provided, allowing more choices of soap engines to be plugged into the run time. With the new dynamic query capability, the web service client can determine at invocation which query should be executed, in addition to the fixed transactions/queries in DADXes. A tracing component has also been added to make it easier to diagnose errors in the provider. Another important feature is the support for web service inspection language (WSIL) so that users can easily explore the web services that are offered in the DB2 web services provider.

Also new in "Stinger" is the significantly reduced footprint of the DB2 Run-Time Client Lite that makes it easier to build and redistribute DB2 applications.

A number of SQL enhancements are available in DB2 UDB "Stinger" that makes the migration of applications built with other database systems much easier. The dependency on a C compiler is removed for building platform independent SQL procedures. The size limit of a single SQL statement has been increased from 64K to 2 MB, which is very desirable for applications such as in BI analysis. These applications often consist of very complex and lengthy statements. Stored procedures can now be called from within a trigger body.

Related concepts:

- "Procedure invocation in a trigger body in DB2 UDB Enterprise Edition" on page 46
- "SQL procedures no longer require a C or C++ compiler" on page 42
- "SQL statement size limit increased to 2 MB" on page 17
- "DB2 Universal JDBC Driver improvements" on page 54
- ".NET common language runtime procedures and UDFs" on page 58
- "DB2 Run-Time Client Lite" on page 59

- “JDK 1.4 supported by DB2 Universal Database” on page 53
- “Development Center supports Java stored procedures on iSeries” on page 63
- “Development Center no longer requires C compiler for building SQL stored procedures” on page 63
- “DB2 .NET Data Provider enhancements” on page 57
- “OLE DB and ODBC .NET data providers” on page 57
- “Development Center supports SQL stored procedure debugging for UNIX 64-bit platforms” on page 63
- “Nested savepoints allow for greater control of rollbacks” on page 61
- “Development Center stored procedure and debugging enhancements” on page 62

Chapter 5. Protecting your valuable information resources

Securing IT systems has become increasingly important for companies of all sizes. The enhancements in DB2[®] Universal Database (DB2 UDB) "Stinger" significantly improve your ability to secure your DB2 environment. DB2 UDB now supports security plug-ins that allow customers to customize DB2 security to meet their company's needs. Support for Kerberos security is provided by a sample security plug-in.

In Windows[®] environments, the restrictions on authorization IDs and group names are now relaxed to be more consistent with Windows standards. Support for Microsoft[®] Active Directory domains has been enhanced. You can now install and run the Windows services for DB2 under the local system account ('SYSTEM').

Related concepts:

- "System security enhancements for Windows improve usability" on page 16
- "Custom security plugins for authentication and group management" on page 81
- "Prevention of access to DB2 system files on Windows-based platforms" on page 80
- "Support for Windows Local System account" on page 20
- "Kerberos authentication support for the UNIX platforms" on page 22

Chapter 6. Extend your solutions to include mobile data with DB2 mobility on demand

You can leverage the mobility on demand capability of DB2[®] Universal Database (DB2 UDB) to extend your current solutions to include mobile data. The mobility on demand capability, based on DB2 Everyplace[®] technology, includes the high performance, robust DB2 Everyplace database as well as a powerful synchronization solution for use with existing DB2 UDB deployments. This capability was introduced with version 8.1.4 of DB2 UDB.

You can download the mobility on demand feature from the DB2 Everyplace Web site if you are an IBM[®] Business Partner or a customer of DB2 UDB Express Edition, DB2 UDB Workgroup Server Edition, DB2 UDB Workgroup Server Unlimited Edition, DB2 UDB Enterprise Server Edition, or DB2 Universal Developer's Edition.

For more information, including instructions on downloading the mobility on demand feature, see <http://www.ibm.com/software/data/db2/everyplace/mobilityondemand.html>

Chapter 7. Find the DB2 documentation you need using the DB2 Information Center

The DB2[®] Information Center for "Stinger" sets a new standard for delivering DB2 Universal Database[™] and related DB2 product information by integrating it into a single information center. The DB2 Information Center gives you access to all of the DB2 information that you need, covering every aspect of DB2 from getting started, to using the DB2 Tools, database connectivity, database administration, query management, business intelligence, application development, and more. The DB2 Information Center also documents major DB2 features and components including replication, data warehousing, and the DB2 extenders. The navigation tree (table of contents) consists primarily of links to task and concept information that aligns with the high-level goals that you and other DB2 users want to achieve. In addition, the navigation tree includes entries for product overviews, reference information, a master index of all topics, and a glossary.

The DB2 Information Center is a fully searchable server that provides quick access to DB2 product information. To refine your search of all the topics in the DB2 Information Center, you can now use wildcards, quotation marks to search for phrases, and Boolean operators (AND, NOT, OR). If a translated version of a topic that you are searching for is not available in the preferred language set in your browser preferences, the DB2 Information Center displays the English version of that topic by default.

Using the DB2 Information Center Setup wizard, you can install the DB2 Information Center locally on your computer. However, if the computer on which you install the DB2 Information Center is connected to a network, other computers can also access it through a Web browser. You can also decide not to install the DB2 Information Center locally, choosing instead to have your DB2 products access a version of the DB2 Information Center at an alternate location, such as the IBM[®] Web site where the most recently released version is available:
<http://publib.boulder.ibm.com/infocenter/db2help/>

Installing the DB2 Information Center locally lets you access DB2 product information without needing Internet or intranet access, which might be useful in a test or application development environment. If you install the DB2 Information Center on a networked computer, multiple users throughout your organization can access the same DB2 Information Center, giving you control over the version of the information that those users can see. This will minimize your update efforts and keep your network traffic within your intranet. If you want users to access the most up-to-date information, and you have access to the Internet, you can choose to configure DB2 products to access the DB2 Information Center on the IBM Web site. This version of the DB2 Information Center on the IBM Web site is updated as necessary to reflect the most recently released version of DB2 Universal Database, and also includes access to online service documents like TechNotes.

Related concepts:

- "DB2 Information Center" on page 84

Related tasks:

- "Displaying topics in your preferred language in the DB2 Information Center" on page 94

- “Invoking the DB2 Information Center” on page 92
- “Installing the DB2 Information Center using the DB2 Setup wizard (UNIX)” on page 88
- “Installing the DB2 Information Center using the DB2 Setup wizard (Windows)” on page 90

Chapter 8. Manageability enhancements

Manageability enhancements summary

The following manageability enhancements have been added for this release:

Related concepts:

- “Activity Monitor” on page 15
- “Index extensions in partitioned environments” on page 16
- “System security enhancements for Windows improve usability” on page 16
- “SQL statement size limit increased to 2 MB” on page 17
- “Log file inclusion in backup images” on page 17
- “Simplified memory configuration and reduced memory errors” on page 19
- “Self-tuning backup and restore operations” on page 17
- “Automating database maintenance” on page 18
- “Health Center enhancements” on page 18
- “DB2 Universal Database is compliant with the FIPS 140-2 (Use of certified Cryptographic Module) government standard” on page 19
- “New command for simplified database recovery” on page 19
- “Advanced alter table functionality from the GUI” on page 20
- “Support for Windows Local System account” on page 20
- “Column default values can be altered” on page 21
- “DB2 embedded application server enhancements” on page 22
- “Control Center for z/OS and OS/390 enhancements” on page 21
- “Kerberos authentication support for the UNIX platforms” on page 22
- “Pattern matching when using the db2look command” on page 22
- “Automatic maintenance available when creating a database” on page 23
- “View historical storage information for all objects in the storage management view” on page 23

Activity Monitor

The Activity Monitor is a tool that assists database administrators in improving the efficiency of database performance monitoring, problem determination, and resolution. The Activity Monitor focuses on monitoring application performance, application concurrency, resource consumption, and SQL statement usage. It will help DBAs to diagnose the cause of database performance problems such as application locking situations, and to tune queries for optimal utilization of database resources.

The Activity Monitor provides easy access to relevant and well-organized monitor data through a set of predefined reports such as *Top CPU-time consuming applications* and *SQL statements with the largest total sort time*. For each predefined report, appropriate actions may be recommended to help resolve resource utilization problems, to optimize performance, or to invoke another tool for further investigation.

Lock monitor data is also supplied to illustrate the details of lock waiting situations. Application lock chains can be displayed to show lock waiting dependencies.

The Activity Monitor is accessible through a GUI interface, the command line processor, and in the form of stored procedures and user-defined functions.

Related tasks:

- “Activity Monitor overview: Activity Monitor help”

Index extensions in partitioned environments

With DB2® Universal Database, "Stinger", you can create an index extension for use in a partitioned environment in addition to support for a single partition environment. This feature is an addition to support that already exists for a single partition. An index extension is an index object for use with indexes that have structured type or distinct type columns. You can create and apply your own logic for the primary components that make up how an index works. The primary components include index maintenance and index exploitation.

Related concepts:

- “Creating a user-defined extended index type” in the *Administration Guide: Implementation*
- “Details on index maintenance” in the *Administration Guide: Implementation*
- “Details on index searching” in the *Administration Guide: Implementation*
- “Details on index exploitation” in the *Administration Guide: Implementation*
- “A scenario for defining an index extension” in the *Administration Guide: Implementation*

System security enhancements for Windows improve usability

DB2® Universal Database, "Stinger", introduces the following enhancements to system security on the Windows® platform:

- Accepting additional special characters in user IDs and security mechanism group names (and consequently in authorization-names and authorization-ids) including &, - , and blank
- Accepting security mechanism group names that are longer than eight characters such as "Sales Representatives"
- Allowing two-part names on CONNECT and ATTACH that contain a Windows domain name and the userid to avoid the network traffic associated with looking up the user name in the trusted domain forest
- Support for Active Directory domains has been enhanced. For example, implicit trusts between domains, domain local groups and nested global groups are supported.

Related concepts:

- “DB2 UDB object naming rules” in the *Administration Guide: Implementation*
- “User, user ID and group naming rules” in the *Administration Guide: Implementation*

SQL statement size limit increased to 2 MB

DB2® Universal Database, "Stinger", provides an increased SQL statement size limit of 2 MB, which can limit the size of some created database objects (for example, CREATE PROCEDURE or CREATE TRIGGER statements). The new 2 MB limit allows you to use large statements when your application logic is in stored procedures or triggers.

The increase in the statement size allows you to migrate a trigger or stored procedure statement from another RDBMS to DB2 UDB.

The new statement limit allows you to record auditing context records that have statement text up to 2 MB.

Related reference:

- "Audit record layout for CONTEXT events" in the *Administration Guide: Implementation*

Log file inclusion in backup images

When creating an online backup image, you can include the log files required to restore and recover a database in the image. This means that if you need to ship backup images to a disaster recovery site, you do not have to send the log files separately or package them together yourself. Further, you do not have to decide which log files are required to guarantee the consistency of an online backup, and you will not be able to delete the log files that pertain to a particular backup image.

To use this feature, specify the INCLUDE LOGS option of the BACKUP DATABASE command. When you specify this option, the backup utility will truncate the currently active log file and copy the necessary set of log extents into the backup image.

Related concepts:

- "Including log files with a backup image" in the *Data Recovery and High Availability Guide and Reference*

Self-tuning backup and restore operations

DB2® will now automatically choose the number of buffers, the buffer size and the parallelism settings for both backup and restore operations. The values chosen are based on the amount of memory available, the number of processors available and the database configuration. The objective is to minimize the amount of time required for backup and restore operations to complete. The BACKUP DATABASE and RESTORE DATABASE commands will automatically choose an optimal value for the following parameters whenever they are not explicitly specified:

- WITH num-buffers BUFFERS
- PARALLELISM n
- BUFFER buffer-size

For restore database operations, a multiple of the buffer size used for the backup operation will always be used.

The values specified by the database manager configuration parameters BACKBUFSZ and RESTBUFSZ are ignored. If you want to use these values, you must explicitly specify them when you issue the BACKUP DATABASE or RESTORE DATABASE command.

Related concepts:

- “Backup overview” in the *Data Recovery and High Availability Guide and Reference*
- “Restore overview” in the *Data Recovery and High Availability Guide and Reference*

Automating database maintenance

You can now automate database maintenance activities to run only when they are needed using the Configure Automatic Maintenance wizard. The activities you can configure for automatic maintenance are database backup, data defragmentation (table or index reorganization), and data access optimization (running statistics). Performing maintenance activities on your databases is essential to ensure that they are optimized for performance and recoverability.

The DB2® optimizer uses catalog statistics to determine the most efficient access plan for any given query. Having out-of-date or incomplete statistics for a table or an index could lead the optimizer to select a plan that is much more inefficient than other alternatives, slowing down query execution. However, deciding which statistics to collect for a given workload is complex, and keeping these statistics up to date can be time-consuming.

With automatic statistics collection, part of DB2's Automatic Table Maintenance feature, you can let DB2 determine which statistics are required by your workload and which statistics need to be updated. With automatic statistics collection enabled, DB2 will automatically run the RUNSTATS utility in the background as required to ensure that the correct statistics are collected and maintained.

Use the Configure Automatic Maintenance wizard to specify your maintenance objectives, including when automatic maintenance can run. DB2 uses the objectives that you have specified to determine whether the maintenance activities need to be done. Then DB2 runs only the required maintenance activities during the next available maintenance window (a user-defined time period for the running of automatic maintenance activities).

Related concepts:

- “Automatic statistics profiling” on page 28
- “Automatic maintenance” in the *Administration Guide: Planning*

Related tasks:

- “Configuring automatic maintenance: Control Center help”

Health Center enhancements

The Health Center for DB2® Universal Database, “Stinger” has a number of new features:

Health Indicator Configuration launchpad

The Health Indicator Configuration launchpad provides easy access to various levels of health indicator settings. This launchpad also helps you to understand how these configurations apply to health monitoring, allowing you to make changes appropriate to your database system environment.

Recommendation advisor

The Recommendation advisor helps you to resolve health alerts on DB2 objects. This advisor provides recommendations that can correct the issue causing a health alert. In addition, the Recommendation advisor helps you to implement the recommendation that you select, whether this requires launching a tool, running a script, or adjusting the configuration parameter settings on an instance or a database.

Troubleshoot Health Alert Notification wizard

The Troubleshoot Health Alert Notification wizard helps you to troubleshoot various problems related to health alert notification, which ensures that notifications are received properly by the correct contacts.

Related tasks:

- “Using the Health Indicator Configuration Launchpad: Health Center help”
- “Resolving health alerts using the Recommendation advisor: Health Center help”
- “Troubleshooting health alert notification: Health Center help”

DB2 Universal Database is compliant with the FIPS 140–2 (Use of certified Cryptographic Module) government standard

DB2[®] Universal Database is compliant with the FIPS 140–2 (Use of certified Cryptographic Module) government standard.

Related concepts:

- “Authentication methods for your server” in the *Administration Guide: Implementation*

New command for simplified database recovery

The new RECOVER DATABASE command combines the functionality of the RESTORE DATABASE and ROLLFORWARD DATABASE commands. When you use this command, you specify the point-in-time to which you want the database recovered. You do not need to indicate which database backup image must be restored or which log files are required to reach the specified point-in-time. The RECOVER DATABASE command also supports recovery operations to the end of the log files.

Related concepts:

- “Managing log files” in the *Data Recovery and High Availability Guide and Reference*

Related reference:

- “RECOVER DATABASE Command” in the *Command Reference*

Simplified memory configuration and reduced memory errors

There are several disadvantages to configuring database heap sizes as maximum hard limits on memory usage:

- Determining the appropriate maximum size for a particular database heap requires extensive knowledge about how the heap will be used by DB2[®].
- Setting a limit on the amount of memory that can be allocated to a heap does not guarantee that memory will be available to that heap when it is required.

- The maximum size of a heap represents a hard limit, so that memory allocation for a heap cannot exceed this limit, even for a short period of time.

To address these limitations, database shared memory heaps and instance shared memory heaps will be configured and interpreted differently. To simplify configuration of these heaps, you will now be able to specify a single value for how much memory DB2 uses for each active database and specify guaranteed minimums reserved for individual heaps. Interpreting configuration parameter values for heap sizes as guaranteed minimums will reduce the number of memory errors experienced during memory usage peaks.

Related reference:

- “dbheap - Database heap configuration parameter” in the *Administration Guide: Performance*
- “instance_memory - Instance memory configuration parameter” in the *Administration Guide: Performance*
- “database_memory - Database shared memory size configuration parameter” in the *Administration Guide: Performance*

Advanced alter table functionality from the GUI

When altering a table from within the Alter Table notebook, you can now make more changes to table columns than you could previously. These changes are possible because the Alter Table notebook now drops and re-creates the table on an as-needed basis.

You can make the following additional changes when altering a table:

- Rename a column
- Drop a column
- Change the data type of a column
- Change the length, scope, or precision values for a column
- Change whether a column is nullable

If the table is dropped and re-created, attempts will be made to restore any dependent objects and to transform the existing data into the target data type of each remaining column.

Related tasks:

- “Alter Table -- Overview: Control Center help”

Support for Windows Local System account

Support for Windows Local System account (LSA) is provided for both the various DB2 services, as well as for applications that access them. The DB2 installation process allows DB2 services to run under SYSTEM (local system account) as an alternative to running under a dedicated user account. You can also perform the DB2 installation itself from a process that is running under the context of the Local System account.

This feature makes it easier to deploy DB2 by avoiding, where appropriate, the complexities associated with handling of userids and passwords. Another enhancement is the ability for DB2 applications that are running in the context of the LSA to access and use the local DB2 server.

Related concepts:

- “Windows local system account support” in the *Administration Guide: Implementation*

Related reference:

- “Required user accounts for installation of DB2 servers (Windows)” in the *Quick Beginnings for DB2 Servers*

Column default values can be altered

There are times when the default values of a column in a database table must be changed. The ALTER TABLE SQL statement allows you to set the default clause of a column in a table. The same default clause restrictions as documented under the CREATE TABLE SQL statement are enforced. Note that generated columns cannot have their default value altered by this statement.

Related reference:

- “ALTER TABLE statement” in the *SQL Reference, Volume 2*

Control Center for z/OS and OS/390 enhancements

Support for DB2® for z/OS V8

The following enhanced commands are now supported by the Control Center for z/OS and OS/390 and allow changes to affected database objects with minimized impact on availability:

- Create Table
- Alter Table
- Create Index
- Alter Index

Extended support of DB2 commands

The following DSN subcommands can now be issued from the Control Center for z/OS and OS/390:

- Bind Package
- Bind Plan
- Free Package
- Free Plan
- Rebind Package
- Rebind Plan
- Rebind Trigger Package

Create Object Maintenance Policy wizard

The new Create Object Maintenance Policy wizard guides you through a series of steps to collect the information necessary for creating the maintenance policy and for generating the JCL job required to perform object maintenance on DB2 for z/OS and OS/390 subsystems.

Secondary authorization ID

Control Center now supports the SET CURRENT SQLID command when administering DB2 for z/OS and OS/390 subsystems. This allows for flexibility in granting access to members of a group, without the complications associated with changes to individual users profiles.

Kerberos authentication support for the UNIX platforms

Kerberos is a third party network authentication protocol that employs a system of shared secret keys to securely authenticate a user in an insecure network environment without flowing a text userid or password. Using Kerberos provides the advantages of single sign-on access to resources, such as a DB2[®] UDB server, and the centralization of user (principal) administration. Kerberos authentication support is extended to DB2 UDB for AIX[®], and for Solaris Operating Environment, in addition to existing support for Windows[®] operating systems that support Active Directory.

This support is provided as a security plug-in.

Related concepts:

- “Authentication methods for your server” in the *Administration Guide: Implementation*
- “Custom security plugins for authentication and group management” on page 81

Pattern matching when using the db2look command

The **db2look** command has a new option, `-tw`, that can be used to generate the DDL for tables whose names match a specified pattern. The underscore (`_`) and percent sign (`%`) wildcard characters can be used in the `-tw` argument to specify the pattern-matching criteria for the tables for which you want the DDL. **db2look** generates the DDL both for all tables with names that match the pattern, and for the dependent objects of these tables. The `-tw` option augments the ability to specify a list of tables (available with the `-t` option), and can provide significant productivity gains in environments with strict table-naming standards when DDL must be extracted from a test environment for promotion to a production environment.

Related reference:

- “db2look - DB2 Statistics and DDL Extraction Tool Command” in the *Command Reference*

DB2 embedded application server enhancements

DB2[®] Universal Database (DB2 UDB) Version 8.1 includes an embedded application server. This enables DB2-supplied web applications to run without needing a separately installed application server. For DB2 UDB “Stinger”, remote administration of the embedded application server is improved. You can now administer the embedded application server using stored procedures.

After a database is enabled, you can administer the embedded application server in the following ways:

- Start and stop the embedded application server. (Use the DB2EAS.SERVER stored procedure.)
- Install a DB2 web application to be served by the embedded application server. (Use the DB2EAS.INSTALLAPP stored procedure.)
- Uninstall a DB2 web application that was served by the embedded application server. (Use the DB2EAS.UNINSTALLAPP stored procedure.)

- Start and stop the DB2 web application being served by the embedded application server. (Use the DB2EAS.APP stored procedure.)

Related tasks:

- “Installing the application server for DB2” in the *Installation and Configuration Supplement*
- “Automatically deploying DB2 Web Tools on the application server for DB2” in the *Installation and Configuration Supplement*
- “Enabling the application server for DB2” in the *Installation and Configuration Supplement*

Automatic maintenance available when creating a database

You can now enable various automatic maintenance features when creating a database (from the Control Center as well as from First Steps). The automatic maintenance features can:

- Create a new database on the disk or directory of your choice
- Assign disk space for the data
- Configure the new database for performance
- Turn on automatic maintenance
- Configure notification by e-mail or pager if the database needs attention

Related tasks:

- “Configuring automatic maintenance: Control Center help”

View historical storage information for all objects in the storage management view

Using the storage management tool you can now view historical storage information for all objects in the storage management view, including the database, table spaces, tables, containers, and indexes.

You can now also view container information, as well as view the percentage of space used in SMS table spaces in single partitioned databases.

A new stored procedure is also available called SYSPROC.GET_DBSIZE_INFO which calculates the database size and capacity.

Related tasks:

- “Storage Management view -- Overview: Control Center help”

Related reference:

- “GET_DBSIZE_INFO procedure” in the *SQL Administrative Routines*

Chapter 9. Performance enhancements

Performance enhancements summary

The following performance enhancements have been added for this release:

Related concepts:

- “Improved RUNSTATS performance through sampling” on page 25
- “Design advisor enhancements for recommending indexes, MQTs, MDC tables and partitions” on page 26
- “Specifying a lock wait mode strategy” on page 27
- “Improved query execution plans through better cardinality estimation” on page 27
- “Reduction of the time and space required to optimize complex queries” on page 28
- “Automatic statistics profiling” on page 28
- “Automating database maintenance” on page 18
- “Throttling of RUNSTATS utility” on page 29
- “Large page support in buffer pools on DB2 Universal Database for Linux” on page 30
- “Multipage file allocation on SMS table spaces enabled by default” on page 30
- “Automatic setting of prefetch size” on page 30
- “XA timeout sample code” on page 30

Improved RUNSTATS performance through sampling

Table statistics are used by the query optimizer in selecting the best access plan for any given query, so it is important that statistics remain current to accurately reflect the state of a table at any given time. As the activity against a table increases, so should the frequency of statistics collection. With the increasing size of databases, it is becoming more important to find efficient ways to collect statistics. Random sampling of table data on which to collect statistics can reduce the amount of time that it takes to collect statistics. For I/O bound or CPU bound systems, the performance benefits can be enormous. The smaller the sample, the faster statistics collection completes.

Starting in "Stinger", the RUNSTATS command provides the option to collect statistics on a sample of the data in the table by using the TABLESAMPLE option. This feature can increase the efficiency of statistics collection since sampling uses only a subset of the data. At the same time, the sampling methods ensure a high level of accuracy.

Related concepts:

- “Data sampling in SQL queries” in the *Administration Guide: Performance*
- “Collecting statistics on a sample of the table data” in the *Administration Guide: Performance*

Related reference:

- “RUNSTATS Command” in the *Command Reference*

Design advisor enhancements for recommending indexes, MQTs, MDC tables and partitions

The DB2® Design advisor is a tool that can help you significantly improve your workload performance. The task of selecting which indexes, clustering dimensions, or partitions to create for a complex workload can be quite daunting. The Design advisor identifies all of the objects needed to improve the performance of your workload. Given a set of SQL statements in a workload, the Design advisor will generate recommendations for:

- new indexes
- new materialized query tables
- conversion to multidimensional clustering tables
- repartitioning of tables
- deletion of objects unused by the specified workload

You can choose to have the Design advisor GUI tool implement some or all of these recommendations immediately or schedule them for a later time.

Using either the Design advisor GUI or the command-line tool, the Design advisor can help simplify the following tasks:

Planning for or setting up a new database or partitioning structure

While designing your database or database partitions, use the Design advisor to:

- Generate design alternatives in a test environment for partitioning, indexes, MQTs, and MDC tables.
- Determine initial database partitioning before loading data into a database.
- Assist in migrating from a non-partitioned DB2 database to a partitioned DB2 database.
- Assist in migrating to DB2 in a partitioned environment from another database product.
- Evaluate indexes, MQTs or partitions that have been generated manually.

Workload performance tuning

After your database is set up, you can use the Design advisor to help you meet the following tuning goals:

- Improve performance of a particular statement or workload.
- Improve general database performance, using the performance of a sample workload as a gauge.
- Improve performance of the most frequently executed queries, for example, as identified by the Activity Monitor.
- Determine how to optimize the performance of a new key query.
- Respond to Health Center recommendations regarding shared memory utility or sort heap problems in a sort-intensive workload.
- Find objects that are not used in a workload.

Related concepts:

- “Automatic summary tables” in the *Administration Guide: Performance*
- “Multidimensional clustering tables” in the *Administration Guide: Planning*

Specifying a lock wait mode strategy

An individual session can now specify a lock wait mode strategy, which is used when the session requires a lock that it cannot obtain immediately. The strategy indicates whether the session will:

- Return an `SQLCODE` and `SQLSTATE` when it cannot obtain a lock
- Wait indefinitely for a lock
- Wait a specified amount of time for a lock
- Use the value of the `locktimeout` database configuration parameter when waiting for a lock

The lock wait mode strategy is specified through the new `SET CURRENT LOCK TIMEOUT` statement, which changes the value of the `CURRENT LOCK TIMEOUT` special register. The `CURRENT LOCK TIMEOUT` special register specifies the number of seconds to wait for a lock before returning an error indicating that a lock cannot be obtained.

Traditional locking approaches can result in applications blocking each other. This happens when one application must wait for another application to release its lock. Strategies to deal with the impact of such blocking usually provide a mechanism to specify the maximum acceptable duration of the block. That is the amount of time that an application will wait prior to returning without a lock. Previously, this was only possible at the database level by changing the value of the `locktimeout` database configuration parameter.

Whereas the value of the `locktimeout` parameter applies to all locks, the lock types that are impacted by this new function include row, table, index key, and multidimensional clustering (MDC) block locks.

Related reference:

- “locktimeout - Lock timeout configuration parameter” in the *Administration Guide: Performance*

Improved query execution plans through better cardinality estimation

Cardinality estimation is the process whereby the optimizer determines the number of qualifying rows after filtering operations, such as predicate application or aggregation, are applied. Accurate data distribution statistics, and sophisticated optimizer algorithms for processing them, are needed to produce accurate cardinality estimates. A query execution plan selected with inaccurate cardinality estimates can perform several orders of magnitude slower than one selected with accurate estimates.

In “Stinger”, the DB2[®] optimizer can exploit distribution statistics from materialized query tables as well as detect correlation from column group statistics to ensure the most accurate cardinality estimates. This helps DB2 select a more efficient access plan for executing the query and improves query performance.

Related concepts:

- “Column correlation for multiple predicates” in the *Administration Guide: Performance*

Related tasks:

- “Collecting distribution statistics for specific columns” in the *Administration Guide: Performance*

Related reference:

- “RUNSTATS Command” in the *Command Reference*

Reduction of the time and space required to optimize complex queries

The size and complexity of tables within databases is increasing. Longer and more complex queries are being used. With the advent of decision support systems it is not unusual to see queries involving 20 to 30 tables. The time required to optimize these complex queries can often become the biggest part of the total response time. In addition, the optimizer can often run out of space while optimizing a complex query, causing the optimizer to have to reoptimize the query with a more space efficient, but less optimal, search algorithm. Use of a less optimal search algorithm often leads to the selection of a poorly performing query execution plan.

Enhancements to the optimizer focused on these areas improves the results for complex queries.

There are no changes you need to make to your queries to take advantage of these improvements.

Related concepts:

- “Query tuning guidelines” in the *Administration Guide: Performance*

Automatic statistics profiling

The RUNSTATS utility provides an option to register and use a statistics profile, which is a set of options that specify which statistics are to be collected on a particular table, for example, table statistics, index statistics, or distribution statistics.

Statistics profiles can now be generated automatically by the DB2[®] automatic statistics profiling feature. When this feature is enabled, information about database activity is collected and stored in a query feedback warehouse. Based on this data, a statistics profile is generated. Enabling this feature can alleviate the problem of uncertainty about which statistics are relevant to a particular workload and permits the collection of the minimal set of statistics to provide optimal database workload performance.

This feature can be used together with the automatic statistics collection feature to decide which statistics to collect on a particular table based on past table activity.

Related concepts:

- “Guidelines for collecting and updating statistics” in the *Administration Guide: Performance*
- “Collecting statistics using a statistics profile” in the *Administration Guide: Performance*
- “Automating database maintenance” on page 18

Automating database maintenance

You can now automate database maintenance activities to run only when they are needed using the Configure Automatic Maintenance wizard. The activities you can configure for automatic maintenance are database backup, data defragmentation (table or index reorganization), and data access optimization (running statistics).

Performing maintenance activities on your databases is essential to ensure that they are optimized for performance and recoverability.

The DB2[®] optimizer uses catalog statistics to determine the most efficient access plan for any given query. Having out-of-date or incomplete statistics for a table or an index could lead the optimizer to select a plan that is much more inefficient than other alternatives, slowing down query execution. However, deciding which statistics to collect for a given workload is complex, and keeping these statistics up to date can be time-consuming.

With automatic statistics collection, part of DB2's Automatic Table Maintenance feature, you can let DB2 determine which statistics are required by your workload and which statistics need to be updated. With automatic statistics collection enabled, DB2 will automatically run the RUNSTATS utility in the background as required to ensure that the correct statistics are collected and maintained.

Use the Configure Automatic Maintenance wizard to specify your maintenance objectives, including when automatic maintenance can run. DB2 uses the objectives that you have specified to determine whether the maintenance activities need to be done. Then DB2 runs only the required maintenance activities during the next available maintenance window (a user-defined time period for the running of automatic maintenance activities).

Related concepts:

- “Automatic statistics profiling” on page 28
- “Automatic maintenance” in the *Administration Guide: Planning*

Related tasks:

- “Configuring automatic maintenance: Control Center help”

Throttling of RUNSTATS utility

Executing resource-intensive utilities like RUNSTATS can negatively impact overall database performance. However, in order to maintain efficient database operation, statistics must be collected regularly, leaving database administrators with the task of identifying periods of time when the impact of utility execution is most tolerated. In many environments, there are no regular windows of reduced database activity.

Throttling of the RUNSTATS utility limits the amount of resources consumed by the utility, based on the current level of database activity. When database activity is low, the utility runs more aggressively; when database activity increases, the resources allocated to executing RUNSTATS are reduced.

Related tasks:

- “Collecting catalog statistics” in the *Administration Guide: Performance*

Related reference:

- “RUNSTATS Command” in the *Command Reference*

Large page support in buffer pools on DB2 Universal Database for Linux

Large page support in buffer pools is supported on DB2[®] Universal Database (DB2 UDB) for Linux in addition to DB2 UDB for AIX[®] 64-bit. Buffer pools are kept in memory and are not swapped out when using large pages. This improves performance of the DB2 UDB for Linux databases and applications.

Related reference:

- “Performance variables” in the *Administration Guide: Performance*

Multipage file allocation on SMS table spaces enabled by default

A Systems Managed Space (SMS) table space file is extended one extent at a time by default when you create a database (whether partitioned or not), unless you have turned on the DB2_NO_MPFA_FOR_NEW_DB registry variable. If you set this registry variable to “Yes”, and then create a database, any SMS table spaces in the database will allocate one page at a time, as needed.

Related concepts:

- “System managed space” in the *Administration Guide: Planning*
- “Extent size” in the *Administration Guide: Planning*

Automatic setting of prefetch size

The prefetch size for a table space determines the number of pages read from a table space when prefetching is being performed. Because prefetching is a means to improve query performance, setting the correct prefetch size is an important step in performance tuning.

When the prefetch size is not specified for a table space, the DB2[®] uses the value for DFT_PREFETCH_SZ configuration parameter as the default. This parameter can now be set to AUTOMATIC, which allows DB2 to calculate an appropriate prefetch size for a table space based on the extent size, the number of containers, and the number of physical spindles per container. This frees the user from having to determine the appropriate value for the table space prefetch size and from having to remember to reset this value when any containers are added or removed from a table space.

Related reference:

- “dft_prefetch_sz - Default prefetch size configuration parameter” in the *Administration Guide: Performance*

XA timeout sample code

See the CLI sample program dbxamon.c for an example of how to show and roll back indoubt transactions.

The sample program allows you to specify a timeout value for the transaction. If the transaction idle time becomes longer than the timeout value then the sample will rollback the transaction.

Related tasks:

- “Manually resolving indoubt transactions” in the *Administration Guide: Planning*

Related samples:

- “dbxamon.c -- Show and roll back indoubt transactions.”

Chapter 10. Availability enhancements

Availability enhancements summary

The following availability enhancements have been added for this release:

Related concepts:

- “High availability disaster recovery (HADR)” on page 33
- “XML Extender platform support for HP-UX 32-bit and 64-bit PA-RISC” on page 35
- “Automatic client reroute” on page 35
- “Utility for demigrating database directories from “Stinger” to version 8.1” on page 35

High availability disaster recovery (HADR)

DB2[®] high availability disaster recovery (HADR) is a data replication feature that provides a high availability solution for both partial and complete site failures. HADR protects against data loss by replicating data changes from a source database, called the primary, to a target database, called the standby.

A partial site failure can be caused by a hardware, network, or software (DB2 or operating system) failure. Without HADR, the database management system (DBMS) server or the machine where the database resides has to be rebooted. The length of time it takes to restart the database and the machine where it resides is unpredictable. It can take several minutes before the database is brought back to a consistent state and made available. With HADR, the standby database can take over in seconds. Further, you can redirect the clients that were using the old primary database to the standby database (new primary database) by using automatic client reroute or retry logic in the application.

A complete site failure can occur when a disaster, such as a fire, causes the entire site to be destroyed. Since HADR uses TCP/IP for communication between the primary and standby databases, the databases can be situated in different locations. If a disaster occurs at the primary site, data availability is maintained by having the remote standby database take over as the primary database with full DB2 functionality. After a takeover operation occurs, you can bring the original primary database back up and return it to its status of primary database; this is known as failback.

With HADR, you can choose the level of protection you want from potential loss of data by specifying one of three synchronization modes: synchronous, near synchronous, or asynchronous.

HADR allows the standby database to take over as the primary database with full DB2 functionality. It is also possible for the original primary database to be brought back up and returned to its status of primary database.

When a failure occurs on the primary database, you can initiate a takeover operation on the standby database which then becomes the new primary. Since the standby database is already online, failover can be accomplished very quickly resulting in minimal down time.

Once the failed old primary is repaired, it can rejoin the HADR pair as a standby database if the two copies of the database can be made consistent. After the original primary database is reintegrated into the HADR pair as the standby database, a failback operation can be performed so that the original primary database is once again the primary database.

Set Up High Availability Disaster Recovery (HADR) wizard:

Use the Set Up High Availability Disaster Recovery (HADR) Databases wizard to set up and configure primary and standby databases for high availability disaster recovery.

The wizard guides you through the following tasks:

- Identifying the HADR pair, specifically the primary and standby databases and instances.
- Setting the logging-related configuration parameters to prepare the primary database for log shipping.
- Updating the HADR-related database configuration parameters on the primary and standby databases.
- Taking a backup of the primary database, copying the backup image to the standby database, and restoring the backup image to a new or existing standby database.
- Moving database objects that are not included in the backup image.
- Updating the services file with the service name and port number pair used by the HADR agent.
- Configuring databases for automatic client reroute.
- Starting HADR as soon as the configuration is completed, or postponing starting HADR to a later time.

Manage High Availability Disaster Recovery window:

Use the Manage High Availability Disaster Recovery window to configure and to check the status of your HADR system. This window will notify you if there are any problems with the HADR configuration. From this window you can perform the following tasks:

- Start HADR on one or both of the databases.
- Stop HADR on one or both of the databases.
- Initialize a takeover operation.
- Check or alter database configuration parameters related to HADR.
- Catalog an HADR database (if required).

Related concepts:

- “High availability disaster recovery overview” in the *Data Recovery and High Availability Guide and Reference*

Related reference:

- “High availability disaster recovery monitor elements” in the *System Monitor Guide and Reference*

XML Extender platform support for HP-UX 32-bit and 64-bit PA-RISC

DB2® XML Extender is now supported by HP-UX Version 11 in both 32-bit and 64-bit PA-RISC platforms. This is the first 64-bit platform supported by XML Extender. The XML Extender MQSeries® functions are not yet available for use on the HP-UX platform.

Related concepts:

- “Introduction to XML Extender” in the *DB2 XML Extender Administration and Programming*

Automatic client reroute

The automatic client reroute feature allows client applications to recover from a loss of communication with the server so that they can continue to work with minimal interruption. After a loss of communication, the client application attempts to reconnect to the server. If this fails, the client is then rerouted to a different server. You can specify an alternate location through the command line processor (CLP), by invoking an application programming interface (API), or when adding a database using the Control Center or the advanced view of the Configuration Assistant.

You can use automatic client reroute with high availability disaster recovery (HADR) to make client applications connect to the new primary database after a takeover operation.

Related concepts:

- “Automatic client reroute implementation” in the *Administration Guide: Implementation*
- “Automatic client reroute and high availability disaster recovery” in the *Data Recovery and High Availability Guide and Reference*
- “Utility for demigrating database directories from “Stinger” to version 8.1” on page 35

Utility for demigrating database directories from “Stinger” to version 8.1

As part of the automatic client reroute feature, the format of database directory files has been changed. The updating of the database directory files is handled automatically when migrating from version 8.1 to “Stinger”. However, to fall back from “Stinger” to version 8.1, you must use the **db2demigdbd** utility to demigrate your database directory files. This utility converts the database directory files from the “Stinger” format to the version 8.1 format.

Related concepts:

- “Automatic client reroute” on page 35

Related reference:

- “db2demigdbd - Demigrate Database Directory Files Command” in the *Command Reference*

Chapter 11. Usability enhancements

Usability enhancements summary

The following usability enhancements have been added for this release:

Related concepts:

- “Command Editor” on page 37
- “Control Center enhancements” on page 37
- “Utility progress monitoring” on page 38
- “Changing generated column properties without having to recreate the table” on page 39

Command Editor

In "Stinger" the Command Center is replaced by the Command Editor. The Command Editor includes all the function previously available with the Command Center, but provides these functions in a simplified interface.

Use the Command Editor to generate, edit, execute, and manipulate SQL statements, IMS[™] commands, and DB2[®] commands; to work with the resulting output; and to view a graphical representation of the access plan for explained SQL statements. You can execute commands and SQL statements on DB2 Universal Database[™] databases, z/OS[™] and OS/390[®] systems and subsystems, and IMSplexes.

The Command Editor is available as two different interfaces. It can be opened as part of the Control Center (embedded) or in a stand-alone view. Both versions offer the same set of functions and both allow you to open multiple Command Editors.

- Using the Command Editor within the Control Center allows you to control the number of windows open on your desktop. The Command Editor opens as a tile inside the Control Center.
- Using the stand-alone Command Editor allows you to execute commands and SQL statements without opening the Control Center.

When building IMS Commands, you can now use the Command Editor's syntax assist feature, which prompts you with valid keyword for the command you are building. In addition, you can define shortcuts for IMS commands and issue commands directly from the results of previous commands.

Control Center enhancements

The "Stinger" Control Center is available in three different views:

- **Basic.** This view provides core DB2[®] UDB functionality, which includes the essential objects, such as databases, tables, and stored procedures.
- **Advanced.** This view displays all objects and actions available in the Control Center. This is the view that you should select if you are working in an enterprise environment and want to connect to DB2 for z/OS[™] or IMS[™].
- **Custom.** This view gives you the ability to tailor the object tree and the object actions to your specific needs.

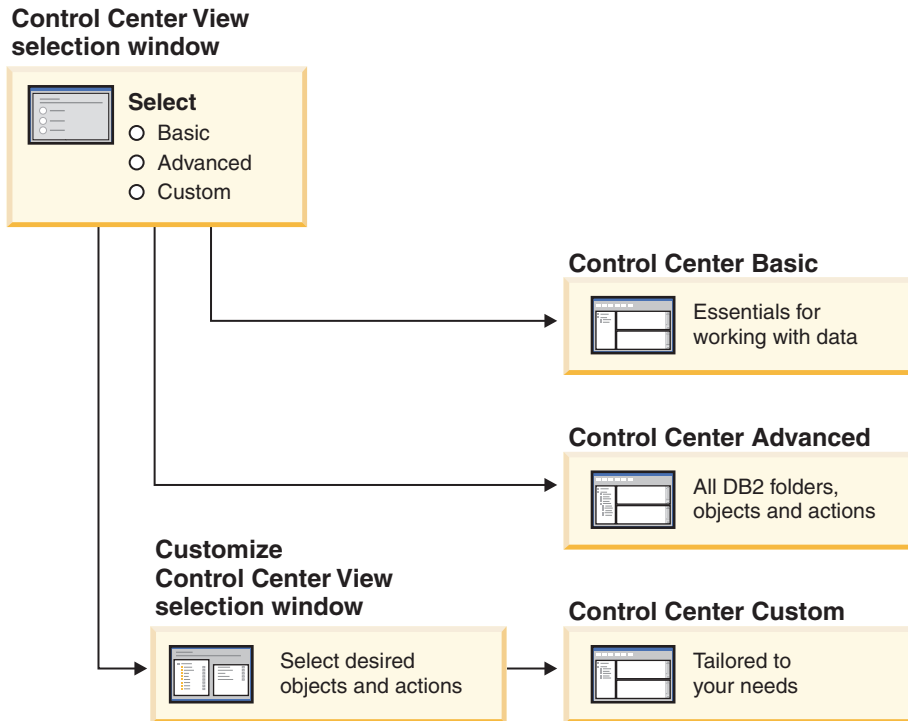


Figure 1. Control Center views

You can change your view if your needs change. Tailoring the Control Center to display the folders and folder objects that you work with regularly (the objects within a folder are called folder objects).

The Control Center now has a third pane called the Object details pane. It displays information about the folder or folder object that you have selected in the Control Center's object tree or contents pane. It also lists some of the actions that you can perform on the selected folder or folder object.

The Control Center now supports additional types of IMSplex members and commands. You can now query or update databases, areas, data groups, or online reorganizations. In addition, IMS command windows make it easier to issue commands.

Utility progress monitoring

You can now monitor the progress of the following operations:

- Backup
- Restore
- Crash recovery
- Load
- Rebalance

Use the SHOW DETAIL option of the LIST UTILITY command to see the status of the operation.

If you use the Control Center, you can also monitor the progress of running utilities in the Utility Status Manager window.

Related tasks:

- “Using Load” in the *Data Movement Utilities Guide and Reference*
- “Monitoring the progress of backup, restore and recovery operations” in the *Data Recovery and High Availability Guide and Reference*
- “Monitoring the status of utilities running on your system: Control Center help”

Related reference:

- “LIST UTILITIES Command” in the *Command Reference*

Changing generated column properties without having to recreate the table

You can now use the ALTER COLUMN clause in the ALTER TABLE statement to alter the various ways that a column value can be generated. More specifically, you can:

- Add the generated expression attribute to an existing non-generated column
- Drop the generated expression attribute from an existing generated expression column
- Add the identity attribute to an existing non-identity column
- Drop the identity attribute from an existing identity column
- Alter a generated column from GENERATED ALWAYS to GENERATED BY DEFAULT, or from GENERATED BY DEFAULT to GENERATED ALWAYS
- Drop the default attribute from a user-defined default column

Previously, generated expression or identity attributes could be assigned to a column only when it was created. Although the expression itself could be changed later, once a generated column was created, it could not be changed to a non-generated column without dropping and then recreating the table. Before dropping the table, you had to export the data in the table and then reload that data into the recreated table with a redefined column.

Related reference:

- “ALTER TABLE statement” in the *SQL Reference, Volume 2*

Chapter 12. SQL enhancements

SQL enhancements summary

The following SQL enhancements have been added for this release:

Related concepts:

- “Procedure invocation in a trigger body in DB2 UDB Enterprise Edition” on page 46
- “SQL procedures no longer require a C or C++ compiler” on page 42
- “Bind option REOPT enables reoptimization of SQL statements” on page 43
- “A new SQL/XML publishing function: XMLSERIALIZE” on page 43
- “Unicode tables in a non-Unicode database supported” on page 44
- “SQL statement size limit increased to 2 MB” on page 17
- “Specifying a lock wait mode strategy” on page 27
- “Improved query execution plans through better cardinality estimation” on page 27
- “Column default values can be altered” on page 21
- “Function XMLNAMESPACES new for “Stinger”” on page 44
- “Changing generated column properties without having to recreate the table” on page 39

Column default values can be altered

There are times when the default values of a column in a database table must be changed. The ALTER TABLE SQL statement allows you to set the default clause of a column in a table. The same default clause restrictions as documented under the CREATE TABLE SQL statement are enforced. Note that generated columns cannot have their default value altered by this statement.

Related reference:

- “ALTER TABLE statement” in the *SQL Reference, Volume 2*

Changing generated column properties without having to recreate the table

You can now use the ALTER COLUMN clause in the ALTER TABLE statement to alter the various ways that a column value can be generated. More specifically, you can:

- Add the generated expression attribute to an existing non-generated column
- Drop the generated expression attribute from an existing generated expression column
- Add the identity attribute to an existing non-identity column
- Drop the identity attribute from an existing identity column
- Alter a generated column from GENERATED ALWAYS to GENERATED BY DEFAULT, or from GENERATED BY DEFAULT to GENERATED ALWAYS
- Drop the default attribute from a user-defined default column

Previously, generated expression or identity attributes could be assigned to a column only when it was created. Although the expression itself could be changed later, once a generated column was created, it could not be changed to a non-generated column without dropping and then recreating the table. Before dropping the table, you had to export the data in the table and then reload that data into the recreated table with a redefined column.

Related reference:

- “ALTER TABLE statement” in the *SQL Reference, Volume 2*

SQL statement size limit increased to 2 MB

DB2® Universal Database, "Stinger", provides an increased SQL statement size limit of 2 MB, which can limit the size of some created database objects (for example, CREATE PROCEDURE or CREATE TRIGGER statements). The new 2 MB limit allows you to use large statements when your application logic is in stored procedures or triggers.

The increase in the statement size allows you to migrate a trigger or stored procedure statement from another RDBMS to DB2 UDB.

The new statement limit allows you to record auditing context records that have statement text up to 2 MB.

Related reference:

- “Audit record layout for CONTEXT events” in the *Administration Guide: Implementation*

SQL procedures no longer require a C or C++ compiler

Beginning with DB2® UDB the creation of SQL procedures does not require a C or C++ compiler on the server; therefore C or C++ compiler setup is not required. When you create an SQL procedure its procedural statements are converted to a native representation that is stored in the database catalogs, as is done with other SQL statements. When an SQL procedure is called, the native representation is loaded from the catalogs and the DB2 engine executes the procedure.

Related tasks:

- “Customizing precompile and bind options for SQL procedures” in the *Application Development Guide: Building and Running Applications*
- “Creating SQL procedures” in the *Application Development Guide: Building and Running Applications*
- “Calling procedures from the Command Line Processor (CLP)” in the *Application Development Guide: Building and Running Applications*
- “Calling SQL procedures with client applications” in the *Application Development Guide: Building and Running Applications*
- “Rebinding SQL procedures” in the *Application Development Guide: Building and Running Applications*

Bind option REOPT enables reoptimization of SQL statements

The bind option REOPT enables reoptimization of static and dynamic SQL statements. Specifically, the access path for an SQL statement containing host variables, special registers, or parameter markers is optimized using the values of these variables, rather than the default estimates determined by the compiler. This optimization takes place at query execution time when these values are available.

You can set the bind option REOPT to one of the following three values:

NONE

The values of any host variables, parameter markers, or special registers in an SQL statement are not used to optimize this statement's access path. The default estimates for these variables are used instead. This is the default behavior.

ONCE When the statement is first executed, the access path for a given SQL statement is optimized using the actual values for any host variables, parameter markers, or special registers. This access path will be used in all subsequent executions of the statement.

ALWAYS

Each time that a statement containing host variables, parameter markers, or special registers is executed, the values of these variables are used to optimize the access path for the statement.

Related concepts:

- "Effects of REOPT on static SQL" in the *Application Development Guide: Programming Client Applications*
- "Effects of REOPT on dynamic SQL" in the *Application Development Guide: Programming Client Applications*

A new SQL/XML publishing function: XMLSERIALIZE

A new standard SQL/XML function, XMLSERIALIZE (with the CONTENT option), allows you to convert an XML data type value into a result string data type that is appropriate for the length of the XML output.

The XML data type is an internal representation of XML and can be used only as input to functions that accept this data type as input. XML is a transient data type that cannot be stored in the database or returned to an application. Until "Stinger", the only supported operation to convert an XML data type value to a string data type value was by serialization using the XML2CLOB function.

Serialization is the inverse operation of parsing; it is the process of converting a parsed XML value into a textual XML value. XMLSERIALIZE converts an XML expression into an SQL string value that, in turn, can be bound out to host character variables.

With XMLSERIALIZE, you can specify a result type like CHAR or VARCHAR, which might be more appropriate and result in better performance than CLOB.

Related reference:

- "Expressions" in the *SQL Reference, Volume 1*
- "XML values" in the *SQL Reference, Volume 1*

Function XMLNAMESPACES new for "Stinger"

An XML namespace is a collection of names that is identified by a uniform reference identifier (URI). Namespaces are used in XML documents as element types and attribute names.

DB2® Universal Database includes several XML functions as part of the SQL/XML language specifications. Among these are XMLFOREST and XMLELEMENT. A new XML function, XMLNAMESPACES, provides XML namespace declarations within the SQL/XML publishing functions XMLELEMENT and XMLFOREST.

Related reference:

- "Expressions" in the *SQL Reference, Volume 1*

Unicode tables in a non-Unicode database supported

You can now create Unicode tables in non-Unicode databases by invoking the CREATE TABLE statement with the CCSID UNICODE clause.

When you create tables in a non-Unicode database, the default value of CCSID is ASCII. When you create tables in a Unicode database, default value of CCSID is UNICODE. You cannot define a table with CCSID ASCII in a Unicode database, which means you cannot create a non-Unicode table in a Unicode database.

To enable support for Unicode tables in a non-Unicode database, you need to activate the new database configuration parameter, *alt_collate*. This parameter specifies the collating sequence for Unicode tables, which presently can be only IDENTITY_16BIT.

In conjunction with support for Unicode tables in non-Unicode databases, DB2® features the addition of a new code page attribute: section code page. The database manager determines code page attributes for all character strings when an application is bound to a database. A section code page is the code page under which an SQL statement runs. The section code page is the database code page, unless the statement references:

- A table that is created with the CCSID UNICODE clause in a non-Unicode database.
- A table function that is defined with PARAMETER CCSID UNICODE in a non-Unicode database.

Related concepts:

- "Derivation of code page values" in the *Application Development Guide: Programming Client Applications*

Related reference:

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

Specifying a lock wait mode strategy

An individual session can now specify a lock wait mode strategy, which is used when the session requires a lock that it cannot obtain immediately. The strategy indicates whether the session will:

- Return an SQLCODE and SQLSTATE when it cannot obtain a lock
- Wait indefinitely for a lock

- Wait a specified amount of time for a lock
- Use the value of the *locktimeout* database configuration parameter when waiting for a lock

The lock wait mode strategy is specified through the new SET CURRENT LOCK TIMEOUT statement, which changes the value of the CURRENT LOCK TIMEOUT special register. The CURRENT LOCK TIMEOUT special register specifies the number of seconds to wait for a lock before returning an error indicating that a lock cannot be obtained.

Traditional locking approaches can result in applications blocking each other. This happens when one application must wait for another application to release its lock. Strategies to deal with the impact of such blocking usually provide a mechanism to specify the maximum acceptable duration of the block. That is the amount of time that an application will wait prior to returning without a lock. Previously, this was only possible at the database level by changing the value of the *locktimeout* database configuration parameter.

Whereas the value of the *locktimeout* parameter applies to all locks, the lock types that are impacted by this new function include row, table, index key, and multidimensional clustering (MDC) block locks.

Related reference:

- “locktimeout - Lock timeout configuration parameter” in the *Administration Guide: Performance*

Improved query execution plans through better cardinality estimation

Cardinality estimation is the process whereby the optimizer determines the number of qualifying rows after filtering operations, such as predicate application or aggregation, are applied. Accurate data distribution statistics, and sophisticated optimizer algorithms for processing them, are needed to produce accurate cardinality estimates. A query execution plan selected with inaccurate cardinality estimates can perform several orders of magnitude slower than one selected with accurate estimates.

In "Stinger", the DB2[®] optimizer can exploit distribution statistics from materialized query tables as well as detect correlation from column group statistics to ensure the most accurate cardinality estimates. This helps DB2 select a more efficient access plan for executing the query and improves query performance.

Related concepts:

- “Column correlation for multiple predicates” in the *Administration Guide: Performance*

Related tasks:

- “Collecting distribution statistics for specific columns” in the *Administration Guide: Performance*

Related reference:

- “RUNSTATS Command” in the *Command Reference*

Procedure invocation in a trigger body in DB2 UDB Enterprise Edition

With DB2® Universal Database Enterprise Edition, you can now invoke procedures from triggers, or any other dynamic compound statement, in single partition environments, by executing a CALL statement that references a procedure within a trigger action. The CALL statement is used to execute SQL and external procedures.

The invocation of a procedure in a trigger-action enables you to encapsulate complex logic in your trigger. The logic can include operations on other tables in the database, or operations external to the database such as sending an e-mail or writing an audit record to a file in the file system of the database server.

Related concepts:

- “Triggered action containing a procedure or function reference” in the *Application Development Guide: Programming Server Applications*

Related tasks:

- “Calling procedures from triggers or SQL routines” in the *Application Development Guide: Programming Server Applications*

Related reference:

- “CREATE TRIGGER statement” in the *SQL Reference, Volume 2*
- “CREATE PROCEDURE (External) statement” in the *SQL Reference, Volume 2*

Chapter 13. Serviceability enhancements

Serviceability enhancements summary

The following serviceability enhancements have been added for this release:

Related concepts:

- “New utility to retrieve statistics from a running DB2 instance or database (similar to 'onstat' for Informix)” on page 47
- “db2diag.log format message enhancements” on page 48
- “Diagnostic log analysis tool for filtering and formatting db2diag.log” on page 48
- “Trace facility improved” on page 49

New utility to retrieve statistics from a running DB2 instance or database (similar to 'onstat' for Informix)

db2pd is a new utility that can be used to retrieve statistics from a running DB2[®] instance or database. It is similar to the **onstat** utility for Informix[®].

The tool can provide a wide range of information useful for troubleshooting and problem determination, performance improvements, and application development design, including:

- locks
- bufferpools
- tablespaces
- containers
- dynamic SQL statements
- agents
- applications
- memory pools and sets
- transactions
- logs
- and others

The tool collects this information without acquiring any latches or using any engine resources. It is therefore possible (and expected) to retrieve information that is changing while **db2pd** is collecting information; hence the data may not be completely accurate. However, two benefits to collecting information without latching include faster retrieval and no competition for engine resources.

Related reference:

- “db2pd - Monitor and Troubleshoot DB2 Command” in the *Command Reference*

db2diag.log format message enhancements

The db2diag.log format has been improved in a number of ways for "Stinger". The log file is now easier to read manually and easier to parse in software. The improvements include:

- Each entry has a complete set of header information
- Message and data portions of the logging record are clearly marked, making the logging record easier to use and to understand
- Timestamps (with time zone) will be included for each record
- Each field has a clear field name in upper-case
- Header and message field line lengths are restricted to 80 characters
- New fields have been added, most notably a severity-level field to help you find the most important entries

Other changes have been made as well, such as changing the database field name to 'DB'.

Related concepts:

- "Diagnostic log analysis tool for filtering and formatting db2diag.log" on page 48

Related reference:

- "db2diag - db2diag.log analysis tool Command" in the *Command Reference*

Diagnostic log analysis tool for filtering and formatting db2diag.log

A new tool for filtering and formatting db2diag.log files (**db2diag**) is now available. You can use this tool to filter diagnostic log files, which use a new message format for "Stinger"

This command line tool has a wide range of options that you can use to modify the output to suit your needs. Among other options, you can indicate which fields to display, use a 'grep' like filter to reduce the number of records, and have the empty fields omitted.

Command line options include:

- **db2diag -help** provides a short description of the options
- **db2diag -h brief** provides descriptions for all options without examples
- **db2diag -h notes** provides usage notes and restrictions
- **db2diag -h examples** provides a small set of examples to get started
- **db2diag -h tutorial** provides examples for all available options
- **db2diag -h all** provides the most complete list of options

Related concepts:

- "db2diag.log format message enhancements" on page 48

Related reference:

- "db2diag - db2diag.log analysis tool Command" in the *Command Reference*

Trace facility improved

The DB2® trace facilities and the **db2trc** command have been improved for "Stinger". The layout of the formatted trace record has been modified to reflect the new, wider range of trace types.

Related concepts:

- "New utility to retrieve statistics from a running DB2 instance or database (similar to 'onstat' for Informix)" on page 47

Related reference:

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

Chapter 14. Data Warehouse Center and Warehouse Manager enhancements

Data Warehouse Center and Warehouse Manager ETL enhancements summary

The following Data Warehouse Center and Warehouse Manager ETL enhancements have been added for this release:

Related concepts:

- “Warehouse Manager ETL (Extract, Transform, and Load) supports Unicode warehouse control databases” on page 51

Warehouse Manager ETL (Extract, Transform, and Load) supports Unicode warehouse control databases

Starting with "Stinger" in the Data Warehouse Center, the warehouse control database must be a UTF-8 (Unicode Transformation Format, or Unicode) database. This requirement provides expanded language support for the Data Warehouse Center. If you try to log on to the Data Warehouse Center using a control database that is not in Unicode format, you will receive an error message that you cannot log on. You can use the Warehouse Control Database Management tool to migrate the metadata from a specified database into a new Unicode database.

Related concepts:

- “Data warehouse objects” in the *Data Warehouse Center Administration Guide*

Chapter 15. Application development enhancements

Application development enhancements summary

The following application development enhancements have been added for this release:

Related concepts:

- "Procedure invocation in a trigger body in DB2 UDB Enterprise Edition" on page 46
- "SQL procedures no longer require a C or C++ compiler" on page 42
- "Simplified association of packages with an application using CURRENT PACKAGE PATH special register" on page 62
- "Bind option REOPT enables reoptimization of SQL statements" on page 43
- "Unicode tables in a non-Unicode database supported" on page 44
- "DB2 Universal JDBC Driver improvements" on page 54
- ".NET common language runtime procedures and UDFs" on page 58
- "IBM DB2 Development Add-In for Microsoft Visual Studio .NET enhancements" on page 58
- "DB2 WebSphere MQ user-defined function 64-bit support" on page 63
- "DB2 Run-Time Client Lite" on page 59
- "JDK 1.4 supported by DB2 Universal Database" on page 53
- "Development Center supports Java stored procedures on iSeries" on page 63
- "Development Center no longer requires C compiler for building SQL stored procedures" on page 63
- "DB2 .NET Data Provider enhancements" on page 57
- "OLE DB and ODBC .NET data providers" on page 57
- "Development Center supports SQL stored procedure debugging for UNIX 64-bit platforms" on page 63
- "Function XMLNAMESPACES new for "Stinger"" on page 44
- "Nested savepoints allow for greater control of rollbacks" on page 61
- "Development Center stored procedure and debugging enhancements" on page 62
- "XA timeout sample code" on page 30

JDK 1.4 supported by DB2 Universal Database

DB2[®] Universal Database (DB2 UDB), "Stinger" supports JDK 1.4 on all supported workstation operating system environments, except on AIX[®] 4.3 and Linux IA64, where IBM[®] Developer Kit 1.3.1 is still required. DB2 UDB for Solaris will ship with JDK 1.4.2. DB2 UDB for all other supported workstation operating systems will ship with JDK 1.4.1 (or the equivalent vendor-named Java[™] developer kit).

When you install DB2 UDB, "Stinger", the latest supported version of the Java developer kit will also be installed if it is not already installed, unless the DB2 UDB installation is an update of a previous DB2 UDB, Version 8 installation. If you are updating a previous DB2 UDB, Version 8 installation, you must install the Java developer kit from the CD.

Related reference:

- “AIX supported development software” in the *Application Development Guide: Building and Running Applications*
- “HP-UX supported development software” in the *Application Development Guide: Building and Running Applications*
- “Linux supported development software” in the *Application Development Guide: Building and Running Applications*
- “Solaris supported development software” in the *Application Development Guide: Building and Running Applications*
- “Windows supported development software” in the *Application Development Guide: Building and Running Applications*

DB2 Universal JDBC Driver improvements

There are numerous improvements to the DB2[®] Universal JDBC Driver for "Stinger" including distributed transaction support and JDBC 3.0 compliance. Other new features increase the flexibility of iterator use in SQLJ and enhance interactivity with host and iSeries[™] databases.

Distributed transaction support:

Support for distributed transaction processing that conforms to the XA specification. This support implements the Java[™] 2 Platform, Enterprise Edition (J2EE) Java Transaction Service (JTS) and Java Transaction API (JTA) specifications (Universal Type 2 Connectivity on DB2 UDB for Linux, UNIX[®], and Windows[®] systems only).

JDBC 3.0 compliance:

The following JDBC 3.0 methods are now part of the DB2 Universal JDBC Driver for "Stinger":

- `Blob.setBytes`
- `Blob.setBinaryStream`
- `Blob.truncate`
- `Clob.setString`
- `Clob.setAsciiStream`
- `Clob.setUnicodeStream`
- `Clob.setCharacterStream`
- `Connection.createStatement(int resultSetType, int resultSetConcurrency, int resultSetHoldability)`
- `Connection.prepareStatement(String sql, int autoGeneratedKeys)`
- `Connection.setSavepoint`
- `Connection.releaseSavepoint`
- `Statement.executeUpdate(String sql, int autoGeneratedKeys)`
- `Statement.execute(String sql, int autoGeneratedKeys)`
- `Statement.getMoreResults(java.sql.Statement.CLOSE_CURRENT_RESULT)`
- `Statement.getMoreResults(java.sql.Statement.KEEP_CURRENT_RESULT)`
- `Statement.getMoreResults(java.sql.Statement.CLOSE_ALL_RESULTS)`

DB2BaseDataSource method: getDB2SystemMonitor:

The method `DB2BaseDataSource.getDB2SystemMonitor` returns an object of the `DB2SystemMonitor` class. This class, for which an interface was added to the DB2 Universal Driver, facilitates database system monitoring from a JDBC application. Each DB2 Universal JDBC Driver connection can have a single system monitor.

DB2BaseDataSource property: `activeServerListJNDIName`:

The property `DB2BaseDataSource.activeServerListJNDIName` identifies a JNDI reference to a `DB2ActiveServerList` instance in a JNDI repository of alternate server information. (`DB2ActiveServerList` is a new class, described later in this topic.) If the value of `activeServerListJNDIName` is not null, connections can failover to an alternate server that is specified in the `DB2ActiveServerList` instance that is referenced by the value. If `activeServerListJNDIName` is null, connections do not failover using alternate server information from a JNDI repository.

DB2BaseDataSource property: `currentFunctionPath`:

The property `DB2BaseDataSource.currentFunctionPath` specifies the SQL path that is used to resolve unqualified data type names and function names in SQL statements that are in JDBC programs. The data type of this property is `String`. For a DB2 UDB for Linux, UNIX and Windows server, the maximum length is 254 bytes. The value is a comma-separated list of schema names. Those names can be ordinary or delimited identifiers.

DB2BaseDataSource property: `currentLockTimeout`:

The property `DB2BaseDataSource.currentLockTimeout` directs DB2 UDB for Linux, UNIX and Windows servers to wait indefinitely for a lock or to wait for the specified number of seconds for a lock when the lock cannot be obtained immediately. The data type of this property is `int`. A value of zero means no wait. A value of -1 means to wait indefinitely. A positive integer indicates the number of seconds to wait for a lock.

DB2BaseDataSource property: `cursorSensitivity`:

The property `DB2BaseDataSource.cursorSensitivity` specifies whether the `java.sql.ResultSet.TYPE_SCROLL_SENSITIVE` value for a JDBC `ResultSet` maps to the `SENSITIVE DYNAMIC` attribute or the `SENSITIVE STATIC` attribute for the underlying DB2 cursor. Possible values are `TYPE_SCROLL_SENSITIVE_STATIC` and `TYPE_SCROLL_SENSITIVE_DYNAMIC`. The default is `TYPE_SCROLL_SENSITIVE_STATIC`. This property is ignored for database servers that do not support sensitive dynamic scrollable cursors.

DB2BaseDataSource property: `jdbcCollection`:

The property `DB2BaseDataSource.jdbcCollection` specifies the collection ID for the packages that are used by an instance of the DB2 Universal JDBC Driver at run time. The data type of `jdbcCollection` is `String`. The default is `NULLID`. This property is used with the `DB2Binder -collection` option. The `DB2Binder` utility must have previously bound DB2 Universal JDBC Driver packages at the server using a `-collection` value that matches the `jdbcCollection` value. The `jdbcCollection` setting does not determine the collection that is used for SQLJ applications. For SQLJ, the collection is determined by the `-collection` option of the SQLJ customizer.

DB2ActiveServerList class:

This new class contains the following methods:

getAlternatePortNumber

Retrieves the port numbers that are associated with the alternate DB2 UDB servers.

getAlternateServerName

Retrieves an array that contains the names of the alternate DB2 UDB servers. These values are IP addresses or DNS server names.

setAlternatePortNumber

Sets the port numbers that are associated with the alternate DB2 UDB servers.

setAlternateServerName

Sets the alternate server names for DB2 UDB servers. These values are IP addresses or DNS server names.

Multiple instances of an iterator can be open in an SQLJ application:

Multiple instances of an iterator can be open concurrently in a single SQLJ application. One application for this ability is to open several instances of an iterator that uses host expressions. Each instance can use a different set of host expression values.

Multiple iterators can be open on the same DB2 table in an SQLJ application:

Multiple iterators on the same DB2 table can be open concurrently. You can perform one operation on a table using one iterator while you perform a different operation on the same table using another iterator.

ROWIDs in SQLJ with the DB2 Universal JDBC Driver allow for unique row identification:

DB2 UDB for z/OS™ and DB2 UDB for iSeries support the ROWID data type for a column in a DB2 table. A ROWID is a value that uniquely identifies a row in a table.

The DB2 Universal JDBC Driver provides the DB2-only class: `com.ibm.db2.jcc.DB2RowID` that you can use in iterators and in CALL statement parameters. For an iterator, you can also use the `byte[]` object type to retrieve ROWID values.

Related concepts:

- “Introduction to Java application support” in the *Application Development Guide: Programming Client Applications*
- “Multiple open instances of an iterator in an SQLJ application” in the *Application Development Guide: Programming Client Applications*
- “Multiple open iterators for the same SQL statement in an SQLJ application” in the *Application Development Guide: Programming Client Applications*
- “ROWIDs in SQLJ with the DB2 Universal JDBC Driver” in the *Application Development Guide: Programming Client Applications*

Related reference:

- “Properties for the DB2 Universal JDBC Driver” in the *Application Development Guide: Programming Client Applications*

- “Summary of DB2 Universal JDBC Driver extensions to JDBC” in the *Application Development Guide: Programming Client Applications*

OLE DB and ODBC .NET data providers

In addition to the DB2[®] .NET Data Provider, DB2 includes an OLE DB .NET Data Provider and a ODBC .NET Data Provider.

The OLE DB .NET Data Provider uses the IBM[®] DB2 OLE DB Driver. Therefore, the connection string keywords supported by the OLE DB .NET Data Provider are the same as those supported by the IBM OLE DB Provider for DB2. Also, the OLE DB .NET Data Provider has the same restrictions as the IBM DB2 OLE DB Provider.

The ODBC .NET Data Provider makes ODBC calls to the DB2 data source using the IBM DB2 CLI Driver. Therefore, the connection string keywords supported by the OLE DB .NET Data Provider are the same as those supported by the IBM DB2 CLI Driver.

Related concepts:

- “OLE DB .NET Data Provider” in the *Application Development Guide: Programming Client Applications*
- “ODBC .NET Data Provider” in the *Application Development Guide: Programming Client Applications*

DB2 .NET Data Provider enhancements

The DB2[®] .NET Data Provider is an extension of the ADO.NET interface that allows .NET applications to access a DB2 database through a secure connection, execute commands, and retrieve results. “Stinger” features the following improvements to the DB2 .NET Data Provider:

- Performance is faster.
- Support for the Microsoft[®] .NET Framework, Version 1.1. This includes support for the following new System.Data features: the `DB2DataReader.HasRows` and `DB2Connection.EnlistDistributedTransaction` methods.
- A new `ConnectionString` keyword to specify the isolation level.
- Four new `DB2Connection` properties enable DBAs to attribute workload to a particular source. The new properties are:
 - `DB2Connection.ClientUser` - client user ID
 - `DB2Connection.ClientWorkStation` - client workstation name
 - `DB2Connection.ClientAccountingInformation` - client accounting string
 - `DB2Connection.ClientApplicationInformation` - client application name
- .NET applications can now access the following database management systems through the DB2 .NET Data Provider:
 - DB2 Universal Database[™] Version 5, Release 1 (or later) for AS/400[®] and iSeries[™], through DB2 Connect[™]
 - DB2 Universal Database Version 7.3 (or later) for VSE & VM, through DB2 Connect

Related concepts:

- “DB2 .NET Data Provider overview” in the *Application Development Guide: Programming Client Applications*

- “DB2 .NET Data Provider” in the *Application Development Guide: Programming Client Applications*

.NET common language runtime procedures and UDFs

DB2[®] Universal Database allows you to create routines (procedures, user-defined functions (UDFs), and methods) in SQL and in other programming languages such as C, C++, Java[™], and OLE. Now you can create procedures and UDFs using any .NET common language runtime (CLR) compatible programming language, including C#, Visual Basic, managed C++, and all other CLR compatible languages.

You can create CLR routines in the same way as other external (non-SQL) routines are created: by executing a CREATE statement that associates a database routine signature with a .NET assembly residing on the database server. You can use the routines to encapsulate commonly used database operations and logic, to extend the functionality of SQL, and to improve the performance of client applications.

Related concepts:

- “Common language runtime (CLR) routines” in the *Application Development Guide: Programming Server Applications*

Related tasks:

- “Creating CLR routines” in the *Application Development Guide: Programming Server Applications*
- “Examples of CLR user-defined functions in C#” in the *Application Development Guide: Programming Server Applications*
- “Examples of CLR procedures in C#” in the *Application Development Guide: Programming Server Applications*
- “Examples of CLR procedures in Visual Basic” in the *Application Development Guide: Programming Server Applications*
- “Examples of CLR user-defined functions in Visual Basic” in the *Application Development Guide: Programming Server Applications*

Related reference:

- “CREATE PROCEDURE (External) statement” in the *SQL Reference, Volume 2*

Related samples:

- “SpCat.db2 -- Drops and creates the procedures implemented in spserver.cs”
- “SpServer.cs -- C# external code implementation of procedures created in spcat.db2”
- “SpCat.db2 -- Drops and creates the procedures implemented in spserver.vb”
- “SpServer.vb -- VB.NET implementation of procedures created in SpCat.db2”

IBM DB2 Development Add-In for Microsoft Visual Studio .NET enhancements

With the IBM[®] DB2[®] Development Add-In for Microsoft[®] Visual Studio .NET, you can now create the following DB2 SQL objects using either a wizard or the built-in code editor:

- tables
- indexes
- views

- triggers

In addition, you can now develop DB2 common language runtime (CLR) procedures from existing methods in a .NET managed language, such as C# and Visual Basic. A DB2 .NET CLR procedure is a user-created DB2 external routine that has its procedure-body implemented by a method in a .NET assembly compiled from source code written in a .NET managed language.

Lastly, you can perform source-level debugging of SQL procedures using the DB2 database project.

DB2 Run-Time Client Lite

The DB2[®] Run-Time Client Lite (DB2 RTCL) is a new installable component that makes it much easier to provide access to DB2 servers from Windows[®]-based applications. The DB2 RTCL is designed to be redistributable by independent software vendors (ISVs) and to be used for application distribution in mass deployment scenarios typical of large enterprises. Similar to the DB2 Run-Time Client component, the new DB2 RTCL provides application interfaces (CLI, ODBC, OLE DB, .NET Data provider, and JDBC) and the network libraries that are required for applications to run.

The main distinguishing features of the DB2 RTCL are:

- A significantly smaller disk footprint
- Shipped as a single executable making it easy to redistribute and deploy
- Windows Installer Merge Module (.msm file) is available which simplifies the integration of the DB2 RTCL code within a larger application
- Terms and conditions pertaining to the redistribution of the DB2 RTCL by ISVs and in-house programmers

Related concepts:

- “DB2 Run-Time Client Lite” in the *Quick Beginnings for DB2 Clients*

A new SQL/XML publishing function: XMLSERIALIZE

A new standard SQL/XML function, XMLSERIALIZE (with the CONTENT option), allows you to convert an XML data type value into a result string data type that is appropriate for the length of the XML output.

The XML data type is an internal representation of XML and can be used only as input to functions that accept this data type as input. XML is a transient data type that cannot be stored in the database or returned to an application. Until "Stinger", the only supported operation to convert an XML data type value to a string data type value was by serialization using the XML2CLOB function.

Serialization is the inverse operation of parsing; it is the process of converting a parsed XML value into a textual XML value. XMLSERIALIZE converts an XML expression into an SQL string value that, in turn, can be bound out to host character variables.

With XMLSERIALIZE, you can specify a result type like CHAR or VARCHAR, which might be more appropriate and result in better performance than CLOB.

Related reference:

- “Expressions” in the *SQL Reference, Volume 1*
- “XML values” in the *SQL Reference, Volume 1*

Procedure invocation in a trigger body in DB2 UDB Enterprise Edition

With DB2® Universal Database Enterprise Edition, you can now invoke procedures from triggers, or any other dynamic compound statement, in single partition environments, by executing a CALL statement that references a procedure within a trigger action. The CALL statement is used to execute SQL and external procedures.

The invocation of a procedure in a trigger-action enables you to encapsulate complex logic in your trigger. The logic can include operations on other tables in the database, or operations external to the database such as sending an e-mail or writing an audit record to a file in the file system of the database server.

Related concepts:

- “Triggered action containing a procedure or function reference” in the *Application Development Guide: Programming Server Applications*

Related tasks:

- “Calling procedures from triggers or SQL routines” in the *Application Development Guide: Programming Server Applications*

Related reference:

- “CREATE TRIGGER statement” in the *SQL Reference, Volume 2*
- “CREATE PROCEDURE (External) statement” in the *SQL Reference, Volume 2*

Unicode tables in a non-Unicode database supported

You can now create Unicode tables in non-Unicode databases by invoking the CREATE TABLE statement with the CCSID UNICODE clause.

When you create tables in a non-Unicode database, the default value of CCSID is ASCII. When you create tables in a Unicode database, default value of CCSID is UNICODE. You cannot define a table with CCSID ASCII in a Unicode database, which means you cannot create a non-Unicode table in a Unicode database.

To enable support for Unicode tables in a non-Unicode database, you need to activate the new database configuration parameter, *alt_collate*. This parameter specifies the collating sequence for Unicode tables, which presently can be only IDENTITY_16BIT.

In conjunction with support for Unicode tables in non-Unicode databases, DB2® features the addition of a new code page attribute: section code page. The database manager determines code page attributes for all character strings when an application is bound to a database. A section code page is the code page under which an SQL statement runs. The section code page is the database code page, unless the statement references:

- A table that is created with the CCSID UNICODE clause in a non-Unicode database.
- A table function that is defined with PARAMETER CCSID UNICODE in a non-Unicode database.

Related concepts:

- “Derivation of code page values” in the *Application Development Guide: Programming Client Applications*

Related reference:

- “CREATE TABLE statement” in the *SQL Reference, Volume 2*

Nested savepoints allow for greater control of rollbacks

DB2[®] Universal Database (DB2 UDB) now supports the nesting of savepoints. This means you can set up a savepoint within another savepoint. DB2 UDB enables you to set as many levels of nested savepoints as you require.

With nested savepoints, an application can have multiple levels of savepoints active simultaneously, and the application can roll back to any active savepoint, as required. A rollback to a particular savepoint statement also releases any active nested savepoints within the savepoint being rolled back.

Related concepts:

- “Transaction management with savepoints” in the *Application Development Guide: Programming Client Applications*
- “Nesting savepoints” in the *Application Development Guide: Programming Client Applications*

Bind option REOPT enables reoptimization of SQL statements

The bind option REOPT enables reoptimization of static and dynamic SQL statements. Specifically, the access path for an SQL statement containing host variables, special registers, or parameter markers is optimized using the values of these variables, rather than the default estimates determined by the compiler. This optimization takes place at query execution time when these values are available.

You can set the bind option REOPT to one of the following three values:

NONE

The values of any host variables, parameter markers, or special registers in an SQL statement are not used to optimize this statement’s access path. The default estimates for these variables are used instead. This is the default behavior.

ONCE When the statement is first executed, the access path for a given SQL statement is optimized using the actual values for any host variables, parameter markers, or special registers. This access path will be used in all subsequent executions of the statement.

ALWAYS

Each time that a statement containing host variables, parameter markers, or special registers is executed, the values of these variables are used to optimize the access path for the statement.

Related concepts:

- “Effects of REOPT on static SQL” in the *Application Development Guide: Programming Client Applications*
- “Effects of REOPT on dynamic SQL” in the *Application Development Guide: Programming Client Applications*

Simplified association of packages with an application using CURRENT PACKAGE PATH special register

The CURRENT PACKAGE PATH special register enables you to specify a list of package qualifiers with which the DB2® server can select a package. This feature is useful when you have multiple variants of a package, that is, when multiple packages share the same name and consistency tokens. Rather than maintain different versions of an application to work with each version of a package, you can use the CURRENT PACKAGE PATH special register to associate a single compiled program with multiple versions of a package.

Related concepts:

- “CURRENT PACKAGE PATH special register for package schemas” in the *Application Development Guide: Programming Client Applications*

SQL procedures no longer require a C or C++ compiler

Beginning with DB2® UDB the creation of SQL procedures does not require a C or C++ compiler on the server; therefore C or C++ compiler setup is not required. When you create an SQL procedure its procedural statements are converted to a native representation that is stored in the database catalogs, as is done with other SQL statements. When an SQL procedure is called, the native representation is loaded from the catalogs and the DB2 engine executes the procedure.

Related tasks:

- “Customizing precompile and bind options for SQL procedures” in the *Application Development Guide: Building and Running Applications*
- “Creating SQL procedures” in the *Application Development Guide: Building and Running Applications*
- “Calling procedures from the Command Line Processor (CLP)” in the *Application Development Guide: Building and Running Applications*
- “Calling SQL procedures with client applications” in the *Application Development Guide: Building and Running Applications*
- “Rebinding SQL procedures” in the *Application Development Guide: Building and Running Applications*

Development Center stored procedure and debugging enhancements

The following stored procedure size limitation and debugging enhancements are available in "Stinger" of the IBM® DB2® Universal Database Development Center:

- Development Center support for DB2 Universal Database™ for z/OS™, Version 8.0 enables you to build Java™ stored procedures up to 10MB in size. Java stored procedures built using the C stored procedure DSNTJSP cannot exceed 32KB.
- A new terminate feature in the debugger enables you to quickly terminate the debugger at any time without waiting for the current debugging session to reach the end of the stored procedure. To use this feature, click the Terminate icon in the debugger.

Development Center no longer requires C compiler for building SQL stored procedures

In IBM® DB2® Universal Database, "Stinger", a C or C++ compiler is no longer required for creating SQL procedures, therefore in the DB2 Development Center, there is no longer a need to configure the C or C++ compiler settings on Windows®, Linux, or UNIX® servers. When an SQL procedure is created the procedural statements are converted to a native representation that is stored in the database catalogs, as is done with other SQL statements.

Default bind options are provided in the SQL build settings, however you can specify different options for each SQL stored procedure.

For all versions of the Development Center earlier than "Stinger", configuration of a C compiler in the stored procedure build settings is still required.

Development Center supports SQL stored procedure debugging for UNIX 64-bit platforms

"Stinger" of the IBM® DB2® Universal Database Development Center supports the debugging of SQL stored procedures for UNIX® 64-bit platforms.

Development Center supports Java stored procedures on iSeries

In "Stinger", the IBM® DB2® Universal Database Development Center supports creating, running, and debugging Java™ stored procedures for iSeries™ V5R3 and later. You can perform these actions for two types of Java stored procedures:

- JDBC (dynamic SQL using JDBC)
- SQLJ (static SQL using SQLJ)

The following restrictions apply to this new feature:

- BINARY and VARBINARY data types are not supported
- PARAMETER STYLE DB2GENERAL is not supported

DB2 WebSphere MQ user-defined function 64-bit support

DB2® WebSphere® MQ user-defined functions (UDFs) are now supported on 64-bit AIX®, HP-UX, Solaris, and Windows®. To run the enable_MQFunctions and disable_MQFunctions commands, "-v" is no longer optional, and "-v 0pc" is required. On UNIX® 64-bit, the runtime library path must be modified to include \$HOME/sqllib/lib32 in order to execute the enable_MQFunction and disable_MQFunctions commands. The following settings include this modification:

AIX

```
LIBPATH=$HOME/sqllib/lib32 enable_MQFunctions -n dbname \  
-u userid -p passwd -v 0pc [-q qMgr -force -noValidate] \  
LIBPATH=$HOME/sqllib/lib32 disable_MQFunctions -n dbname \  
-u userid -p passwd -v 0pc
```

HP-UX

```
SHLIB_PATH=$HOME/sqllib/lib32 enable_MQFunctions -n dbname \  
-u userid -p passwd -v 0pc [-q qMgr -force -noValidate] \  
SHLIB_PATH=$HOME/sqllib/lib32 disable_MQFunctions -n dbname \  
-u userid -p passwd -v 0pc
```

Solaris

```
LD_LIBRARY_PATH=$HOME/sqllib/lib32 enable_MQFunctions -n dbname \  
-u userID -p passwd -v 0pc [-q qMgr -force -noValidate]  
LD_LIBRARY_PATH=$HOME/sqllib/lib32 disable_MQFunctions -n dbname \  
-u userID -p passwd -v 0pc
```

Related concepts:

- “MQSeries Enablement” in the *Application Development Guide: Programming Client Applications*
- “WebSphere MQ Functional Overview” in the *Application Development Guide: Programming Client Applications*

Related reference:

- “enable_MQFunctions” in the *Command Reference*
- “disable_MQFunctions” in the *Command Reference*

Advanced alter table functionality from the GUI

When altering a table from within the Alter Table notebook, you can now make more changes to table columns than you could previously. These changes are possible because the Alter Table notebook now drops and re-creates the table on an as-needed basis.

You can make the following additional changes when altering a table:

- Rename a column
- Drop a column
- Change the data type of a column
- Change the length, scope, or precision values for a column
- Change whether a column is nullable

If the table is dropped and re-created, attempts will be made to restore any dependent objects and to transform the existing data into the target data type of each remaining column.

Related tasks:

- “Alter Table -- Overview: Control Center help”

XA timeout sample code

See the CLI sample program `dbxamon.c` for an example of how to show and roll back indoubt transactions.

The sample program allows you to specify a timeout value for the transaction. If the transaction idle time becomes longer than the timeout value then the sample will rollback the transaction.

Related tasks:

- “Manually resolving indoubt transactions” in the *Administration Guide: Planning*

Related samples:

- “`dbxamon.c` -- Show and roll back indoubt transactions.”

Chapter 16. Business intelligence enhancements

Business intelligence enhancements summary

The following business intelligence enhancements have been added for this release:

Related concepts:

- “Online import” on page 65
- “Improved RUNSTATS performance through sampling” on page 25
- “Design advisor enhancements for recommending indexes, MQTs, MDC tables and partitions” on page 26
- “DB2 Data Warehouse Edition enhancements including integrated installation” on page 67

4 Online import

4 The import utility now supports two locking modes: offline mode (ALLOW NO
4 ACCESS) and online mode (ALLOW WRITE ACCESS). Offline mode is the default.
4 Previous to Version 8.1.4, offline mode was the only import behavior available.

4 The online mode provides better availability of the target table and potentially
4 increases the performance of the import utility.

4 Online import is especially valuable when range-clustered tables (RCT) are used.
4 Load into RCTs is not supported, and therefore performing multiple concurrent
4 imports into a single target table is the most efficient way to populate the database
4 using supported DB2 utilities.

Improved RUNSTATS performance through sampling

Table statistics are used by the query optimizer in selecting the best access plan for any given query, so it is important that statistics remain current to accurately reflect the state of a table at any given time. As the activity against a table increases, so should the frequency of statistics collection. With the increasing size of databases, it is becoming more important to find efficient ways to collect statistics. Random sampling of table data on which to collect statistics can reduce the amount of time that it takes to collect statistics. For I/O bound or CPU bound systems, the performance benefits can be enormous. The smaller the sample, the faster statistics collection completes.

Starting in "Stinger", the RUNSTATS command provides the option to collect statistics on a sample of the data in the table by using the TABLESAMPLE option. This feature can increase the efficiency of statistics collection since sampling uses only a subset of the data. At the same time, the sampling methods ensure a high level of accuracy.

Related concepts:

- “Data sampling in SQL queries” in the *Administration Guide: Performance*
- “Collecting statistics on a sample of the table data” in the *Administration Guide: Performance*

Related reference:

- “RUNSTATS Command” in the *Command Reference*

Design advisor enhancements for recommending indexes, MQTs, MDC tables and partitions

The DB2® Design advisor is a tool that can help you significantly improve your workload performance. The task of selecting which indexes, clustering dimensions, or partitions to create for a complex workload can be quite daunting. The Design advisor identifies all of the objects needed to improve the performance of your workload. Given a set of SQL statements in a workload, the Design advisor will generate recommendations for:

- new indexes
- new materialized query tables
- conversion to multidimensional clustering tables
- repartitioning of tables
- deletion of objects unused by the specified workload

You can choose to have the Design advisor GUI tool implement some or all of these recommendations immediately or schedule them for a later time.

Using either the Design advisor GUI or the command-line tool, the Design advisor can help simplify the following tasks:

Planning for or setting up a new database or partitioning structure

While designing your database or database partitions, use the Design advisor to:

- Generate design alternatives in a test environment for partitioning, indexes, MQTs, and MDC tables.
- Determine initial database partitioning before loading data into a database.
- Assist in migrating from a non-partitioned DB2 database to a partitioned DB2 database.
- Assist in migrating to DB2 in a partitioned environment from another database product.
- Evaluate indexes, MQTs or partitions that have been generated manually.

Workload performance tuning

After your database is set up, you can use the Design advisor to help you meet the following tuning goals:

- Improve performance of a particular statement or workload.
- Improve general database performance, using the performance of a sample workload as a gauge.
- Improve performance of the most frequently executed queries, for example, as identified by the Activity Monitor.
- Determine how to optimize the performance of a new key query.
- Respond to Health Center recommendations regarding shared memory utility or sort heap problems in a sort-intensive workload.
- Find objects that are not used in a workload.

Related concepts:

- “Automatic summary tables” in the *Administration Guide: Performance*

- “Multidimensional clustering tables” in the *Administration Guide: Planning*

DB2 Data Warehouse Edition enhancements including integrated installation

Users can install the various products that make up the DB2® Data Warehouse Edition through a single and convenient installation interface. The integrated installer allows you to pick and choose which components you want installed. The components include DB2 Cube Views, DB2 Intelligent Miner™, DB2 Warehouse Manager ETL (Extract, Transform, and Load), DB2 Query Patroller™, and DB2 Office Connect. This installation option is available on UNIX®, Linux, and Windows® operating systems.

In addition to the integrated installer, DB2 Data Warehouse Edition features a number of enhancements in its Business Intelligence components. These will be described in a separate announcement.

Related concepts:

- “Query Patroller installation environment overview” in the *DB2 Query Patroller Guide: Installation, Administration and Usage*
- “Warehouse Manager ETL (Extract, Transform, and Load) supports Unicode warehouse control databases” on page 51

Related tasks:

- “Query Patroller administration tasks overview” in the *DB2 Query Patroller Guide: Installation, Administration and Usage*

Chapter 17. DB2 family enhancements

DB2 family enhancements summary

The following DB2 family enhancements have been added for this release:

Related concepts:

- “DB2 Connect Linux zSeries 64-bit support” on page 74
- “DB2 Run-Time Client Lite” on page 59
- “DB2 Data Warehouse Edition enhancements including integrated installation” on page 67
- “DB2 Geodetic Extender” on page 74
- Chapter 7, “Find the DB2 documentation you need using the DB2 Information Center,” on page 13
- “Asynchronous I/O support on Linux (2.6 kernel exploitation)” on page 71
- “DB2 Query Patroller 64-bit support for UNIX, Linux, and Windows” on page 71
- “Licensing updates” on page 72
- “Function XMLNAMESPACES new for “Stinger”” on page 44
- “DB2 for Linux 64-bit support on the zSeries platform” on page 73
- “XML Extender platform support for HP-UX 32-bit and 64-bit PA-RISC” on page 35
- “Spatial Extender enhancements” on page 74
- “DB2 Express support on Linux PowerPC 32-bit” on page 73
- “DB2 for LinuxPPC 64-bit support” on page 73
- “DB2 for Linux 64-bit support on iSeries and pSeries platforms” on page 74

DB2 Run-Time Client Lite

The DB2[®] Run-Time Client Lite (DB2 RTCL) is a new installable component that makes it much easier to provide access to DB2 servers from Windows[®]-based applications. The DB2 RTCL is designed to be redistributable by independent software vendors (ISVs) and to be used for application distribution in mass deployment scenarios typical of large enterprises. Similar to the DB2 Run-Time Client component, the new DB2 RTCL provides application interfaces (CLI, ODBC, OLE DB, .NET Data provider, and JDBC) and the network libraries that are required for applications to run.

The main distinguishing features of the DB2 RTCL are:

- A significantly smaller disk footprint
- Shipped as a single executable making it easy to redistribute and deploy
- Windows Installer Merge Module (.msm file) is available which simplifies the integration of the DB2 RTCL code within a larger application
- Terms and conditions pertaining to the redistribution of the DB2 RTCL by ISVs and in-house programmers

Related concepts:

- “DB2 Run-Time Client Lite” in the *Quick Beginnings for DB2 Clients*

DB2 Data Warehouse Edition enhancements including integrated installation

Users can install the various products that make up the DB2® Data Warehouse Edition through a single and convenient installation interface. The integrated installer allows you to pick and choose which components you want installed. The components include DB2 Cube Views, DB2 Intelligent Miner™, DB2 Warehouse Manager ETL (Extract, Transform, and Load), DB2 Query Patroller™, and DB2 Office Connect. This installation option is available on UNIX®, Linux, and Windows® operating systems.

In addition to the integrated installer, DB2 Data Warehouse Edition features a number of enhancements in its Business Intelligence components. These will be described in a separate announcement.

Related concepts:

- “Query Patroller installation environment overview” in the *DB2 Query Patroller Guide: Installation, Administration and Usage*
- “Warehouse Manager ETL (Extract, Transform, and Load) supports Unicode warehouse control databases” on page 51

Related tasks:

- “Query Patroller administration tasks overview” in the *DB2 Query Patroller Guide: Installation, Administration and Usage*

Find the DB2 documentation you need using the DB2 Information Center

The DB2® Information Center for “Stinger” sets a new standard for delivering DB2 Universal Database™ and related DB2 product information by integrating it into a single information center. The DB2 Information Center gives you access to all of the DB2 information that you need, covering every aspect of DB2 from getting started, to using the DB2 Tools, database connectivity, database administration, query management, business intelligence, application development, and more. The DB2 Information Center also documents major DB2 features and components including replication, data warehousing, and the DB2 extenders. The navigation tree (table of contents) consists primarily of links to task and concept information that aligns with the high-level goals that you and other DB2 users want to achieve. In addition, the navigation tree includes entries for product overviews, reference information, a master index of all topics, and a glossary.

The DB2 Information Center is a fully searchable server that provides quick access to DB2 product information. To refine your search of all the topics in the DB2 Information Center, you can now use wildcards, quotation marks to search for phrases, and Boolean operators (AND, NOT, OR). If a translated version of a topic that you are searching for is not available in the preferred language set in your browser preferences, the DB2 Information Center displays the English version of that topic by default.

Using the DB2 Information Center Setup wizard, you can install the DB2 Information Center locally on your computer. However, if the computer on which

you install the DB2 Information Center is connected to a network, other computers can also access it through a Web browser. You can also decide not to install the DB2 Information Center locally, choosing instead to have your DB2 products access a version of the DB2 Information Center at an alternate location, such as the IBM® Web site where the most recently released version is available:
<http://publib.boulder.ibm.com/infocenter/db2help/>

Installing the DB2 Information Center locally lets you access DB2 product information without needing Internet or intranet access, which might be useful in a test or application development environment. If you install the DB2 Information Center on a networked computer, multiple users throughout your organization can access the same DB2 Information Center, giving you control over the version of the information that those users can see. This will minimize your update efforts and keep your network traffic within your intranet. If you want users to access the most up-to-date information, and you have access to the Internet, you can choose to configure DB2 products to access the DB2 Information Center on the IBM Web site. This version of the DB2 Information Center on the IBM Web site is updated as necessary to reflect the most recently released version of DB2 Universal Database, and also includes access to online service documents like TechNotes.

Related concepts:

- “DB2 Information Center” on page 84

Related tasks:

- “Displaying topics in your preferred language in the DB2 Information Center” on page 94
- “Invoking the DB2 Information Center” on page 92
- “Installing the DB2 Information Center using the DB2 Setup wizard (UNIX)” on page 88
- “Installing the DB2 Information Center using the DB2 Setup wizard (Windows)” on page 90

Asynchronous I/O support on Linux (2.6 kernel exploitation)

Asynchronous I/O (AIO) support is now available on Linux (2.6 kernel) for raw devices and O_DIRECT filesystems. Asynchronous I/O improves the performance of activities such as page cleaning and prefetching. You can enable and disable AIO on Linux at run time by issuing the **db2set** command.

To use AIO, users must install libaio-0.3.96 or later, have a kernel that supports AIO (e.g. version 2.6), run the **db2set DB2NOLIOAIO=false** command, and restart DB2®.

Related concepts:

- “I/O server configuration for prefetching and parallelism” in the *Administration Guide: Performance*

DB2 Query Patroller 64-bit support for UNIX, Linux, and Windows

DB2® Query Patroller 64-bit support is now available for UNIX® (AIX, Sun, HP-UX), Linux (IA64), and Windows® (Server 2003 and XP Professional) operating systems.

DB2 Query Patroller™ greatly improves data warehouse scalability, allowing hundreds of users to safely submit queries on databases that contain several terabytes of data. With this tool, administrators can automatically manage and control all aspects of query submission. DB2 Query Patroller prioritizes and schedules queries, and controls the number of queries running on a system at one time. As a result, query completion is more predictable, and computing resources are used more efficiently.

Related concepts:

- “Query Patroller” in the *DB2 Query Patroller Guide: Installation, Administration and Usage*

Related reference:

- “Installation requirements for Query Patroller server (Windows)” in the *DB2 Query Patroller Guide: Installation, Administration and Usage*
- “Installation requirements for Query Patroller server (UNIX)” in the *DB2 Query Patroller Guide: Installation, Administration and Usage*

Licensing updates

Starting with "Stinger", two additional licenses are available:

- Database partitioned feature (DPF) license
- DB2® Geodetic Extender license

Use the License Center to add these two licenses.

Database Partitioned Feature (DPF) license:

IBM® DB2 Universal Database™ Enterprise Server Edition (ESE) is distributed with a generic license key that allows you to start the DB2 ESE engine, create databases, connect to servers, and perform other tasks. To start partitioned instances, you will need to add the database partitioning feature (DPF) license key.

Note:

- Evaluation copies of DB2 Universal Database allow evaluation of all ESE features, including the DPF feature.
- Without the DPF license key, partitioned instances will be started, but violation messages will be logged in the administration tools log (also called the admin log).
- If you migrate pre-"Stinger" partitioned instances to "Stinger" and you do not have a DPF license key, violation messages will be logged in the admin log until you obtain and add a DPF license key.

DB2 Geodetic Extender license:

The geodetic license key allows you to use the DB2 Geodetic Extender functionality.

DB2 Geodetic Extender is a separately-priced feature, and requires that you purchase a separate license key. Refer to the *DB2 Spatial Extender and Geodetic Extender User's Guide and Reference* for details on how to enable DB2 Geodetic Extender.

Related concepts:

- “Data partitioning” in the *Administration Guide: Planning*
- “DB2 Geodetic Extender” in the *IBM DB2 Spatial Extender and Geodetic Extender User’s Guide and Reference*

Related tasks:

- “Enabling data partitioning in a database” in the *Administration Guide: Implementation*
- “Registering the DB2 product license key using the db2licm command” in the *Installation and Configuration Supplement*
- “Setting up and enabling DB2 Geodetic Extender” in the *IBM DB2 Spatial Extender and Geodetic Extender User’s Guide and Reference*
- “Adding a database partitioning feature (DPF) license: License Center help”
- “Adding a DB2 Geodetic Extender license: License Center help”

Related reference:

- “DB2 product license files” in the *Installation and Configuration Supplement*

Function XMLNAMESPACES new for "Stinger"

An XML namespace is a collection of names that is identified by a uniform reference identifier (URI). Namespaces are used in XML documents as element types and attribute names.

DB2® Universal Database includes several XML functions as part of the SQL/XML language specifications. Among these are XMLFOREST and XMLELEMENT. A new XML function, XMLNAMESPACES, provides XML namespace declarations within the SQL/XML publishing functions XMLELEMENT and XMLFOREST.

Related reference:

- “Expressions” in the *SQL Reference, Volume 1*

DB2 Express support on Linux PowerPC 32-bit

DB2® Express is now supported on Linux PowerPC® 32-bit.

DB2 for Linux 64-bit support on the zSeries platform

DB2® for Linux390 64-bit support is now available on the zSeries® platform. A 64-bit instance will be created when you install DB2 for Linux390 on zSeries by using the DB2 Setup wizard GUI installation, by using a response file installation, or by issuing the **db2icrt** command.

Related concepts:

- “DB2 Connect Linux zSeries 64-bit support” on page 74

DB2 for LinuxPPC 64-bit support

DB2® for LinuxPPC 64-bit support is now available, which includes the generation of refreshed hybrid DB2 images. You can create a 32-bit or a 64-bit instance for most DB2 products (excluding DB2 Universal Database Express Edition and DB2 Universal Database™ Workgroup Server Edition, which only supports the creation

of a 32-bit instance). Supported distributions include Red Hat Enterprise Linux 3 (RHEL3), SuSE SLES8 and SLES9 for 32-bit, and Red Hat Enterprise Linux 3 (RHEL3) and SuSE SLES9 for 64-bit.

DB2 for Linux 64-bit support on iSeries and pSeries platforms

DB2[®] for Linux 64-bit coverage is now available on iSeries[™] and pSeries[™] platforms. The benefits of 64-bit coverage include more address space for code, variables, shared memory, and for bufferpools. Supported distributions include Red Hat Enterprise Linux 3 (RHEL3) and SuSE SLES9.

DB2 Connect Linux zSeries 64-bit support

DB2[®] Connect on Linux zSeries[®] now offers 64-bit support.

DB2 extender enhancements

DB2 Geodetic Extender

DB2[®] Geodetic Extender uses the same spatial data types and functions as DB2 Spatial Extender to store and manipulate geographic data in a DB2 database. DB2 Geodetic Extender treats the Earth as a globe so that you can run seamless queries on geographic coordinates around the poles and coordinates that cross the date line. Distance and area calculations are precise, regardless of location on the surface of the Earth.

DB2 Geodetic Extender is a separately-priced feature, and requires that you purchase a separate license key.

Related concepts:

- “DB2 Geodetic Extender” in the *IBM DB2 Spatial Extender and Geodetic Extender User's Guide and Reference*

Related tasks:

- “Setting up and enabling DB2 Geodetic Extender” in the *IBM DB2 Spatial Extender and Geodetic Extender User's Guide and Reference*

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 within a database management system (DBMS). DB2 Spatial Extender conforms to the OpenGIS Consortium (OGC) and to ISO standards. DB2 Spatial Extender "Stinger" provides the following enhancements:

- With "Stinger", there is no longer a separate license fee for DB2 Spatial Extender, an add-on feature of DB2 Universal Database[™] Enterprise Server Edition.
- Spatial grid indexes help improve the performance of queries on spatial data. The spatial grid Index Advisor helps you to determine appropriate grid sizes, based on the minimum bounding rectangles (MBRs) of the geometries in the spatial column. MBR calculations and histogram outputs are improved. Multiple query window sizes enable the Index Advisor to recommend more accurate grid sizes.

Related concepts:

- “The purpose of DB2 Spatial Extender” in the *IBM DB2 Spatial Extender and Geodetic Extender User’s Guide and Reference*
- “System requirements for installing Spatial Extender” in the *IBM DB2 Spatial Extender and Geodetic Extender User’s Guide and Reference*
- “Tuning spatial grid indexes with the Index Advisor—Overview” in the *IBM DB2 Spatial Extender and Geodetic Extender User’s Guide and Reference*

XML Extender platform support for HP–UX 32–bit and 64–bit PA–RISC

DB2® XML Extender is now supported by HP-UX Version 11 in both 32-bit and 64-bit PA-RISC platforms. This is the first 64-bit platform supported by XML Extender. The XML Extender MQSeries® functions are not yet available for use on the HP-UX platform.

Related concepts:

- “Introduction to XML Extender” in the *DB2 XML Extender Administration and Programming*

Chapter 18. National language enhancements

National language enhancements summary

The following National language enhancements have been added for this release:

Related concepts:

- “Thai and Unicode collation algorithm differences” on page 77
- “New collators for Unicode databases” on page 77

Thai and Unicode collation algorithm differences

The collation algorithm used in the Thai Industrial Standard (TIS) TIS620-1 (code page 874) Thai database with the NLSCHAR collation option is similar to (but not identical to) the collation algorithm used in a Unicode database with the UCA400_LTH collation option. The differences identified in the standard include sorting character weight, as well as detailed differences in the weights of the space, hyphen, and full-stop characters. The standard also includes a discussion of specific Thai characters.

Users of Thai and Unicode databases should investigate and test the algorithm differences to ensure that collation is occurring according to their business practices.

Related concepts:

- “Unicode implementation in DB2 Universal Database” in the *Administration Guide: Planning*
- “New collators for Unicode databases” on page 77

New collators for Unicode databases

Two new collators for Unicode databases are now supported: UCA400_NO and UCA400_LTH.

Related concepts:

- “Unicode implementation in DB2 Universal Database” in the *Administration Guide: Planning*
- “Thai and Unicode collation algorithm differences” on page 77

Chapter 19. Security enhancements

Security enhancements summary

The following security enhancement has been added for this release:

Related concepts:

- “Common Criteria certification of DB2 Universal Database products” on page 79
- “System security enhancements for Windows improve usability” on page 16
- “Custom security plugins for authentication and group management” on page 81
- “Prevention of access to DB2 system files on Windows-based platforms” on page 80
- “Access token on Windows provides users' group information” on page 80
- “Encryption of user data using two new authentication types” on page 81

Common Criteria certification of DB2 Universal Database products

DB2 Universal Database is being evaluated for certification under the Common Criteria at evaluation assurance level 4 (EAL4). For more information about Common Criteria, see the Common Criteria web site at: <http://niap.nist.gov/cc-scheme/>.

System security enhancements for Windows improve usability

DB2[®] Universal Database, "Stinger", introduces the following enhancements to system security on the Windows[®] platform:

- Accepting additional special characters in user IDs and security mechanism group names (and consequently in authorization-names and authorization-ids) including &, - , and blank
- Accepting security mechanism group names that are longer than eight characters such as "Sales Representatives"
- Allowing two-part names on CONNECT and ATTACH that contain a Windows domain name and the userid to avoid the network traffic associated with looking up the user name in the trusted domain forest
- Support for Active Directory domains has been enhanced. For example, implicit trusts between domains, domain local groups and nested global groups are supported.

Related concepts:

- “DB2 UDB object naming rules” in the *Administration Guide: Implementation*
- “User, user ID and group naming rules” in the *Administration Guide: Implementation*

Support for Windows Local System account

Support for Windows Local System account (LSA) is provided for both the various DB2 services, as well as for applications that access them. The DB2 installation process allows DB2 services to run under SYSTEM (local system account) as an alternative to running under a dedicated user account. You can also perform the DB2 installation itself from a process that is running under the context of the Local System account.

This feature makes it easier to deploy DB2 by avoiding, where appropriate, the complexities associated with handling of userids and passwords. Another enhancement is the ability for DB2 applications that are running in the context of the LSA to access and use the local DB2 server.

Related concepts:

- “Windows local system account support” in the *Administration Guide: Implementation*

Related reference:

- “Required user accounts for installation of DB2 servers (Windows)” in the *Quick Beginnings for DB2 Servers*

Prevention of access to DB2 system files on Windows-based platforms

The new DB2_EXTSECURITY registry variable prevents unauthorized access to DB2 by locking DB2 system files. By default, this registry variable is set to ON.

Access token on Windows provides users' group information

An access token is an object that describes the security context of a process or thread. The information in the token includes the identity and privileges of the user account associated with the process or thread. When a user logs on, the system verifies the user's password by comparing it with information stored in a security database. If the password is authenticated, the system produces an access token. Every process executed on behalf of this user has a copy of this access token.

The access token contains information such as all the groups to which the user belongs including local groups and various domain groups (global groups, domain local groups, and universal groups).

The access token can also be acquired based on credentials cached by the operating system. In this case, the access token of the last logon can be referenced when the machine cannot contact the domain controller. For example, a laptop in a company's domain environment can still be logged on at home using the domain account, and all of the group's information of the domain account can still be referenced by applications as if the laptop was still connected to the company's domain.

Related concepts:

- “Security issues when installing DB2 Universal Database” in the *Administration Guide: Implementation*

Custom security plugins for authentication and group management

You can now create your own authentication and group management mechanisms in the form of loadable plugins that DB2[®] will load and access to perform user authentication. Now implemented within this plugin framework are the existing operating system-based and Kerberos authentication methods and the existing operating system-based group management mechanisms that DB2 provides. If you continue to use these existing methods, no changes are required. However, if you want to use the Kerberos authentication method for a non-Windows platform client, you will need to ensure that Kerberos is configured on that system and then update the required database manager configuration parameters.

The ability to create custom security plugins presents you with alternatives to the authentication methods currently provided by DB2. You can write your own plugins, and then compile and install them. Installation is simple: the plugins must be copied to the correct directories and then specific database manager configuration parameters must be updated.

To help you write your own security plugins, the `sqllib/samples/security/plugins` directory contains sample source files for various types of plugins.

Related concepts:

- “Kerberos authentication support for the UNIX platforms” on page 22

Kerberos authentication support for the UNIX platforms

Kerberos is a third party network authentication protocol that employs a system of shared secret keys to securely authenticate a user in an insecure network environment without flowing a text userid or password. Using Kerberos provides the advantages of single sign-on access to resources, such as a DB2[®] UDB server, and the centralization of user (principal) administration. Kerberos authentication support is extended to DB2 UDB for AIX[®], and for Solaris Operating Environment, in addition to existing support for Windows[®] operating systems that support Active Directory.

This support is provided as a security plug-in.

Related concepts:

- “Authentication methods for your server” in the *Administration Guide: Implementation*
- “Custom security plugins for authentication and group management” on page 81

Encryption of user data using two new authentication types

To improve the security of user data, there are two new authentication types that allow you to manage the encryption of user data. One new authentication type, `SQL_AUTHENTICATION_DATAENC`, requires connections to use the data encryption. The other new authentication type, `SQL_AUTHENTICATION_DATAENC_CMP`, allows for a compatibility mode with downlevel products that do not support the new authentication type. For such products, the use of this authentication type allows a connection using

SERVER_ENCRYPT where the encryption of user data is not required. Any product that supports SQL_AUTHENTICATION_DATAENC will be forced to use it.

Related concepts:

- “Authentication methods for your server” in the *Administration Guide: Implementation*

Appendix A. DB2 Universal Database technical information

DB2 documentation and help

DB2[®] technical information is available through the following tools and methods:

- DB2 Information Center
 - Topics
 - Help for DB2 tools
 - Sample programs
 - Tutorials
- Downloadable PDF files, PDF files on CD, and printed books
 - Guides
 - Reference manuals
- Command line help
 - Command help
 - Message help
 - SQL state help
- Installed source code
 - Sample programs

You can access additional DB2 Universal Database[™] technical information such as technotes, white papers, and Redbooks[™] online at ibm.com[®]. Access the DB2 Information Management software library site at www.ibm.com/software/data/pubs/.

DB2 documentation updates

IBM[®] may periodically make documentation FixPaks and other documentation updates to the DB2 Information Center available. If you access the DB2 Information Center at <http://publib.boulder.ibm.com/infocenter/db2help/>, you will always be viewing the most up-to-date information. If you have installed the DB2 Information Center locally, then you need to install any updates manually before you can view them. Documentation updates allow you to update the information that you installed from the *DB2 Information Center CD* when new information becomes available.

The Information Center is updated more frequently than either the PDF or the hardcopy books. To get the most current DB2 technical information, install the documentation updates as they become available or go to the DB2 Information Center at the www.ibm.com site.

Related concepts:

- “CLI sample programs” in the *CLI Guide and Reference, Volume 1*
- “Java sample programs” in the *Application Development Guide: Building and Running Applications*
- “DB2 Information Center” on page 84

Related tasks:

- “Invoking contextual help from a DB2 tool” on page 101

- “Updating the DB2 Information Center installed on your computer or intranet server” on page 93
- “Invoking message help from the command line processor” on page 102
- “Invoking command help from the command line processor” on page 103
- “Invoking SQL state help from the command line processor” on page 103

Related reference:

- “DB2 PDF and printed documentation” on page 95

DB2 Information Center

The DB2[®] Information Center gives you access to all of the information you need to take full advantage of DB2 family products, including DB2 Universal Database[™], DB2 Connect[™], DB2 Information Integrator and DB2 Query Patroller[™]. The DB2 Information Center also contains information for major DB2 features and components including replication, data warehousing, and the DB2 extenders.

The DB2 Information Center has the following features if you view it in Mozilla 1.0 or later or Microsoft[®] Internet Explorer 5.5 or later. Some features require you to enable support for JavaScript[™]:

Flexible installation options

You can choose to view the DB2 documentation using the option that best meets your needs:

- To effortlessly ensure that your documentation is always up to date, you can access all of your documentation directly from the DB2 Information Center hosted on the IBM[®] Web site at <http://publib.boulder.ibm.com/infocenter/db2help/>
- To minimize your update efforts and keep your network traffic within your intranet, you can install the DB2 documentation on a single server on your intranet
- To maximize your flexibility and reduce your dependence on network connections, you can install the DB2 documentation on your own computer

Search

You can search all of the topics in the DB2 Information Center by entering a search term in the **Search** text field. You can retrieve exact matches by enclosing terms in quotation marks, and you can refine your search with wildcard operators (*, ?) and Boolean operators (AND, NOT, OR).

Task-oriented table of contents

You can locate topics in the DB2 documentation from a single table of contents. The table of contents is organized primarily by the kind of tasks you may want to perform, but also includes entries for product overviews, goals, reference information, an index, and a glossary.

- Product overviews describe the relationship between the available products in the DB2 family, the features offered by each of those products, and up to date release information for each of these products.
- Goal categories such as installing, administering, and developing include topics that enable you to quickly complete tasks and develop a deeper understanding of the background information for completing those tasks.

- Reference topics provide detailed information about a subject, including statement and command syntax, message help, and configuration parameters.

Show current topic in table of contents

You can show where the current topic fits into the table of contents by clicking the **Refresh / Show Current Topic** button in the table of contents frame or by clicking the **Show in Table of Contents** button in the content frame. This feature is helpful if you have followed several links to related topics in several files or arrived at a topic from search results.

Index You can access all of the documentation from the index. The index is organized in alphabetical order by index term.

Glossary

You can use the glossary to look up definitions of terms used in the DB2 documentation. The glossary is organized in alphabetical order by glossary term.

Integrated localized information

The DB2 Information Center displays information in the preferred language set in your browser preferences. If a topic is not available in your preferred language, the DB2 Information Center displays the English version of that topic.

For iSeries™ technical information, refer to the IBM eServer™ iSeries information center at www.ibm.com/eserver/series/infocenter/.

Related concepts:

- “DB2 Information Center installation scenarios” on page 85

Related tasks:

- “Updating the DB2 Information Center installed on your computer or intranet server” on page 93
- “Displaying topics in your preferred language in the DB2 Information Center” on page 94
- “Invoking the DB2 Information Center” on page 92
- “Installing the DB2 Information Center using the DB2 Setup wizard (UNIX)” on page 88
- “Installing the DB2 Information Center using the DB2 Setup wizard (Windows)” on page 90

DB2 Information Center installation scenarios

Different working environments can pose different requirements for how to access DB2® information. The DB2 Information Center can be accessed on the IBM® Web site, on a server on your organization’s network, or on a version installed on your computer. In all three cases, the documentation is contained in the DB2 Information Center, which is an architected web of topic-based information that you view with a browser. By default, DB2 products access the DB2 Information Center on the IBM Web site. However, if you want to access the DB2 Information Center on an intranet server or on your own computer, you must install the DB2 Information Center using the DB2 Information Center CD found in your product Media Pack. Refer to the summary of options for accessing DB2 documentation which follows, along with the three installation scenarios, to help determine which

method of accessing the DB2 Information Center works best for you and your work environment, and what installation issues you might need to consider.

Summary of options for accessing DB2 documentation:

The following table provides recommendations on which options are possible in your work environment for accessing the DB2 product documentation in the DB2 Information Center.

Internet access	Intranet access	Recommendation
Yes	Yes	Access the DB2 Information Center on the IBM Web site, or access the DB2 Information Center installed on an intranet server.
Yes	No	Access the DB2 Information Center on the IBM Web site.
No	Yes	Access the DB2 Information Center installed on an intranet server.
No	No	Access the DB2 Information Center on a local computer.

Scenario: Accessing the DB2 Information Center on your computer:

Tsu-Chen owns a factory in a small town that does not have a local ISP to provide him with Internet access. He purchased DB2 Universal Database™ to manage his inventory, his product orders, his banking account information, and his business expenses. Never having used a DB2 product before, Tsu-Chen needs to learn how to do so from the DB2 product documentation.

After installing DB2 Universal Database on his computer using the typical installation option, Tsu-Chen tries to access the DB2 documentation. However, his browser gives him an error message that the page he tried to open cannot be found. Tsu-Chen checks the installation manual for his DB2 product and discovers that he has to install the DB2 Information Center if he wants to access DB2 documentation on his computer. He finds the *DB2 Information Center CD* in the media pack and installs it.

From the application launcher for his operating system, Tsu-Chen now has access to the DB2 Information Center and can learn how to use his DB2 product to increase the success of his business.

Scenario: Accessing the DB2 Information Center on the IBM Web site:

Colin is an information technology consultant with a training firm. He specializes in database technology and SQL and gives seminars on these subjects to businesses all over North America using DB2 Universal Database. Part of Colin’s seminars includes using DB2 documentation as a teaching tool. For example, while teaching courses on SQL, Colin uses the DB2 documentation on SQL as a way to teach basic and advanced syntax for database queries.

Most of the businesses at which Colin teaches have Internet access. This situation influenced Colin’s decision to configure his mobile computer to access the DB2 Information Center on the IBM Web site when he installed the latest version of DB2 Universal Database. This configuration allows Colin to have online access to the latest DB2 documentation during his seminars.

However, sometimes while travelling Colin does not have Internet access. This posed a problem for him, especially when he needed to access to DB2 documentation to prepare for seminars. To avoid situations like this, Colin installed a copy of the DB2 Information Center on his mobile computer.

Colin enjoys the flexibility of always having a copy of DB2 documentation at his disposal. Using the **db2set** command, he can easily configure the registry variables on his mobile computer to access the DB2 Information Center on either the IBM Web site, or his mobile computer, depending on his situation.

Scenario: Accessing the DB2 Information Center on an intranet server:

Eva works as a senior database administrator for a life insurance company. Her administration responsibilities include installing and configuring the latest version of DB2 Universal Database on the company's UNIX[®] database servers. Her company recently informed its employees that, for security reasons, it would not provide them with Internet access at work. Because her company has a networked environment, Eva decides to install a copy of the DB2 Information Center on an intranet server so that all employees in the company who use the company's data warehouse on a regular basis (sales representatives, sales managers, and business analysts) have access to DB2 documentation.

Eva instructs her database team to install the latest version of DB2 Universal Database on all of the employee's computers using a response file, to ensure that each computer is configured to access the DB2 Information Center using the host name and the port number of the intranet server.

However, through a misunderstanding Migual, a junior database administrator on Eva's team, installs a copy of the DB2 Information Center on several of the employee computers, rather than configuring DB2 Universal Database to access the DB2 Information Center on the intranet server. To correct this situation Eva tells Migual to use the **db2set** command to change the DB2 Information Center registry variables (DB2_DOCHOST for the host name, and DB2_DOCPORT for the port number) on each of these computers. Now all of the appropriate computers on the network have access to the DB2 Information Center, and employees can find answers to their DB2 questions in the DB2 documentation.

Related concepts:

- "DB2 Information Center" on page 84

Related tasks:

- "Updating the DB2 Information Center installed on your computer or intranet server" on page 93
- "Installing the DB2 Information Center using the DB2 Setup wizard (UNIX)" on page 88
- "Installing the DB2 Information Center using the DB2 Setup wizard (Windows)" on page 90
- "Setting the location for accessing the DB2 Information Center: Common GUI help"

Related reference:

- "db2set - DB2 Profile Registry Command" in the *Command Reference*

Installing the DB2 Information Center using the DB2 Setup wizard (UNIX)

DB2 product documentation can be accessed in three ways: on the IBM Web site, on an intranet server, or on a version installed on your computer. By default, DB2 products access DB2 documentation on the IBM Web site. If you want to access the DB2 documentation on an intranet server or on your own computer, you must install the documentation from the *DB2 Information Center CD*. Using the DB2 Setup wizard, you can define your installation preferences and install the DB2 Information Center on a computer that uses a UNIX operating system.

Prerequisites:

This section lists the hardware, operating system, software, and communication requirements for installing the DB2 Information Center on UNIX computers.

- **Hardware requirements**

You require one of the following processors:

- PowerPC (AIX)
- HP 9000 (HP-UX)
- Intel 32-bit (Linux)
- Solaris UltraSPARC computers (Solaris Operating Environment)

- **Operating system requirements**

You require one of the following operating systems:

- IBM AIX 5.1 (on PowerPC)
- HP-UX 11i (on HP 9000)
- Red Hat Linux 8.0 (on Intel 32-bit)
- SuSE Linux 8.1 (on Intel 32-bit)
- Sun Solaris Version 8 (on Solaris Operating Environment UltraSPARC computers)

Note: The DB2 Information Center runs on a subset of the UNIX operating systems on which DB2 clients are supported. It is therefore recommended that you either access the DB2 Information Center from the IBM Web site, or that you install and access the DB2 Information Center on an intranet server.

- **Software requirements**

- The following browser is supported:
 - Mozilla Version 1.0 or greater

- The DB2 Setup wizard is a graphical installer. You must have an implementation of the X Window System software capable of rendering a graphical user interface for the DB2 Setup wizard to run on your computer. Before you can run the DB2 Setup wizard you must ensure that you have properly exported your display. For example, enter the following command at the command prompt:
`export DISPLAY=9.26.163.144:0.`

- **Communication requirements**

- TCP/IP

Procedure:

To install the DB2 Information Center using the DB2 Setup wizard:

1. Log on to the system.
2. Insert and mount the DB2 Information Center product CD on your system.
3. Change to the directory where the CD is mounted by entering the following command:

```
cd /cd
```

where */cd* represents the mount point of the CD.

4. Enter the **./db2setup** command to start the DB2 Setup wizard.
5. The IBM DB2 Setup Launchpad opens. To proceed directly to the installation of the DB2 Information Center, click **Install Product**. Online help is available to guide you through the remaining steps. To invoke the online help, click **Help**. You can click **Cancel** at any time to end the installation.
6. On the **Select the product you would like to install** page, click **Next**.
7. Click **Next** on the **Welcome to the DB2 Setup wizard** page. The DB2 Setup wizard will guide you through the program setup process.
8. To proceed with the installation, you must accept the license agreement. On the **License Agreement** page, select **I accept the terms in the license agreement** and click **Next**.
9. Select **Install DB2 Information Center on this computer** on the **Select the installation action** page. If you want to use a response file to install the DB2 Information Center on this or other computers at a later time, select **Save your settings in a response file**. Click **Next**.
10. Select the languages in which the DB2 Information Center will be installed on **Select the languages to install** page. Click **Next**.
11. Configure the DB2 Information Center for incoming communication on the **Specify the DB2 Information Center port** page. Click **Next** to continue the installation.
12. Review the installation choices you have made in the **Start copying files** page. To change any settings, click **Back**. Click **Install** to copy the DB2 Information Center files onto your computer.

You can also install the DB2 Information Center using a response file.

The installation logs `db2setup.his`, `db2setup.log`, and `db2setup.err` are located, by default, in the `/tmp` directory.

The `db2setup.log` file captures all DB2 product installation information, including errors. The `db2setup.his` file records all DB2 product installations on your computer. DB2 appends the `db2setup.log` file to the `db2setup.his` file. The `db2setup.err` file captures any error output that is returned by Java, for example, exceptions and trap information.

When the installation is complete, the DB2 Information Center will be installed in one of the following directories, depending upon your UNIX operating system:

- AIX: `/usr/opt/db2_08_01`
- HP-UX: `/opt/IBM/db2/V8.1`
- Linux: `/opt/IBM/db2/V8.1`
- Solaris Operating Environment: `/opt/IBM/db2/V8.1`

Related concepts:

- “DB2 Information Center” on page 84
- “DB2 Information Center installation scenarios” on page 85

Related tasks:

- “Installing DB2 using a response file (UNIX)” in the *Installation and Configuration Supplement*
- “Updating the DB2 Information Center installed on your computer or intranet server” on page 93
- “Displaying topics in your preferred language in the DB2 Information Center” on page 94
- “Invoking the DB2 Information Center” on page 92
- “Installing the DB2 Information Center using the DB2 Setup wizard (Windows)” on page 90

Installing the DB2 Information Center using the DB2 Setup wizard (Windows)

DB2 product documentation can be accessed in three ways: on the IBM Web site, on an intranet server, or on a version installed on your computer. By default, DB2 products access DB2 documentation on the IBM Web site. If you want to access the DB2 documentation on an intranet server or on your own computer, you must install the DB2 documentation from the *DB2 Information Center CD*. Using the DB2 Setup wizard, you can define your installation preferences and install the DB2 Information Center on a computer that uses a Windows operating system.

Prerequisites:

This section lists the hardware, operating system, software, and communication requirements for installing the DB2 Information Center on Windows.

- **Hardware requirements**

You require one of the following processors:

- 32-bit computers: a Pentium or Pentium compatible CPU

- **Operating system requirements**

You require one of the following operating systems:

- Windows 2000
- Windows XP

Note: The DB2 Information Center runs on a subset of the Windows operating systems on which DB2 clients are supported. It is therefore recommended that you either access the DB2 Information Center on the IBM Web site, or that you install and access the DB2 Information Center on an intranet server.

- **Software requirements**

- The following browsers are supported:
 - Mozilla 1.0 or greater
 - Internet Explorer Version 5.5 or 6.0 (Version 6.0 for Windows XP)

- **Communication requirements**

- TCP/IP

Restrictions:

- You require an account with administrative privileges to install the DB2 Information Center.

Procedure:

To install the DB2 Information Center using the DB2 Setup wizard:

1. Log on to the system with the account that you have defined for the DB2 Information Center installation.
2. Insert the CD into the drive. If enabled, the auto-run feature starts the IBM DB2 Setup Launchpad.
3. The DB2 Setup wizard determines the system language and launches the setup program for that language. If you want to run the setup program in a language other than English, or the setup program fails to auto-start, you can start the DB2 Setup wizard manually.

To start the DB2 Setup wizard manually:

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

```
x:\setup.exe /i 2-letter language identifier
```

where *x*: represents your CD drive, and *2-letter language identifier* represents the language in which the setup program will be run.

- c. Click **OK**.
4. The IBM DB2 Setup Launchpad opens. To proceed directly to the installation of the DB2 Information Center, click **Install Product**. Online help is available to guide you through the remaining steps. To invoke the online help, click **Help**. You can click **Cancel** at any time to end the installation.
 5. On the **Select the product you would like to install** page, click **Next**.
 6. Click **Next** on the **Welcome to the DB2 Setup wizard** page. The DB2 Setup wizard will guide you through the program setup process.
 7. To proceed with the installation, you must accept the license agreement. On the **License Agreement** page, select **I accept the terms in the license agreement** and click **Next**.
 8. Select **Install DB2 Information Center on this computer** on the **Select the installation action** page. If you want to use a response file to install the DB2 Information Center on this or other computers at a later time, select **Save your settings in a response file**. Click **Next**.
 9. Select the languages in which the DB2 Information Center will be installed on **Select the languages to install** page. Click **Next**.
 10. Configure the DB2 Information Center for incoming communication on the **Specify the DB2 Information Center port** page. Click **Next** to continue the installation.
 11. Review the installation choices you have made in the **Start copying files** page. To change any settings, click **Back**. Click **Install** to copy the DB2 Information Center files onto your computer.

You can install the DB2 Information Center using a response file. You can also use the **db2rspgn** command to generate a response file based on an existing installation.

For information on errors encountered during installation, see the db2.log and db2wi.log files located in the 'My Documents'\DB2LOG\ directory. The location of the 'My Documents' directory will depend on the settings on your computer.

The db2wi.log file captures the most recent DB2 installation information. The db2.log captures the history of DB2 product installations.

Related concepts:

- “DB2 Information Center” on page 84
- “DB2 Information Center installation scenarios” on page 85

Related tasks:

- “Installing a DB2 product using a response file (Windows)” in the *Installation and Configuration Supplement*
- “Updating the DB2 Information Center installed on your computer or intranet server” on page 93
- “Displaying topics in your preferred language in the DB2 Information Center” on page 94
- “Invoking the DB2 Information Center” on page 92
- “Installing the DB2 Information Center using the DB2 Setup wizard (UNIX)” on page 88

Related reference:

- “db2rspgn - Response File Generator Command (Windows)” in the *Command Reference*

Invoking the DB2 Information Center

The DB2 Information Center gives you access to all of the information that you need to use DB2 products for Linux, UNIX, and Windows operating systems such as DB2 Universal Database, DB2 Connect, DB2 Information Integrator, and DB2 Query Patroller.

You can invoke the DB2 Information Center from one of the following places:

- Computers on which a DB2 UDB client or server is installed
- An intranet server or local computer on which the DB2 Information Center installed
- The IBM Web site

Prerequisites:

Before you invoke the DB2 Information Center:

- *Optional:* Configure your browser to display topics in your preferred language
- *Optional:* Configure your DB2 client to use the DB2 Information Center installed on your computer or intranet server

Procedure:

To invoke the DB2 Information Center on a computer on which a DB2 UDB client or server is installed:

- From the Start Menu (Windows operating system): Click **Start** → **Programs** → **IBM DB2** → **Information** → **Information Center**.
- From the command line prompt:
 - For Linux and UNIX operating systems, issue the **db2icdocs** command.
 - For the Windows operating system, issue the **db2icdocs.exe** command.

To open the DB2 Information Center installed on an intranet server or local computer in a Web browser:

- Open the Web page at <http://<host-name>:<port-number>/>, where <host-name> represents the host name and <port-number> represents the port number on which the DB2 Information Center is available.

To open the DB2 Information Center on the IBM Web site in a Web browser:

- Open the Web page at publib.boulder.ibm.com/infocenter/db2help/.

Related concepts:

- “DB2 Information Center” on page 84
- “DB2 Information Center installation scenarios” on page 85

Related tasks:

- “Displaying topics in your preferred language in the DB2 Information Center” on page 94
- “Invoking contextual help from a DB2 tool” on page 101
- “Updating the DB2 Information Center installed on your computer or intranet server” on page 93
- “Invoking command help from the command line processor” on page 103
- “Setting the location for accessing the DB2 Information Center: Common GUI help”

Related reference:

- “HELP Command” in the *Command Reference*

Updating the DB2 Information Center installed on your computer or intranet server

The DB2 Information Center available from <http://publib.boulder.ibm.com/infocenter/db2help/> will be periodically updated with new or changed documentation. IBM may also make DB2 Information Center updates available to download and install on your computer or intranet server. Updating the DB2 Information Center does not update DB2 client or server products.

Prerequisites:

You must have access to a computer that is connected to the Internet.

Procedure:

To update the DB2 Information Center installed on your computer or intranet server:

1. Open the DB2 Information Center hosted on the IBM Web site at: <http://publib.boulder.ibm.com/infocenter/db2help/>
2. In the Downloads section of the welcome page under the Service and Support heading, click the **DB2 Universal Database documentation** link.
3. Determine if the version of your DB2 Information Center is out of date by comparing the latest refreshed documentation image level to the documentation level you have installed. The documentation level you have installed is listed on the DB2 Information Center welcome page.

4. If a more recent version of the DB2 Information Center is available, download the latest refreshed *DB2 Information Center* image applicable to your operating system.
5. To install the refreshed *DB2 Information Center* image, follow the instructions provided on the Web page.

Related concepts:

- “DB2 Information Center installation scenarios” on page 85

Related tasks:

- “Invoking the DB2 Information Center” on page 92
- “Installing the DB2 Information Center using the DB2 Setup wizard (UNIX)” on page 88
- “Installing the DB2 Information Center using the DB2 Setup wizard (Windows)” on page 90

Displaying topics in your preferred language in the DB2 Information Center

The DB2 Information Center attempts to display topics in the language specified in your browser preferences. If a topic has not been translated into your preferred language, the DB2 Information Center displays the topic in English.

Procedure:

To display topics in your preferred language in the Internet Explorer browser:

1. In Internet Explorer, click the **Tools** —> **Internet Options** —> **Languages...** button. The Language Preferences window opens.
2. Ensure your preferred language is specified as the first entry in the list of languages.
 - To add a new language to the list, click the **Add...** button.

Note: Adding a language does not guarantee that the computer has the fonts required to display the topics in the preferred language.

- To move a language to the top of the list, select the language and click the **Move Up** button until the language is first in the list of languages.
3. Refresh the page to display the DB2 Information Center in your preferred language.

To display topics in your preferred language in the Mozilla browser:

1. In Mozilla, select the **Edit** —> **Preferences** —> **Languages** button. The Languages panel is displayed in the Preferences window.
2. Ensure your preferred language is specified as the first entry in the list of languages.
 - To add a new language to the list, click the **Add...** button to select a language from the Add Languages window.
 - To move a language to the top of the list, select the language and click the **Move Up** button until the language is first in the list of languages.
3. Refresh the page to display the DB2 Information Center in your preferred language.

Related concepts:

- “DB2 Information Center” on page 84

DB2 PDF and printed documentation

The following tables provide official book names, form numbers, and PDF file names. To order hardcopy books, you must know the official book name. To print a PDF file, you must know the PDF file name.

The DB2 documentation 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, or to print or view the PDF 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

Core DB2 information

The information in these books is fundamental to all DB2 users; you will find this information useful whether you are a programmer, a database administrator, or someone who works with DB2 Connect, DB2 Warehouse Manager, or other DB2 products.

Table 1. Core DB2 information

Name	Form Number	PDF File Name
<i>IBM DB2 Universal Database Command Reference</i>	SC09-4828	db2n0x81
<i>IBM DB2 Universal Database Glossary</i>	No form number	db2t0x81
<i>IBM DB2 Universal Database Message Reference, Volume 1</i>	GC09-4840, not available in hardcopy	db2m1x81
<i>IBM DB2 Universal Database Message Reference, Volume 2</i>	GC09-4841, not available in hardcopy	db2m2x81
<i>IBM DB2 Universal Database What's New</i>	SC09-4848	db2q0x81

Administration information

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

Table 2. Administration information

Name	Form number	PDF file name
<i>IBM DB2 Universal Database Administration Guide: Planning</i>	SC09-4822	db2d1x81
<i>IBM DB2 Universal Database Administration Guide: Implementation</i>	SC09-4820	db2d2x81
<i>IBM DB2 Universal Database Administration Guide: Performance</i>	SC09-4821	db2d3x81
<i>IBM DB2 Universal Database Administrative API Reference</i>	SC09-4824	db2b0x81
<i>IBM DB2 Universal Database Data Movement Utilities Guide and Reference</i>	SC09-4830	db2dmx81
<i>IBM DB2 Universal Database Data Recovery and High Availability Guide and Reference</i>	SC09-4831	db2hax81
<i>IBM DB2 Universal Database Data Warehouse Center Administration Guide</i>	SC27-1123	db2ddx81
<i>IBM DB2 Universal Database SQL Reference, Volume 1</i>	SC09-4844	db2s1x81
<i>IBM DB2 Universal Database SQL Reference, Volume 2</i>	SC09-4845	db2s2x81
<i>IBM DB2 Universal Database System Monitor Guide and Reference</i>	SC09-4847	db2f0x81

Application development information

The information in these books is of special interest to application developers or programmers working with DB2 Universal Database (DB2 UDB). You will find information about supported languages and compilers, as well as the documentation required to access DB2 UDB using the various supported programming interfaces, such as embedded SQL, ODBC, JDBC, SQLJ, and CLI. If you are using the DB2 Information Center, you can also access HTML versions of the source code for the sample programs.

Table 3. Application development information

Name	Form number	PDF file name
<i>IBM DB2 Universal Database Application Development Guide: Building and Running Applications</i>	SC09-4825	db2axx81
<i>IBM DB2 Universal Database Application Development Guide: Programming Client Applications</i>	SC09-4826	db2a1x81
<i>IBM DB2 Universal Database Application Development Guide: Programming Server Applications</i>	SC09-4827	db2a2x81

Table 3. Application development information (continued)

Name	Form number	PDF file name
<i>IBM DB2 Universal Database Call Level Interface Guide and Reference, Volume 1</i>	SC09-4849	db211x81
<i>IBM DB2 Universal Database Call Level Interface Guide and Reference, Volume 2</i>	SC09-4850	db212x81
<i>IBM DB2 Universal Database Data Warehouse Center Application Integration Guide</i>	SC27-1124	db2adx81
<i>IBM DB2 XML Extender Administration and Programming</i>	SC27-1234	db2sxx81

Business intelligence information

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

Table 4. Business intelligence information

Name	Form number	PDF file name
<i>IBM DB2 Warehouse Manager Standard Edition Information Catalog Center Administration Guide</i>	SC27-1125	db2dix81
<i>IBM DB2 Warehouse Manager Standard Edition Installation Guide</i>	GC27-1122	db2idx81
<i>IBM DB2 Warehouse Manager Standard Edition Managing ETI Solution Conversion Programs with DB2 Warehouse Manager</i>	SC18-7727	iwhe1mstx80

DB2 Connect information

The information in this category describes how to access data on mainframe and midrange servers using DB2 Connect Enterprise Edition or DB2 Connect Personal Edition.

Table 5. DB2 Connect information

Name	Form number	PDF file name
<i>IBM Connectivity Supplement</i>	No form number	db2h1x81
<i>IBM DB2 Connect Quick Beginnings for DB2 Connect Enterprise Edition</i>	GC09-4833	db2c6x81
<i>IBM DB2 Connect Quick Beginnings for DB2 Connect Personal Edition</i>	GC09-4834	db2c1x81
<i>IBM DB2 Connect User's Guide</i>	SC09-4835	db2c0x81

Getting started information

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

Table 6. Getting started information

Name	Form number	PDF file name
<i>IBM DB2 Universal Database Quick Beginnings for DB2 Clients</i>	GC09-4832, not available in hardcopy	db2itx81
<i>IBM DB2 Universal Database Quick Beginnings for DB2 Servers</i>	GC09-4836	db2isx81
<i>IBM DB2 Universal Database Quick Beginnings for DB2 Personal Edition</i>	GC09-4838	db2i1x81
<i>IBM DB2 Universal Database Installation and Configuration Supplement</i>	GC09-4837, not available in hardcopy	db2iyx81
<i>IBM DB2 Universal Database Quick Beginnings for DB2 Data Links Manager</i>	GC09-4829	db2z6x81

Tutorial information

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

Table 7. Tutorial information

Name	Form number	PDF file name
<i>Business Intelligence Tutorial: Introduction to the Data Warehouse</i>	No form number	db2tux81
<i>Business Intelligence Tutorial: Extended Lessons in Data Warehousing</i>	No form number	db2tax81
<i>Information Catalog Center Tutorial</i>	No form number	db2aix81
<i>Video Central for e-business Tutorial</i>	No form number	db2twx81
<i>Visual Explain Tutorial</i>	No form number	db2tvx81

Optional component information

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

Table 8. Optional component information

Name	Form number	PDF file name
<i>IBM DB2 Cube Views Guide and Reference</i>	SC18-7298	db2aax81

Table 8. Optional component information (continued)

Name	Form number	PDF file name
IBM DB2 Query Patroller Guide: Installation, Administration and Usage Guide	GC09-7658	db2dwx81
IBM DB2 Spatial Extender and Geodetic Extender User's Guide and Reference	SC27-1226	db2sbx81
IBM DB2 Universal Database Data Links Manager Administration Guide and Reference	SC27-1221	db2z0x82
DB2 Net Search Extender Administration and User's Guide	SH12-6740	N/A

Note: HTML for this document is *not* installed from the HTML documentation CD.

Release notes

The release notes provide additional information specific to your product's release and FixPak level. The release notes also provide summaries of the documentation updates incorporated in each release, update, 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.	Not available.

Note: The Release Notes are available in:

- XHTML and Text format, on the product CDs
- PDF format, on the PDF Documentation CD

In addition the portions of the Release Notes that discuss *Known Problems and Workarounds* and *Incompatibilities Between Releases* also appear in the DB2 Information Center.

To view the Release Notes in text format 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:

- For AIX operating systems: /usr/opt/db2_08_01
- For all other UNIX-based operating systems: /opt/IBM/db2/V8.1

Related concepts:

- "DB2 documentation and help" on page 83

Related tasks:

- "Printing DB2 books from PDF files" on page 100
- "Ordering printed DB2 books" on page 100
- "Invoking contextual help from a DB2 tool" on page 101

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 installed. If you need to install 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. Open `index.htm`. The file opens in a browser window.
3. Click on the title of the PDF you want to see. The PDF will open in Acrobat Reader.
4. Select **File** → **Print** to print any portions of the book that you want.

Related concepts:

- “DB2 Information Center” on page 84

Related tasks:

- “Mounting the CD-ROM (AIX)” in the *Quick Beginnings for DB2 Servers*
- “Mounting the CD-ROM (HP-UX)” in the *Quick Beginnings for DB2 Servers*
- “Mounting the CD-ROM (Linux)” in the *Quick Beginnings for DB2 Servers*
- “Ordering printed DB2 books” on page 100
- “Mounting the CD-ROM (Solaris Operating Environment)” in the *Quick Beginnings for DB2 Servers*

Related reference:

- “DB2 PDF and printed documentation” on page 95

Ordering printed DB2 books

If you prefer to use hardcopy books, you can order them in one of three ways.

Procedure:

Printed books can be ordered in some countries or regions. Check the IBM Publications website for your country or region to see if this service is available in your country or region. When the publications are available for ordering, you can:

- 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 <http://www.ibm.com/shop/publications/order>. The ability to order books from the IBM Publications Center may not be available in all countries.

At the time the DB2 product becomes available, the printed books are the same as those that are available in PDF format on the *DB2 PDF Documentation CD*. Content in the printed books that appears in the *DB2 Information Center CD* is also the same. However, there is some additional content available in DB2 Information Center CD that does not appear anywhere in the PDF books (for example, SQL Administration routines and HTML samples). Not all books available on the DB2 PDF Documentation CD are available for ordering in hardcopy.

Note: The DB2 Information Center is updated more frequently than either the PDF or the hardcopy books; install documentation updates as they become available or refer to the DB2 Information Center at <http://publib.boulder.ibm.com/infocenter/db2help/> to get the most current information.

Related tasks:

- “Printing DB2 books from PDF files” on page 100

Related reference:

- “DB2 PDF and printed documentation” on page 95

Invoking contextual help from a DB2 tool

Contextual help provides information about the tasks or controls that are associated with a particular window, notebook, wizard, or advisor. Contextual help is available from DB2 administration and development tools that have graphical user interfaces. There are two types of contextual help:

- Help accessed through the **Help** button that is located on each window or notebook
- Infopops, which are pop-up information windows displayed when the mouse cursor is placed over a field or control, or when a field or control is selected in a window, notebook, wizard, or advisor and F1 is pressed.

The **Help** button gives you access to overview, prerequisite, and task information. The infopops describe the individual fields and controls.

Procedure:

To invoke contextual help:

- For window and notebook help, start one of the DB2 tools, then open any window or notebook. Click the **Help** button at the bottom right corner of the window or notebook to invoke the contextual help.

You can also access the contextual help from the **Help** menu item at the top of each of the DB2 tools centers.

Within wizards and advisors, click on the Task Overview link on the first page to view contextual help.

- For infopop help about individual controls on a window or notebook, click the control, then click **F1**. Pop-up information containing details about the control is displayed in a yellow window.

Note: To display infopops simply by holding the mouse cursor over a field or control, select the **Automatically display infopops** check box on the **Documentation** page of the Tool Settings notebook.

Similar to infopops, diagnosis pop-up information is another form of context-sensitive help; they contain data entry rules. Diagnosis pop-up information is displayed in a purple window that appears when data that is not valid or that is insufficient is entered. Diagnosis pop-up information can appear for:

- Compulsory fields.
- Fields whose data follows a precise format, such as a date field.

Related tasks:

- “Invoking the DB2 Information Center” on page 92
- “Invoking message help from the command line processor” on page 102
- “Invoking command help from the command line processor” on page 103
- “Invoking SQL state help from the command line processor” on page 103
- “Access to the DB2 Information Center: Concepts help”
- “How to use the DB2 UDB help: Common GUI help”
- “Setting the location for accessing the DB2 Information Center: Common GUI help”
- “Setting up access to DB2 contextual help and documentation: Common GUI help”

Invoking message help from the command line processor

Message help describes the cause of a message and describes any action you should take in response to the error.

Procedure:

To invoke message help, open the command line processor and enter:

```
? XXXnnnnn
```

where *XXXnnnnn* represents a valid message identifier.

For example, ? SQL30081 displays help about the SQL30081 message.

Related concepts:

- “Introduction to messages” in the *Message Reference Volume 1*

Related reference:

- “db2 - Command Line Processor Invocation Command” in the *Command Reference*

Invoking command help from the command line processor

Command help explains the syntax of commands in the command line processor.

Procedure:

To invoke command help, open the command line processor and enter:

```
? command
```

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

For example, ? catalog displays help for all of the CATALOG commands, while ? catalog database displays help only for the CATALOG DATABASE command.

Related tasks:

- “Invoking contextual help from a DB2 tool” on page 101
- “Invoking the DB2 Information Center” on page 92
- “Invoking message help from the command line processor” on page 102
- “Invoking SQL state help from the command line processor” on page 103

Related reference:

- “db2 - Command Line Processor Invocation Command” in the *Command Reference*

Invoking SQL state help from the command line processor

DB2 Universal Database returns an SQLSTATE value for conditions that could be the result of an SQL statement. SQLSTATE help explains the meanings of SQL states and SQL state class codes.

Procedure:

To invoke SQL state 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, and ? 08 displays help for the 08 class code.

Related tasks:

- “Invoking the DB2 Information Center” on page 92
- “Invoking message help from the command line processor” on page 102
- “Invoking command help from the command line processor” on page 103

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:

You can view the XHTML versions of the tutorials from the Information Center at <http://publib.boulder.ibm.com/infocenter/db2help/>.

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

DB2 Universal Database tutorials:

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.

Information Catalog Center Tutorial

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

Visual Explain Tutorial

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

DB2 troubleshooting information

A wide variety of troubleshooting and problem determination information is available to assist you in using DB2® products.

DB2 documentation

Troubleshooting information can be found throughout the DB2 Information Center, as well as throughout the PDF books that make up the DB2 library. You can refer to the "Support and troubleshooting" branch of the DB2 Information Center navigation tree (in the left pane of your browser window) to see a complete listing of the DB2 troubleshooting documentation.

DB2 Technical Support Web site

Refer to the DB2 Technical Support Web site if you are experiencing problems and want help finding possible causes and solutions. The Technical Support site has links to the latest DB2 publications, TechNotes, Authorized Program Analysis Reports (APARs), FixPaks and the latest listing of internal DB2 error codes, and other resources. You can search through this knowledge base to find possible solutions to your problems.

Access the DB2 Technical Support Web site at <http://www.ibm.com/software/data/db2/udb/winos2unix/support>

DB2 Problem Determination Tutorial Series

Refer to the DB2 Problem Determination Tutorial Series Web site to find information on how to quickly identify and resolve problems you might encounter while working with DB2 products. One tutorial introduces you to the DB2 problem determination facilities and tools available, and helps you decide when to use them. Other tutorials deal with related topics, such

as "Database Engine Problem Determination", "Performance Problem Determination", and "Application Problem Determination".

See the full set of DB2 problem determination tutorials on the DB2 Technical Support site at <http://www.ibm.com/software/data/support/pdm/db2tutorials.html>

Related concepts:

- "DB2 Information Center" on page 84
- "Introduction to problem determination - DB2 Technical Support tutorial" in the *Troubleshooting Guide*

Accessibility

Accessibility features help users with physical disabilities, such as restricted mobility or limited vision, to use software products successfully. The following list specifies the major accessibility features in DB2[®] Version 8 products:

- All DB2 functionality is available using the keyboard for navigation instead of the mouse. For more information, see "Keyboard input and navigation."
- You can customize the size and color of the fonts on DB2 user interfaces. For more information, see "Accessible display."
- DB2 products support accessibility applications that use the Java[™] Accessibility API. For more information, see "Compatibility with assistive technologies" on page 106.
- DB2 documentation is provided in an accessible format. For more information, see "Accessible documentation" on page 106.

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 operations that can also be done using a mouse. Standard operating system keystrokes are used for standard operating system operations.

For more information about using keys or key combinations to perform operations, see Keyboard shortcuts and accelerators: Common GUI help.

Keyboard navigation

You can navigate the DB2 tools user interface using keys or key combinations.

For more information about using keys or key combinations to navigate the DB2 Tools, see Keyboard shortcuts and accelerators: Common GUI help.

Keyboard focus

In UNIX[®] operating systems, the area of the active window where your keystrokes will have an effect is highlighted.

Accessible display

The DB2 tools have features that improve accessibility for users with low vision or other visual impairments. These accessibility enhancements include support for customizable font properties.

Font settings

You can select the color, size, and font for the text in menus and dialog windows, using the Tools Settings notebook.

For more information about specifying font settings, see [Changing the fonts for menus and text: Common GUI help](#).

Non-dependence on color

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

Compatibility with assistive technologies

The DB2 tools interfaces support the Java Accessibility API, which enables you to use screen readers and other assistive technologies with DB2 products.

Accessible documentation

Documentation for DB2 is provided in XHTML 1.0 format, which is viewable in most Web browsers. XHTML 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.

Syntax diagrams are provided in dotted decimal format. This format is available only if you are accessing the online documentation using a screen-reader.

Related concepts:

- “Dotted decimal syntax diagrams” on page 106

Related tasks:

- “Keyboard shortcuts and accelerators: Common GUI help”
- “Changing the fonts for menus and text: Common GUI help”

Dotted decimal syntax diagrams

Syntax diagrams are provided in dotted decimal format for users accessing the Information Center using a screen reader.

In dotted decimal format, each syntax element is written on a separate line. If two or more syntax elements are always present together (or always absent together), they can appear on the same line, because they can be considered as a single compound syntax element.

Each line starts with a dotted decimal number; for example, 3 or 3.1 or 3.1.1. To hear these numbers correctly, make sure that your screen reader is set to read out punctuation. All the syntax elements that have the same dotted decimal number (for example, all the syntax elements that have the number 3.1) are mutually exclusive alternatives. If you hear the lines 3.1 USERID and 3.1 SYSTEMID, you know that your syntax can include either USERID or SYSTEMID, but not both.

The dotted decimal numbering level denotes the level of nesting. For example, if a syntax element with dotted decimal number 3 is followed by a series of syntax elements with dotted decimal number 3.1, all the syntax elements numbered 3.1 are subordinate to the syntax element numbered 3.

Certain words and symbols are used next to the dotted decimal numbers to add information about the syntax elements. Occasionally, these words and symbols might occur at the beginning of the element itself. For ease of identification, if the word or symbol is a part of the syntax element, it is preceded by the backslash (\) character. The * symbol can be used next to a dotted decimal number to indicate that the syntax element repeats. For example, syntax element *FILE with dotted decimal number 3 is given the format 3 * FILE. Format 3* FILE indicates that syntax element FILE repeats. Format 3* * FILE indicates that syntax element * FILE repeats.

Characters such as commas, which are used to separate a string of syntax elements, are shown in the syntax just before the items they separate. These characters can appear on the same line as each item, or on a separate line with the same dotted decimal number as the relevant items. The line can also show another symbol giving information about the syntax elements. For example, the lines 5.1*, 5.1 LASTRUN, and 5.1 DELETE mean that if you use more than one of the LASTRUN and DELETE syntax elements, the elements must be separated by a comma. If no separator is given, assume that you use a blank to separate each syntax element.

If a syntax element is preceded by the % symbol, this indicates a reference that is defined elsewhere. The string following the % symbol is the name of a syntax fragment rather than a literal. For example, the line 2.1 %OP1 means that you should refer to separate syntax fragment OP1.

The following words and symbols are used next to the dotted decimal numbers:

- ? means an optional syntax element. A dotted decimal number followed by the ? symbol indicates that all the syntax elements with a corresponding dotted decimal number, and any subordinate syntax elements, are optional. If there is only one syntax element with a dotted decimal number, the ? symbol is displayed on the same line as the syntax element, (for example 5? NOTIFY). If there is more than one syntax element with a dotted decimal number, the ? symbol is displayed on a line by itself, followed by the syntax elements that are optional. For example, if you hear the lines 5 ?, 5 NOTIFY, and 5 UPDATE, you know that syntax elements NOTIFY and UPDATE are optional; that is, you can choose one or none of them. The ? symbol is equivalent to a bypass line in a railroad diagram.
- ! means a default syntax element. A dotted decimal number followed by the ! symbol and a syntax element indicates that the syntax element is the default option for all syntax elements that share the same dotted decimal number. Only one of the syntax elements that share the same dotted decimal number can specify a ! symbol. For example, if you hear the lines 2? FILE, 2.1! (KEEP), and 2.1 (DELETE), you know that (KEEP) is the default option for the FILE keyword. In this example, if you include the FILE keyword but do not specify an option, default option KEEP will be applied. A default option also applies to the next higher dotted decimal number. In this example, if the FILE keyword is omitted, default FILE(KEEP) is used. However, if you hear the lines 2? FILE, 2.1, 2.1.1! (KEEP), and 2.1.1 (DELETE), the default option KEEP only applies to the next higher dotted decimal number, 2.1 (which does not have an associated keyword), and does not apply to 2? FILE. Nothing is used if the keyword FILE is omitted.
- * means a syntax element that can be repeated 0 or more times. A dotted decimal number followed by the * symbol indicates that this syntax element can be used zero or more times; that is, it is optional and can be repeated. For example, if you hear the line 5.1* data area, you know that you can include one

data area, more than one data area, or no data area. If you hear the lines 3*, 3 HOST, and 3 STATE, you know that you can include HOST, STATE, both together, or nothing.

Notes:

1. If a dotted decimal number has an asterisk (*) next to it and there is only one item with that dotted decimal number, you can repeat that same item more than once.
 2. If a dotted decimal number has an asterisk next to it and several items have that dotted decimal number, you can use more than one item from the list, but you cannot use the items more than once each. In the previous example, you could write HOST STATE, but you could not write HOST HOST.
 3. The * symbol is equivalent to a loop-back line in a railroad syntax diagram.
- + means a syntax element that must be included one or more times. A dotted decimal number followed by the + symbol indicates that this syntax element must be included one or more times; that is, it must be included at least once and can be repeated. For example, if you hear the line 6.1+ data area, you must include at least one data area. If you hear the lines 2+, 2 HOST, and 2 STATE, you know that you must include HOST, STATE, or both. Similar to the * symbol, the + symbol can only repeat a particular item if it is the only item with that dotted decimal number. The + symbol, like the * symbol, is equivalent to a loop-back line in a railroad syntax diagram.

Related concepts:

- “Accessibility” on page 105

Related tasks:

- “Keyboard shortcuts and accelerators: Common GUI help”

Related reference:

- “How to read the syntax diagrams” in the *SQL Reference, Volume 2*

Common Criteria certification of DB2 Universal Database products

DB2 Universal Database is being evaluated for certification under the Common Criteria at evaluation assurance level 4 (EAL4). For more information about Common Criteria, see the Common Criteria web site at: <http://niap.nist.gov/cc-scheme/>.

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

Numerics

- 64-bit support
 - DB2 Connect
 - Linux zSeries 74

A

- accessibility
 - dotted decimal syntax diagrams 106
 - features 105
- asynchronous I/O, page cleaner
- performance
 - Linux 71

B

- Business Intelligence Power Pack
 - installing 67, 70

C

- command help
 - invoking 103
- commands
 - db2look
 - enhancements 22

D

- DB2 books
 - printing PDF files 100
- DB2 clients
 - Run-Time Client Lite 59, 69
- DB2 Connect
 - Linux zSeries
 - 64-bit support 74
- DB2 Information Center 84
 - invoking 13, 70, 92
- DB2 tutorials 103
- DB2 Universal Database
 - prevention of access to DB2 system files on Windows platforms 80
- db2look command
 - enhancements 22
- disability 105
- documentation
 - displaying 92
- dotted decimal syntax diagrams 106

G

- Geodetic Extender
 - new extender 74

H

- help
 - displaying 92, 94

- help (*continued*)
 - for commands
 - invoking 103
 - for messages
 - invoking 102
 - for SQL statements
 - invoking 103
- HTML documentation
 - updating 93

I

- IMPORT utility
 - online 65
- Information Center
 - installing 85, 88, 90
- installing
 - Business Intelligence Power Pack 67, 70
 - Information Center 85, 88, 90
- invoking
 - command help 103
 - DB2 Information Center 13, 70
 - message help 102
 - SQL statement help 103
- iSeries Java stored procedures 63

J

- Java stored procedures for iSeries 63

K

- keyboard shortcuts
 - support for 105

L

- Linux
 - asynchronous I/O, page cleaner
 - performance 71
 - zSeries support 73
- Linux zSeries
 - DB2 Connect
 - 64-bit support 74

M

- message help
 - invoking 102

O

- online
 - help, accessing 101
 - import 65
- ordering DB2 books 100

P

- printed books, ordering 100
- printing
 - PDF files 100
- problem determination
 - online information 104
 - tutorials 104

Q

- Query Patroller
 - Support for 64-bit 71

R

- Run-Time Client Lite 59, 69
- RUNSTATS
 - with data sampling 25, 65

S

- sampling
 - with statistics collection 25, 65
- Spatial Extender
 - licensing changes 74
- Spatial Index Advisor
 - enhancements 74
- SQL statement help
 - invoking 103
- statistics collection
 - with data sampling 25, 65
- stored procedures, Java 63
- Support for 64-bit
 - Query Patroller 71

T

- troubleshooting
 - online information 104
 - tutorials 104
- tutorials 103
 - troubleshooting and problem determination 104

U

- Updating
 - HTML documentation 93

Z

- zSeries support
 - Linux 73

Contacting IBM

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

- 1-800-IBM-SERV (1-800-426-7378) 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 <http://www.ibm.com/planetwide>

Product information

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

This site contains the latest information on the technical library, ordering books, product 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



Printed in USA

SC09-4848-01



Spine information:



IBM® DB2 Universal Database™

What's New

"Singer"