

DB2

IBM

DB2 Version 9
for Linux, UNIX, and Windows




What's New

DB2®

IBM

DB2 Version 9
for Linux, UNIX, and Windows



What's New

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

Edition Notice

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 1993, 2006. All rights reserved.

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

DB2 9 BETA

Contents

| | |
|--------------------------------------|------------|
| About this book. | vii |
| Who should use this book | vii |
| How this book is structured. | vii |

Part 1. What's new for Version 9.1 . . . 1

Chapter 1. Highlights of Version 9.1. . . 3

| | |
|-----------------------------------------------------------------------|----|
| Highlights of Version 9.1 summary | 3 |
| Information as a service (native XML data store) | 6 |
| More agile development | 9 |
| Autonomic computing. | 10 |
| Improved large database management using table partitioning | 12 |
| Secure and resilient. | 15 |
| Manage your business, not your database | 15 |
| Hardware and operating system support | 17 |

Chapter 2. DB2 Connect enhancements 19

| | |
|-----------------------------------------------------------------------|----|
| DB2 Connect enhancements summary | 19 |
| DB2 Client easier to upgrade to DB2 Connect Personal Edition. | 21 |
| 32-bit DB2 Connect instance support changes | 22 |

Chapter 3. Product, packaging, and terminology changes 25

| | |
|----------------------------------------------------------------------------|----|
| Product, packaging, and terminology changes summary | 25 |
| New name for DB2 Universal Database for Linux, UNIX, and Windows | 25 |
| New names for DB2 Information Integrator products | 25 |
| DB2 Version 9.1 component replacements and name changes. | 26 |
| DB2 product-line changes | 27 |
| Distribution key terminology change | 27 |

Chapter 4. Native XML data store support 29

| | |
|------------------------------------------------------------------------------------------------|----|
| Native XML data store support summary | 29 |
| XML schema, DTD, and external entity management using the XML schema repository (XSR). | 30 |
| XQuery language support | 30 |
| XQuery builder | 31 |
| DB2 Net Search Extender support for the XML data type | 32 |
| Application development support - native XML data store. | 32 |
| Application development support summary - native XML data store. | 32 |
| Application programming language support for XML. | 32 |
| XML data type support in SQL and external procedures | 33 |
| XML support in Developer Workbench | 34 |

| | |
|-----------------------------------------------------------------------------------------------------|----|
| XML support in SQL statements and SQL/XML functions | 35 |
| XML support in SQL Assist | 36 |
| Annotated XML schema decomposition | 37 |
| Performance enhancements - native XML data store | 37 |
| Performance enhancements summary - native XML data store | 37 |
| Indexes over XML data | 37 |
| Performance guidelines for XML indexing | 38 |
| RUNSTATS command support for the native XML data store | 38 |
| Optimizer support for the native XML data store | 39 |
| Administration tools support - native XML data store. | 39 |
| Administration tools support summary - native XML data store | 39 |
| Control Center support for native XML data store | 39 |
| Explain and Visual Explain support for SQL/XML and XQuery statements | 40 |
| Command line processor (CLP) and command line tool support - native XML data store | 40 |
| Command line processor (CLP) and command line tool support summary - native XML data store. | 40 |
| DB2 command line processor support for the native XML data store. | 40 |
| db2batch command changes for native XML data store. | 41 |
| db2look command changes for native XML data store. | 41 |
| Import and export utility support for the native XML data store | 42 |

Chapter 5. Application development enhancements 43

| | |
|-------------------------------------------------------------------------------------------------|----|
| Application development enhancements summary | 43 |
| BINARY, VARBINARY, and DECFLOAT data type support in .NET and CLI client applications | 43 |
| Client support for trusted connections to DB2 for z/OS databases | 45 |
| Command line processor (CLP) 64 KB limit for SQL statements has been removed | 45 |
| DB2 .NET Data Provider enhancements and support for the .NET Framework 2.0 | 45 |
| Developer Workbench replaces the Development Center | 47 |
| External table function support across database partitions | 49 |
| IBM Database Add-Ins for Microsoft Visual Studio 2005 enhancements | 49 |
| Java routine class loader enhancements | 50 |
| JDBC and SQLJ enhancements | 52 |
| New and changed development software support | 53 |
| New samples. | 55 |
| SAMPLE database enhancements | 56 |

TRIM and STRIP scalar functions added. 57

Chapter 6. Backup, logging, and recovery enhancements 59

Backup, logging, and recovery enhancements summary 59
Continue a recover operation that ended during the rollforward phase 59
Redirected restore operation using an automatically generated script 59
Rebuild database function provides new restore options 60
Tivoli Storage Manager support added (Windows x64) 61

Chapter 7. Client and connectivity enhancements 63

Client and connectivity enhancements summary . . 63
Connection timeout support for database applications added 63
A single client (DB2 Client) is provided for application development and administration . . . 64
DB2 Runtime Client enhancements (Windows) . . 64
Internet Protocol Version 6 (IPv6) support added . 65
ODBC and CLI applications can now run without a DB2 client 65

Chapter 8. Federation enhancements 67

Federation enhancements summary 67
User mapping retrieval from an external repository is supported 67
Statement level isolation for nicknames 68
Two-phase commit for multivendor data sources . . 68

Chapter 9. Installation, migration, and fix pack enhancements. 69

Installation, migration, and fix pack enhancements summary 69
Coexistence of multiple DB2 versions and fix packs enhancements (Linux and UNIX) 69
Coexistence of multiple DB2 versions and fix packs now supported (Windows) 70
Installation of DB2 products without using an Administrator user ID now supported (Windows) . 72
Manual installation enhancements (Linux and UNIX) 73
Migration resources for DB2 database systems added 73
New DB2 product uninstall features using the db2unins command (Windows). 74
Response file enhancements 74
Update capabilities added to the DB2 Information Center 75

Chapter 10. Manageability enhancements 77

Manageability enhancements summary 77
Adaptive, self-tuning memory allocation 77
ALTER TABLE statement enhancements 78

Automated evaluation of object maintenance policies by the DB2 UDB for z/OS health monitor . 78
Automatic storage enhancements 79
Automatic table and index reorganization enhancements 80
Buffer pool identifiers displayed in database memory output 81
Copy schema support 81
Enhanced access to DB2 administration commands through SQL 81
EXEC SQL utility support for loading data into DB2 UDB for z/OS tables 82
Fast communications manager (FCM) enhancements 82
First Steps interface enhancements. 83
Identifier length limits increased to 128 bytes for additional objects and names 84
Load support for partitioned tables added 84
New SQL Monitor Area (SQLMA) for monitoring information flow in partitioned database environments. 85

Chapter 11. National language enhancements 87

National language enhancements summary. . . . 87
Unicode support added for character-based string functions 87

Chapter 12. Performance enhancements 89

Performance enhancements summary. 89
Enhanced query performance using statistical views 89
Faster data loading using SOURCEUSEREXIT customizable user exit. 90
Load from cursor with remote fetch 90
Materialized query table (MQT) enhancements . . 90
Online processing of the SET INTEGRITY statement 91
Optimizer registry variables and keywords added 91
Row compression support added 92
64 KB page size support added for POWER5+ processors (AIX). 93

Chapter 13. Scalability enhancements 95

Scalability enhancements summary 95
Increased maximum number of index columns and maximum size of index keys 95
Larger system and user temporary tables 95
Table size limits increased to 1.1 trillion rows and 16 terabytes 96

Chapter 14. Security enhancements . . 97

Security enhancements summary 97
Data access security improved through label-based access control (LBAC) 97
RESTRICT option added to CREATE DATABASE statement 98
Security administrator (SECADM) authority added to centralize security privileges. 99
SETSESSIONUSER privilege added 99
TRANSFER OWNERSHIP SQL statement added 100

Chapter 15. Troubleshooting and problem determination enhancements. 101

| | |
|--------------------------------------------------------------------------|-----|
| Troubleshooting and problem determination enhancements summary | 101 |
| Data collection improvements for database system failures | 101 |
| High water mark option added for db2pd -fcm command | 101 |
| Trace mask support added to the db2trc command | 102 |

Part 2. What's changed from Version 8.2. 103

Chapter 16. Changes in existing functionality 105

| | |
|-------------------------------------------------------------------------------------------------------|-----|
| Changes in existing functionality summary | 105 |
| Installation, migration, and fix pack changes | 105 |
| Installation, migration, and fix pack changes summary | 105 |
| Application and routine migration changes | 106 |
| db2_deinstall command parameter changes (Linux and UNIX) | 106 |
| DB2 install images package format changes (Linux and UNIX) | 106 |
| Documentation updates are available between product releases | 107 |
| Installation CD changes for national language versions of DB2 products (Linux and UNIX) | 108 |
| Licensing support changes | 108 |
| Locally installed Information Center installation restrictions | 108 |
| Migration support changes during installation (Windows) | 109 |
| Migration support changes for DB2 database systems | 110 |
| Multiple CD installation now required for some products (Linux and UNIX) | 111 |
| Database setup changes | 112 |
| Database setup changes summary | 112 |
| Authority required to start the database manager has changed (Windows) | 113 |
| Automatic configuration of prefetchers and page cleaners enabled by default | 113 |
| Automatic self-tuning memory enabled by default during database creation | 113 |
| Automatic statistics collection enabled by default during database creation | 114 |
| Automatic storage enabled by default during database creation | 115 |
| Configuration Advisor enabled by default during database creation | 116 |
| Configuration parameters default value changes | 117 |
| Databases require additional space | 119 |
| database_memory configuration parameter change | 119 |
| Increased log, table space, and memory requirements due to larger record identifiers (RIDs) | 120 |
| Lock lists require additional space | 120 |

| | |
|---------------------------------------------------------------------------------------------------|-----|
| Registry and environment variables default value changes | 121 |
| 32-bit DB2 database instance support changes | 121 |
| Administration changes | 123 |
| Administration changes summary | 123 |
| Application ID format changed | 123 |
| Backup image naming convention has changed (Windows) | 124 |
| Catalog table changes in support of table partitioning | 124 |
| Collating sequence IDENTITY replaces BINARY in command output | 125 |
| Data movement utilities changes | 125 |
| db2mtrk command changes (Windows) | 126 |
| DMS table space type default value changed to LARGE | 126 |
| Handling of new database connections has changed | 127 |
| Load from cursor incompatibilities with previous releases | 127 |
| Location of diagnostic messages for automatic maintenance | 127 |
| New DB2 agents may affect | |
| MAX_CONNECTIONS configuration and application monitor output | 128 |
| REORGCHK command output changes | 129 |
| ROLLFORWARD DATABASE output has changed | 129 |
| SET SESSION AUTHORIZATION requires SETSESSIONUSER privilege | 129 |
| sheapthres_shr parameter change | 130 |
| Table space-level point-in-time rollforward operations restrictions | 131 |
| TSM management class no longer used for filtering when restoring or retrieving logs | 131 |
| Write-to-table event monitor changes | 131 |
| Application development changes | 132 |
| Application development changes summary | 132 |
| Application and routine feature support changes | 132 |
| Changed parameters and output for the db2batch command | 133 |
| Cursor blocking is disabled in SQL procedures | 134 |
| DB2 Embedded Application Server updated | 134 |
| IBM Software Development Kit (SDK) for Java 5.x support added (AIX, Linux, and Windows) | 135 |
| Removal of mutating table conflicts when calling procedures from triggers | 135 |
| Reverse scans enabled by default for indexes, primary keys and unique keys | 136 |
| Sample build scripts updated | 136 |

Chapter 17. Deprecated functionality 139

| | |
|-------------------------------------------------------------------------------------------|-----|
| Deprecated functionality summary | 139 |
| Check pending table state is replaced and iCheckPending parameter is deprecated | 139 |
| CLISchema CLI keyword no longer supported | 140 |
| COLNAMES column in SYSCAT.INDEXES is deprecated | 140 |
| db2ilist command options are deprecated (Linux and UNIX) | 140 |
| DB2LINUXAIO registry variable is deprecated | 141 |

| | |
|------------------------------------------------------------------------------------|-----|
| db2secv82 command is deprecated | 141 |
| External routines now require an explicit entry point specification | 141 |
| Database logging using raw devices is deprecated | 142 |
| Some SQL administrative routines have been replaced | 143 |
| ADD PARTITIONING KEY clause of the ALTER TABLE statement is deprecated | 143 |
| DROP PARTITIONING KEY clause of the ALTER TABLE statement is deprecated | 144 |
| Type 1 indexes are deprecated | 144 |
| DB2 JDBC Type 2 driver is deprecated | 145 |
| The -w option of the db2icrt, db2ilist, and db2iupdt commands is ignored | 145 |

Chapter 18. Discontinued functionality 147

| | |
|------------------------------------------------------------------------------------------------|-----|
| Discontinued functionality summary | 147 |
| Alternate FixPak images are discontinued (UNIX) | 147 |
| Autoloader utility (db2atld) is no longer supported | 148 |
| Audio, Image, and Video (AIV) Extenders are no longer supported | 148 |
| Data Links Manager no longer supported | 148 |
| The Data Warehouse Center and the Information Catalog Center are no longer included | 150 |
| db2reg2large utility for converting DMS table space size is discontinued | 150 |
| db2profc and db2profp utilities are discontinued | 150 |
| DB2_SCATTERED_IO registry variable is discontinued (Linux) | 151 |
| Desktop icon and folder making utilities are no longer supported (Linux) | 151 |
| Extended Storage option for buffer pools is discontinued | 151 |
| Fast communications manager (FCM) no longer uses virtual interface (VI) architecture | 152 |

| | |
|--------------------------------------------------------------------------------|-----|
| NetBIOS and SNA communication protocols are no longer supported | 152 |
| DB2 Administration Tools are no longer supported on some platforms | 152 |
| Text Extender is no longer supported | 153 |
| Type 3 JDBC support is discontinued | 153 |
| Vendor load API (sqlvltld) is discontinued | 153 |
| VSE and VM objects are no longer supported in the DB2 Control Center | 154 |

Appendix A. DB2 Database technical information 155

| | |
|---------------------------------------------------------------------------------------------|-----|
| Overview of the DB2 technical information | 155 |
| Documentation feedback | 155 |
| DB2 technical library in PDF format | 156 |
| Ordering printed DB2 books | 158 |
| Displaying SQL state help from the command line processor | 159 |
| Accessing different versions of the DB2 Information Center | 159 |
| Displaying topics in your preferred language in the DB2 Information Center | 160 |
| Updating the DB2 Information Center installed on your computer or intranet server | 161 |
| DB2 Visual Explain tutorial | 162 |
| DB2 troubleshooting information | 163 |
| Terms and Conditions | 163 |

Appendix B. Notices 165

| | |
|----------------------|-----|
| Trademarks | 167 |
|----------------------|-----|

Index 169

Contacting IBM 181

About this book

This book provides information about the new and changed functionality in DB2[®] Database for Linux[®], UNIX[®], and Windows[®] Version 9.1 and DB2 Connect[™] Version 9.1.

Who should use this book

This book is for database administrators, application programmers, and other DB2 database users who want to quickly find out what new enhancements are available in DB2 Database for Linux, UNIX, and Windows Version 9.1 and DB2 Connect Version 9.1, and what differences exist between this version and Version 8.2 of those products. This book provides overview information and does not contain detailed instruction for using the features described. To get additional information, you should use the references that are provided.

If you are interested in features and enhancements introduced in Version 9.1, you should read Part 1, “What’s new for Version 9.1,” on page 1.

If you are interested in changed, deprecated, or discontinued functionality, you should read Part 2, “What’s changed from Version 8.2,” on page 103. This information will point out important changes that you need to know before using Version 9.1.

How this book is structured

The following topics are covered:

Part 1, What’s new for Version 9.1

Chapter 1, “Highlights of Version 9.1”

This chapter introduces important new features and enhancements in DB2 Version 9.1.

Chapter 2, “DB2 Connect enhancements”

This chapter describes the several enhancements and changes in DB2 Version 9.1 which affect the functionality and capabilities of DB2 Connect.

Chapter 3, “Product, packaging, and terminology changes”

This chapter describes the product, packaging and terminology changes introduced in Version 9.1, including component replacements, name changes, and product-line changes.

Chapter 4, “Native XML data store support”

The native XML data store enables well-formed XML documents to be stored in their hierarchical form within columns of a table. This chapter introduces the new XML data type and related native XML data store support available in Version 9.1.

Chapter 5, “Application development enhancements”

This chapter describes the enhancements introduced in Version 9.1 to simplify database application development.

Chapter 6, “Backup, logging, and recovery enhancements”

This chapter describes the backup, logging, and recovery enhancements, which will help you ensure the availability of your data.

Chapter 7, “Client and connectivity enhancements”

This chapter describes enhancements that ensure that you have flexible and effective methods of accessing data from client systems and applications, such as Internet Protocol Version 6 (IPv6) communication protocol support, support for running ODBC and CLI applications without a DB2 Client, and connection timeout support.

Chapter 8, “Federation enhancements”

This chapter describes the federation enhancements introduced in DB2 Version 9.1. With federation, you can send distributed requests to multiple data sources using a single SQL statement.

Chapter 9, “Installation, migration, and fix pack enhancements”

This chapter describes the new features and enhancements, such as coexistence of multiple DB2 database system versions and fix packs on Windows and support for concurrent copies of DB2 database systems on Linux and UNIX, that allow your IT staff to spend more time supporting business goals and less time installing and configuring database systems.

Chapter 10, “Manageability enhancements”

This chapter introduces several new features, including default enablement of autonomic computing features, automatic storage support for multi-partition databases, and support for copying schemas between databases. These features and others will allow you to spend less time managing your databases and more time managing your business.

Chapter 11, “National language enhancements”

This chapter describes the enhancements for national languages, including Unicode support for character-based string functions.

Chapter 12, “Performance enhancements”

This chapter describes the enhancements that ensure the highest performance when accessing and updating data, including data row compression, enhanced query performance using statistical views, and faster data loading capabilities.

Chapter 13, “Scalability enhancements”

This chapter introduces enhancements that provide the scalability you need for your growing business, such as larger system temporary tables, improved fast communications manager (FCM), and indexes that you can define on a greater number of columns.

Chapter 14, “Security enhancements”

This chapter describes security enhancements and other features that help you protect and manage your sensitive data.

Chapter 15, “Troubleshooting and problem determination enhancements”

This chapter describes enhancements that provide increased control over the set of diagnostic information produced when you encounter problems.

Part 2, What’s changed from Version 8.2**Chapter 16, “Changes in existing functionality”**

This chapter outlines the changes to existing DB2 database system functionality, including changes related to installation, migration, and fix packs; database setup; database administration; and application development.

Chapter 17, “Deprecated functionality”

This chapter lists the deprecated functionality, which refers to specific

functions or features that are supported but are no longer recommended and might be removed in a future release.

Chapter 18, “Discontinued functionality”

This chapter lists functions and features that are no longer supported in DB2 Version 9.1.

Additional information

Appendix A, “DB2 Database technical information”

This appendix contains information about accessing and using the latest documentation for your DB2 database system.

Appendix B, “Notices”

This appendix contains the legal requirements and limitations for using the DB2 database product and its documentation.

Part 1. What's new for Version 9.1

Read this part if you are interested in features and enhancements introduced in Version 9.1.

This part covers the following topics:

Chapter 1, “Highlights of Version 9.1”

This chapter introduces important new features and enhancements in DB2 Version 9.1.

Chapter 2, “DB2 Connect enhancements”

This chapter describes the several enhancements and changes in DB2 Version 9.1 which affect the functionality and capabilities of DB2 Connect.

Chapter 3, “Product, packaging, and terminology changes”

This chapter describes the product, packaging and terminology changes introduced in Version 9.1, including component replacements, name changes, and product-line changes.

Chapter 4, “Native XML data store support”

The native XML data store enables well-formed XML documents to be stored in their hierarchical form within columns of a table. This chapter introduces the new XML data type and related native XML data store support available in Version 9.1.

Chapter 5, “Application development enhancements”

This chapter describes the enhancements introduced in Version 9.1 to simplify database application development.

Chapter 6, “Backup, logging, and recovery enhancements”

This chapter describes the backup, logging, and recovery enhancements, which will help you ensure the availability of your data.

Chapter 7, “Client and connectivity enhancements”

This chapter describes enhancements that ensure that you have flexible and effective methods of accessing data from client systems and applications, such as Internet Protocol Version 6 (IPv6) communication protocol support, support for running ODBC and CLI applications without a DB2 Client, and connection timeout support.

Chapter 8, “Federation enhancements”

This chapter describes the federation enhancements introduced in DB2 Version 9.1. With federation, you can send distributed requests to multiple data sources using a single SQL statement.

Chapter 9, “Installation, migration, and fix pack enhancements”

This chapter describes the new features and enhancements, such as coexistence of multiple DB2 database system versions and fix packs on Windows and support for concurrent copies of DB2 database systems on Linux and UNIX, that allow your IT staff to spend more time supporting business goals and less time installing and configuring database systems.

Chapter 10, “Manageability enhancements”

This chapter introduces several new features, including default enablement of autonomic computing features, automatic storage support for multi-partition databases, and support for copying schemas between

databases. These features and others will allow you to spend less time managing your databases and more time managing your business.

Chapter 11, “National language enhancements”

This chapter describes the enhancements for national languages, including Unicode support for character-based string functions.

Chapter 12, “Performance enhancements”

This chapter describes the enhancements that ensure the highest performance when accessing and updating data, including data row compression, enhanced query performance using statistical views, and faster data loading capabilities.

Chapter 13, “Scalability enhancements”

This chapter introduces enhancements that provide the scalability you need for your growing business, such as larger system temporary tables, enhancements to fast communications manager (FCM), and indexes that you can define on a greater number of columns.

Chapter 14, “Security enhancements”

This chapter describes security enhancements and other features that help you protect and manage your sensitive data.

Chapter 15, “Troubleshooting and problem determination enhancements”

This chapter describes enhancements that provide increased control over the set of diagnostic information produced when you encounter problems.

Chapter 1. Highlights of Version 9.1

Highlights of Version 9.1 summary

DB2 Version 9.1 for Linux, UNIX, and Windows delivers new features that address the needs of your business today, whether those needs are integrating business data from across your organization, reducing IT costs, focusing limited IT resource on creating business value, or providing a secure and resilient information management system for your company's valuable information assets. To meet your business needs, DB2 Version 9.1 introduces important new features and enhancements.

Information as a service, using new hybrid relational and XML data server:

New features include:

- A new XML data type that allows you to store well-formed XML documents in their hierarchical form within columns of a table.
- Support for the XML data type in SQL statements and SQL/XML functions.
- Support for the new XQuery language developed by the World Wide Web Consortium (W3C). DB2 Version 9.1 allows you to invoke the XQuery language directly, calling functions that extract XML data from DB2 tables and views.
- New tools such as XQuery builder to create queries against XML data.
- Support for indexing over XML data, which improves the efficiency of queries that you issue against XML documents.
- Access and management of XML data by the DB2 data server. Existing DB2 tools such as the Control Center, command line processor (CLP), the **db2look** command, and Visual Explain are enhanced to support XML data.
- XML support in SQL and external procedures.
- Support for XML in many DB2-supported programming languages, which enables applications to combine XML and relational data access and storage.

For more information, see "Information as a service (native XML data store)" on page 6.

More agile application development, using an enhanced set of application development tools:

New features include:

- Application development tools and support for the new native XML data store feature, which enable applications to access and store both XML and relational data.
- An enhanced DB2 Driver for JDBC and SQLJ that complies with JDBC 3.0 specifications and includes support for SQLJ statements that perform functions equivalent to most JDBC methods. Other key features include support for many new data types, new DB2-only methods to support trusted connections to DB2 Version 9.1 for z/OS® (DB2 for z/OS) data servers, and heterogeneous pooling and connection reuse.
- The DB2 Developer Workbench, which is a comprehensive development environment for creating, editing, debugging, testing, and deploying DB2 stored

procedures and user-defined functions. You can also use Developer Workbench to develop SQLJ applications and to create, edit, and run SQL statements and XML queries.

- Additional support for Visual Studio 2005, which includes support for Web services, full support for the native XML data store, and the ability to build applications and Web sites without writing code.
- An enhanced DB2 Runtime Client which includes 64-bit support, coexistence with other DB2 products on the same computer, and free-distribution licensing with applications that you sell to other companies.

For more information, see “More agile development” on page 9.

New automated database administration features that improve productivity:

New features include:

- Adaptive, self-tuning memory allocation, which helps reduce or eliminate the task of configuring your DB2 server by continuously updating configuration parameters, resizing buffer pools and dynamically determining the total amount of memory to be used by the database.
- Automated health monitoring for DB2 UDB for z/OS objects, which allows you to automate object maintenance policy evaluations for DB2 UDB for z/OS objects through the Control Center’s Create and Change Object Policy wizard.
- Automatic storage support, which automatically grows the size of your database across disk and file systems, is now available for multi-partition databases.
- Automated statistics collection, which is enabled by default when you create a new database.
- Automatic configuration of prefetchers and page cleaners based on DB2 database system environment characteristics.
- New policy options for automated table and index reorganization, which provide your database administrator with more capabilities for managing table and index reorganizations.

For more information, see “Autonomic computing” on page 10.

Improved large database management, using table partitioning:

Table partitioning is a data organization scheme in which table data is divided across multiple storage objects called table partitions or ranges according to values in one or more table columns. These storage objects can be in different table spaces, in the same table space, or a combination of both.

Benefits of this new feature include:

- The ability to create very large tables. A partitioned table can contain vastly more data than an ordinary table. By dividing table data across multiple storage objects, you can significantly increase the size of a table.
- More flexible administration capabilities. You can now perform administrative tasks on individual data partitions, breaking down time-consuming maintenance operations into a series of smaller operations.
- More granular control of index placement. You can place indexes in different table spaces and manage them individually.
- Fast, easy roll in or roll out of data. This ability can be particularly useful in a data warehouse environment where you often move data in and out to run decision-support queries.

- Improved query performance. Separating data by using table partitioning allows you to improve query processing performance by avoiding scans of irrelevant data.

For more information, see “Improved large database management using table partitioning” on page 12.

New features that improve database security and resiliency:

Security features and enhancements include:

- Improved data access control at the row and column level, using label-based access control (LBAC).
- A new security administrator (SECADM) authority level that collects several security-related privileges under one privilege, providing greater control over access to information assets.
- A new RESTRICT option for the CREATE DATABASE statement, which provides greater control over granting database permissions.
- The SETSESSIONUSER privilege, which provides more control over who has authority to switch session user identities. It allows the holder to switch identities to any of the authorization IDs on which the privilege was granted.
- The TRANSFER OWNERSHIP SQL statement, which provides the ability to change the ownership of a database object.

Resiliency enhancements include:

- The ability to restart interrupted recovery operations, which can save precious time and effort in database recovery situations.
- Support for performing redirected restore operations with scripts automatically generated from existing backup images.
- The ability to rebuild databases from table space backup images. This functionality makes DB2 recovery more robust and versatile and provides you with a more complete recovery solution.

For more information, see “Secure and resilient” on page 15.

New performance, scalability, manageability, and installation enhancements help reduce the time you spend managing your database:

Performance and scalability enhancements include:

- The ability to compress table data objects using data row compression, which can help you achieve disk storage space savings, disk I/O savings, and quicker data access times.
- Statistical data for views, which can provide better access plans for improving query performance.
- Faster data-loading capabilities using customized scripts or programs.
- Materialized query table (MQT) enhancements that provide better support for designing MQTs, better query performance, and improved MQT maintenance.
- Larger record identifiers (RIDs), which allow more data pages per object and more records per page for system and user temporary tables used by the database manager while performing operations such as sorts and joins.
- Index keys that can include up to 64 columns and can be up to 8 KB in size.

Manageability enhancements include:

- Simpler memory management using adaptive, self-tuning memory allocation. Self-tuning memory provides a configuration that is dynamic and responsive to significant changes in workload characteristics.
- Automatic statistics collection enabled by default when you create databases. With automatic statistics collection enabled, DB2 Version 9.1 automatically runs the RUNSTATS utility in the background to ensure that the correct statistics are collected and maintained.
- Automatic storage support for multi-partition databases. This feature automatically grows the size of your database across disk and file systems, as required.
- ALTER TABLE statement enhancements that provide you to change some attributes of tables without having to drop and recreate the tables.
- New policy options that provide your database administrators with new automatic table and index reorganization capabilities.
- The ability to copy database schemas and create model schemas. Once your database administrators establish a model schema, they can use it as a template for creating new versions.
- New administrative SQL routines and views that provide a primary, easy-to-use programmatic interface to administer DB2 through SQL.
- Dynamic fast communication manager (FCM) buffers and new configuration parameters that can be tuned automatically by the DB2 database manager.

Installation enhancements include:

- A new CLI driver (the IBM® DB2 Driver for ODBC and CLI), which you can now install without installing a DB2 client.
- More flexible and efficient management of product licenses.
- The ability to install multiple DB2 versions and fix packs on the same computer.
- New response file keywords that allow your IT staff to set up DB2 products without end-user interaction.
- Non-administrator installation of DB2 products on the Windows operating system using the Windows elevated privileges feature.

For more information, see “Manage your business, not your database” on page 15.

Related concepts:

- “More agile development” on page 9
- “Autonomic computing” on page 10
- “Hardware and operating system support” on page 17
- “Improved large database management using table partitioning” on page 12
- “Information as a service (native XML data store)” on page 6
- “Manage your business, not your database” on page 15
- “Secure and resilient” on page 15

Information as a service (native XML data store)

DB2 Version 9.1 is a hybrid relational and XML data server that provides the ability to store both relational and XML data. DB2 Version 9.1 introduces a new native XML data store that is fully integrated into the DB2 database system, thus allowing you to access and manage XML data by leveraging DB2 functionality.

Native XML data store support in DB2 Version 9.1 includes:

- Integration with the DB2 database system, which includes:
 - Support for a new XML data type. The native XML data store enables you to store well-formed XML documents in their hierarchical forms within columns of a table. You define XML columns with the new XML data type. For more information, see “Native XML data store support summary” on page 29.
 - Support for the XQuery language. XQuery is a functional programming language that was designed by the World Wide Web Consortium (W3C) to meet specific requirements for querying XML data. DB2 Version 9.1 allows you to invoke XQuery directly, calling functions that extract XML data from DB2 tables and views. You can also invoke XQuery from an SQL query. For more information, see “XQuery language support” on page 30.
 - Support for the XML data type in SQL statements and SQL/XML functions. This support enables you to perform many common database operations. For more information, see “XML support in SQL statements and SQL/XML functions” on page 35.
 - Support for indexing XML data. The use of indexes over XML data improves the efficiency of queries that you issue against XML documents. For more information, see “Indexes over XML data” on page 37.

For more information about the DB2 native XML data store, see “Native XML data store support summary” on page 29.

- Enhanced and new tools to access and manage XML data, which include:
 - The new XQuery builder, which provides a graphical interface to help you to create and test XML queries without needing to understand XQuery syntax. For more information, see “XQuery builder” on page 31.
 - The Developer Workbench (which replaces the Version 8 Development Center), which contains support for XML functions, the XML data type, and XML schema registration. The XQuery builder is part of the Developer Workbench. For more information, see “XML support in Developer Workbench” on page 34.
 - The DB2 Command Line Processor (CLP), which you can use as an interface for interacting with DB2 instances and databases. The CLP provides support for the XML data type. For more information, see “DB2 command line processor support for the native XML data store” on page 40.
 - The Control Center, which helps you to administer DB2 databases and perform a variety of tasks, including creating objects and monitoring performance. The Control Center supports the XML data type in many of its administrative functions. This allows you to work with XML data and relational data using a single GUI tool. For more information, see “Control Center support for native XML data store” on page 39.
 - The Explain facility and Visual Explain GUI tool, which show you how DB2 evaluates query statements. Both support SQL/XML functions and XQuery statements. For more information, see “Explain and Visual Explain support for SQL/XML and XQuery statements” on page 40.

For information about other tools enhanced to support the native XML data store, see “Command line processor (CLP) and command line tool support summary - native XML data store” on page 40.

- Application development support, which includes:
 - XML support for programming languages, which enables applications to access and store both XML and relational data. For more information, see “Application programming language support for XML” on page 32.
 - XML support in SQL and external procedures, which enables XML data to be passed to SQL and external procedures by including parameters of data type

XML in CREATE PROCEDURE parameter signatures. As of DB2 UDB Version 8, procedures support SQL statements that produce or use XML values, as well as the temporary storage of XML data values in variables. For more information, see “XML data type support in SQL and external procedures” on page 33.

For more information about application development support for the native XML data store, see “Application development support summary - native XML data store” on page 32.

Benefits of the native XML data store include:

- A powerful mechanism for integrating and storing data from various data sources, such as eForms, documents, XML messages, or sources of other business-critical data - a feature not found in a traditional relational data server. The ability to integrate business data from multiple sources and services is key to making insightful decisions in today’s competitive marketplace.
- Storage of XML data using a hierarchical format representing the XML data model, instead of a relational model.
- High speed search retrieval with the introduction of new XML index types.
- Protection of the integrity of your XML data. Shredding XML data into relational tables compromises the digital signatures and other critical metadata that accompany your data. Because the native XML data store does not shred or decompose your XML data, your original XML document, including digital signatures, is protected. The native XML data store also allows you to avoid the resource and performance costs associated with rebuilding an XML document every time that it is retrieved.
- A robust and flexible foundation upon which you can build service-oriented applications. One of the key benefits of XML is its ability to standardize information, which allows for seamless communication with vendors, partners, and customers.
- Flexible schema capabilities that allow you to seamlessly and cost-effectively modify application structures without disrupting your data server.
- Flexible access to XML data using the new XQuery builder, XPath, SQL, and standard reporting tools.
- DB2 GUI tools to easily create and manage XML structures and build XQuery and SQL statements.
- The security and stability of DB2 Version 9.1.

You never have to compromise. The DB2 server incorporates the best XML and relational technologies into one server without forcing your XML developers to think like relational developers. For the licensing details of the native XML data store, see the DB2 Database for Linux, UNIX, and Windows home page at www.ibm.com/db2/udb.

Related concepts:

- “Administration tools support summary - native XML data store” on page 39
- “Application development support summary - native XML data store” on page 32
- “Command line processor (CLP) and command line tool support summary - native XML data store” on page 40
- “Highlights of Version 9.1 summary” on page 3
- “Native XML data store support summary” on page 29
- “Performance enhancements summary - native XML data store” on page 37

More agile development

DB2 Version 9.1 provides new features and enhancements that simplify database application development and ease application deployment.

These features and enhancements include:

- Support for a new native XML data store, which includes:
 - Application development support for the native XML data store, which enables applications to access and store both XML and relational data. For more information, see “Application programming language support for XML” on page 32.
 - Support for the new XQuery language developed by the World Wide Web Consortium (W3C). DB2 Version 9.1 allows you to invoke XQuery directly, calling functions that extract XML data from DB2 tables and views. For more information, see “XQuery language support” on page 30.
 - Support for XML in SQL statements and SQL/XML functions, which allows you to perform many common database operations using XML data. For more information, see “XML support in SQL statements and SQL/XML functions” on page 35.
 - The new XQuery builder, which helps you to build and test XML queries without needing to understand the semantics of the XQuery language. For more information, see “XQuery builder” on page 31.
 - Support for XML type in SQL and external procedures. For more information, see “XML data type support in SQL and external procedures” on page 33.
 - Annotated XML schema decomposition. For more information, see “Annotated XML schema decomposition” on page 37.
- An enhanced DB2 Driver for JDBC and SQLJ, which includes:
 - Support for SQLJ statements that perform functions equivalent to most JDBC methods
 - Support for many new data types
 - New DB2-only methods to support trusted connections to DB2 for z/OS database servers
 - Heterogeneous pooling and connection reuseFor more information about the DB2 Driver for JDBC and SQLJ, see “JDBC and SQLJ enhancements” on page 52.
- Developer Workbench, which includes:
 - The XQuery builder graphical tool, which helps you to build and test XML queries
 - Integrated stored procedure debugging capabilities
 - Support for developing SQLJ applications
 - Support for XML functions
 - Support for change management systems that enable you to share projectsFor more information about Developer Workbench, see “Developer Workbench replaces the Development Center” on page 47.
- Additional support for Microsoft® Visual Studio 2005, which includes:
 - Support for building applications and Web sites without writing code
 - Wide-ranging support for Web services
 - Ease-of-use improvements for working with database objects
 - Full support for the DB2 native XML data store

For more information about the additional support for Visual Studio, see “IBM Database Add-Ins for Microsoft Visual Studio 2005 enhancements” on page 49.

- Enhancements to the DB2 Runtime Client, which include:
 - A 64-bit version
 - The capability for the DB2 Runtime Client to coexist with other DB2 products on the same computer
 - Licensing changes, allowing you to freely distribute the DB2 Runtime Client
- For more information about enhancements to the DB2 Runtime Client, see “DB2 Runtime Client enhancements (Windows)” on page 64.

For information about other DB2 Version 9.1 application development enhancements, follow the related links below.

Related concepts:

- “Connection timeout support for database applications added” on page 63
- “A single client (DB2 Client) is provided for application development and administration” on page 64
- “External table function support across database partitions” on page 49
- “Highlights of Version 9.1 summary” on page 3
- “Java routine class loader enhancements” on page 50
- “New and changed development software support” on page 53
- “New samples” on page 55
- “SAMPLE database enhancements” on page 56
- “Application development enhancements summary” on page 43
- “Application development support summary - native XML data store” on page 32
- “Client support for trusted connections to DB2 for z/OS databases” on page 45

Autonomic computing

In DB2 Version 8.2, IBM introduced several autonomic computing features to make the database administrators' job easier. The features included were as follows:

- Design Advisor enhancements for recommending indexes, MQTs, MDC tables, and partitions.
- The Configure Automatic Maintenance wizard for automating database maintenance activities, such as backup, table defragmentation, and table statistics gathering.
- Self-healing features such as the Health Center Recommendation advisor and automated log file management.
- Self-tuning backup and restore operations.
- Ability to throttle backup operations and statistics collection.
- Automatic statistics profiling.
- New RECOVER DATABASE command for simplified database recovery.
- Automatic setting of prefetch size (DFT_PREFETCH_SZ configuration parameter).

Follow the related links for details about the Version 8.2 autonomic computing enhancements.

Building on the enhancements introduced in Version 8.2, new automated database administration features in DB2 Version 9.1 continue to help improve the productivity and effectiveness of your database administrators. Some of the key features include:

Adaptive, self-tuning memory allocation

This new feature helps reduce or eliminate the task of configuring your DB2 server by continuously updating configuration parameters and resizing buffer pools. When enabled, this feature dynamically distributes available memory resources between different memory consumers. On Windows and AIX® operating systems, the self-tuning memory feature can also determine overall database memory requirements and automatically tune the total database shared memory usage. For more information about this new feature, see “Adaptive, self-tuning memory allocation” on page 77.

Automated health monitoring for DB2 UDB for z/OS objects

Through the Control Center’s Create and Change Object Policy wizard, you can automate object maintenance policy evaluations for DB2 UDB for z/OS objects so that they are performed at scheduled times and intervals. For objects that create a policy alert, notifications are sent to health alert contacts specified in the object’s maintenance policy. For more information about health monitoring for DB2 UDB for z/OS objects, see “Automated evaluation of object maintenance policies by the DB2 UDB for z/OS health monitor” on page 78.

Automatic storage support

Automatic storage automatically grows the size of your database across disk and file systems. It eliminates the need to manage storage containers while taking advantage of the performance and flexibility of database managed storage. In DB2 Version 9.1, automatic storage support has been added for multi-partition databases. In addition, automatic storage is now enabled by default when you create new databases. For additional information about automatic storage enhancements, see the following topics:

- “Automatic storage enhancements” on page 79
- “Automatic storage enabled by default during database creation” on page 115

Automated statistics collection

In DB2 Version 9.1, automatic statistics collection using RUNSTATS is enabled by default when you create a new database. The DB2 server collects statistical information about your data in a background process when required. The DB2 optimizer uses this information to ensure the most efficient retrieval of information from the database. For more information, see “Automatic statistics collection enabled by default during database creation” on page 114.

Automatic configuration of prefetchers and page cleaners

In DB2 Version 9.1, the number of prefetchers and page cleaners can be automatically determined by the DB2 database system based on environment characteristics such as number of CPUs, number of database partitions, and parallelism settings of the table spaces in the database. For more information about enabling this new feature, see “Automatic configuration of prefetchers and page cleaners enabled by default” on page 113.

Automatic table and index reorganization enhancements

New policy options for automated table and index reorganization provide your database administrator with more capabilities for managing table and index reorganizations. For information about specific enhancements, see “Automatic table and index reorganization enhancements” on page 80.

For the complete list of DB2 Version 9.1 enhancements, follow the links under the ‘Related concepts’ section.

Related concepts:

- “Highlights of Version 9.1 summary” on page 3
- “Manageability enhancements summary” on page 77
- “Optimizing restore performance” in *Data Recovery and High Availability Guide and Reference*
- “Automatic statistics profiling” in *Performance Guide*
- “The Design Advisor” in *Performance Guide*

Related tasks:

- “Enabling automatic table and index reorganization” in *Performance Guide*
- “Resolving health monitor alerts using the Health Center” in *System Monitor Guide and Reference*

Related reference:

- “BACKUP DATABASE command” in *Command Reference*
- “RECOVER DATABASE command” in *Command Reference*
- “auto_maint - Automatic maintenance configuration parameter” in *Performance Guide*
- “dft_prefetch_sz - Default prefetch size configuration parameter” in *Performance Guide*

Improved large database management using table partitioning

Table partitioning is a data organization scheme in which table data is divided across multiple storage objects, called data partitions or ranges, according to values in one or more table columns. Each data partition is stored separately. These storage objects can be in different table spaces, in the same table space, or a combination of both.

The ability to partition table data across multiple storage objects provides your database administrators with greater scalability, flexibility, control, and performance. Specific benefits and features of table partitioning include:

- The ability to create very large tables. By dividing table data across multiple storage objects, you can significantly increase the size of a table.
- More flexible administration capabilities. Administration is more flexible because you can perform administrative tasks on individual data partitions, breaking down time-consuming maintenance operations into a series of smaller operations. For example, you can back up and restore individual data partitions instead of entire tables. For more information, see “Table partitioning” in *Administration Guide: Planning*.
- More granular control of index placement. You can place indexes in different table spaces and manage them individually. For more information, see “Understanding index behavior on partitioned tables” in *Performance Guide*.

- Fast, easy roll in or roll out of data using the ATTACH PARTITION and DETACH PARTITION clauses of the ALTER TABLE statement. This ability can be particularly useful in a data warehouse environment where you often need to load or delete data to run decision-support queries. For more information see "Attaching a data partition" in *Administration Guide: Implementation* and "Detaching a data partition" in *Administration Guide: Implementation*.
- Improved query performance. Separating data using table partitioning allows you to improve query processing performance by avoiding scans of irrelevant data. The DB2 optimizer eliminates irrelevant partitions from a query operation. For more information, see "Optimization strategies for partitioned tables" in *Performance Guide*.
- The ability to combine table partitioning with other data organization schemes. Using table partitioning with the Data Partitioning Feature (DPF), you can spread ranges of data evenly across database partitions to take advantage of the intra-query parallelism and database partition load balancing features of DPF. When used with multi-dimensional clustering (MDC), table partitioning allows you to group rows with similar values on multiple dimensions in the same table extent.
- DB2 Control Center support. You can use the Create Table wizard in the DB2 Control Center as an alternative to the CREATE TABLE statement for creating partitioned tables. You can also migrate an existing table or view to a partitioned table. For more information, see "Creating a table using the Create Table wizard" in *Administration Guide: Implementation*.
- Load support for partitioned tables. The load utility inserts data records into the correct data partition without the need to use an external utility to partition the input data before loading. For more information, see "Load support for partitioned tables added" on page 84.
- Performance and cost optimization of storage. By only using your fastest and most expensive storage hardware for only the most active table partitions, you can optimize your overall storage costs. If most of your queries only run against the last three months of data you have to option to assign slower and cheaper storage hardware to older data.

Frequently asked questions about table partitioning:

1. Is table partitioning similar to equivalent functionality in DB2 for z/OS or Informix[®] Dynamic Server and Informix Extended Parallel Server?

The functionality for attach, detach, and combined roll in and roll out of table data is extremely similar. The DB2 for z/OS, DB2 UDB for iSeries and DB2 Database for Linux, UNIX, and Windows use a common base syntax, but implement a different subset. DB2 Version 9.1 uses an abbreviated syntax that is actually much more concise than the syntax by either of the two products.

2. Is using table partitioning the same as using DPF?

No. Table partitioning divides table data by allowing you to define ranges of data for a table so that each range is stored separately. Each range, known as a data partition, corresponds to a single storage object. These storage objects can be in different table spaces, in the same table space, or a combination of both. By contrast, DPF distributes table data evenly across database partitions.

The DB2 Version 9.1 organization schemes include:

- DISTRIBUTE BY HASH
 - PARTITION BY RANGE
 - ORGANIZE BY DIMENSIONS
3. Does table partitioning work with Database Partitioning feature (DPF)?

Absolutely. You still get the parallelism of DPF and you also get the data partition elimination of table partitioning to boost query performance. By specifying the DISTRIBUTE BY and PARTITION BY clauses of the CREATE TABLE statement, you can spread data across database partitions spanning multiple table spaces.

4. How does table partitioning improve query performance?

Table partitioning improves query performance through data partition elimination. Data partition elimination refers to the database servers ability to determine, based on the query predicates, that only a subset of the data partitions of a table need to be accessed to answer a query. Data partition elimination offers particular benefits when you are running decision-support queries against a partitioned table.

5. How long does a reorganization take after an attach or a detach of a data partition?

There is no reorganization needed during an attach or a detach operation because there is no data movement. The pre-existing table is logically linked up to the partitioned table. The downtime is extremely small (less than 10 seconds).

6. What should I do with my existing UNION ALL views?

Get rid of them! You don't need UNION ALL views anymore. Table partitioning does everything that a UNION ALL view can, including rolling in and rolling out table data.

7. How many data partitions can I have?

You can create a partitioned table with thousands of data partitions. A partitioned table can contain vastly more data than an ordinary table.

Related concepts:

- "Data organization schemes" in *Administration Guide: Planning*
- "Partitioned database environments" in *Administration Guide: Planning*
- "Table partitioning" in *Administration Guide: Planning*
- "Table partitioning keys" in *Administration Guide: Planning*
- "Data organization schemes in DB2 and Informix databases" in *Administration Guide: Planning*
- "Highlights of Version 9.1 summary" on page 3
- "Load support for partitioned tables added" on page 84

Related tasks:

- "Creating partitioned tables" in *Administration Guide: Implementation*
- "Adding data partitions to partitioned tables" in *Administration Guide: Implementation*
- "Approaches to migrating existing tables and views to partitioned tables" in *Administration Guide: Implementation*
- "Creating a table using the Create Table wizard" in *Administration Guide: Implementation*

Related reference:

- "Examples of rolling in and rolling out partitioned table data" in *Administration Guide: Implementation*
- "CREATE TABLE statement" in *SQL Reference, Volume 2*

Secure and resilient

IT security is a major concern for organizations today. It has become crucial to ensure that sensitive data is protected. IT Administrators need the ability to efficiently manage system security, quickly analyze their security environment, and monitor access to data.

DB2 Version 9.1 provides several new features and enhancements that contribute to providing a secure and resilient environment for your data. Version 9.1 introduces security enhancements to ensure that your sensitive data continues to be even better protected; and backup, logging, and recovery enhancements to help ensure that your data is available during all hours of the day.

- Security enhancements include:
 - Support for data access control using label-based access control (LBAC). For more information, see “Data access security improved through label-based access control (LBAC)” on page 97.
 - A new security administrator authority level (SECADM), which provides greater control over access to information assets and improved reporting capabilities for monitoring access to sensitive data. For more information, see “Security administrator (SECADM) authority added to centralize security privileges” on page 99.
 - A new RESTRICT option for the CREATE DATABASE statement, which provides greater control over database privileges. For more information, see “RESTRICT option added to CREATE DATABASE statement” on page 98.

For the complete list of Version 9.1 security enhancements, see “Security enhancements summary” on page 97.

- Backup and recovery enhancements include:
 - The ability to restart interrupted recovery operations. For more information, see “Continue a recover operation that ended during the rollforward phase” on page 59.
 - Support for performing redirected restore operations by generating scripts from existing backup images. Redirected restore operations let you redefine table space containers. For more information, see “Redirected restore operation using an automatically generated script” on page 59.
 - The ability to rebuild databases from table space backup images. For more information, see “Rebuild database function provides new restore options” on page 60.

Related concepts:

- “Backup, logging, and recovery enhancements summary” on page 59
- “Highlights of Version 9.1 summary” on page 3
- “Security enhancements summary” on page 97

Manage your business, not your database

Working with companies such as SAP, IBM has made many improvements in DB2 Version 9.1 that will allow your IT staff to spend more time supporting your business needs and less time installing and managing database systems. Version 9.1 introduces performance and scalability enhancements to help them achieve the highest performance when accessing and updating your data; manageability

enhancements to enable them to reduce the time required to administer and tune your database systems; and installation enhancements to allow them to set up and deploy your applications more quickly.

- Performance and scalability enhancements include:
 - The ability to compress table data objects using data row compression. For more information, see “Row compression support added” on page 92.
 - Improved access plans for queries using statistical views. For more information, see “Enhanced query performance using statistical views” on page 89.
 - Faster data-loading capabilities using customized scripts or programs. For more information, see “Faster data loading using SOURCEUSEREXIT customizable user exit” on page 90.
 - Enhanced query execution for materialized query tables. For more information, see “Materialized query table (MQT) enhancements” on page 90.
 - Larger record identifiers (RIDs), which allow more data pages per object and more records per page. For more information, see “Table size limits increased to 1.1 trillion rows and 16 terabytes” on page 96.
 - Index keys that can include up to 64 columns and that are up to 8 KB in size. For more information, see “Increased maximum number of index columns and maximum size of index keys” on page 95.

For more information about performance enhancements, see “Performance enhancements summary” on page 89.

- Manageability enhancements include:
 - Simpler memory management using adaptive, self-tuning memory allocation. Self-tuning memory provides a configuration that is dynamic and responsive to significant changes in workload characteristics. For more information, see “Adaptive, self-tuning memory allocation” on page 77.
 - Automatic statistic collection enabled by default when creating databases. With automatic statistics collection enabled, DB2 automatically runs the RUNSTATS utility in the background to ensure the correct statistics are collected and maintained. For more information, see “Automatic statistics collection enabled by default during database creation” on page 114.
 - Automatic storage support for partitioned databases. For more information, see “Automatic storage enhancements” on page 79.
 - The ability to change some attributes of tables without having to drop and recreate the tables. For more information, see “ALTER TABLE statement enhancements” on page 78.
 - New policy options that provide database administrators with more table and index reorganization capabilities. For more information, see “Automatic table and index reorganization enhancements” on page 80.
 - The ability to copy database schemas and create model schemas. Once your database administrators establish a model schema, they can use it as a template for creating new versions. For more information, see “Copy schema support” on page 81.
 - New administrative SQL routines and views. The administrative routines and views provide a primary, easy-to-use programmatic interface to administer DB2 through SQL. For more information, see “Enhanced access to DB2 administration commands through SQL” on page 81.

- Dynamic fast communication manager (FCM) buffers and new configuration parameters that can be tuned automatically by the DB2 database manager. For more information, see “Fast communications manager (FCM) enhancements” on page 82.

For the complete list of manageability enhancements, see “Manageability enhancements summary” on page 77.

- Installation enhancements include:
 - Easier management of product licenses using the License Center and the db2licm command. For more information, see “Licensing support changes” on page 108.
 - The ability to install multiple DB2 versions and fix packs on the same computer. For more information, see “Coexistence of multiple DB2 versions and fix packs enhancements (Linux and UNIX)” on page 69 and “Coexistence of multiple DB2 versions and fix packs now supported (Windows)” on page 70.
 - New response file keywords that allow your IT staff to setup DB2 products without end-user interaction. For more information, see “Response file enhancements” on page 74.
 - Licensing changes for the DB2 Runtime Client allowing staff to freely distribute it. For more information, see “DB2 Runtime Client enhancements (Windows)” on page 64.

For the complete list of installation enhancements, see “Installation, migration, and fix pack enhancements summary” on page 69.

Related concepts:

- “Client and connectivity enhancements summary” on page 63
- “Highlights of Version 9.1 summary” on page 3
- “Installation, migration, and fix pack enhancements summary” on page 69
- “Manageability enhancements summary” on page 77
- “Performance enhancements summary” on page 89
- “Scalability enhancements summary” on page 95

Hardware and operating system support

Growing business data processing needs for faster, more scalable applications are driving requirements for increased processing power from hardware. 64-bit server support paves the way towards delivering greater processing power, memory, and application performance. The AIX, HP, Solaris, Linux on POWER™, Linux on zSeries®, Linux for AMD64 and Intel® EM64T, Linux on IA64, Windows on X64, and Windows on IA64 operating systems all support 64-bit processors with default 64-bit kernels and default 64-bit user spaces.

In response to market demand while also building on the 64-bit DB2 server support introduced in DB2 UDB Version 8, IBM is phasing out 32-bit DB2 database server support on several platforms and making DB2 database server support on 64-bit hardware and operating systems a higher priority.

For a complete list of the supported environments for 32-bit and 64-bit DB2 server instances, as well as for 32-bit and 64-bit client instances, refer to the DB2 Database for Linux, UNIX, and Windows system requirements Web page at www.ibm.com/software/data/db2/udb/sysreqs.html.

Related concepts:

- “Highlights of Version 9.1 summary” on page 3
- “Installation, migration, and fix pack enhancements summary” on page 69

Chapter 2. DB2 Connect enhancements

DB2 Connect enhancements summary

DB2 Connect provides fast and robust connectivity to IBM mainframe databases for e-business and other applications running on Linux, UNIX, and Windows operating systems. DB2 Connect has several connection solutions, including DB2 Connect Personal Edition, and a number of DB2 Connect server products. A DB2 Connect server is a connectivity server that concentrates and manages connections from multiple desktop clients and Web applications to DB2 database servers running on host or iSeries™ systems. IBM's DB2 UDB for iSeries, DB2 UDB for OS/390®, DB2 UDB for z/OS, and DB2 Server for VSE & VM continue to be the systems of choice for managing most critical data for the world's largest organizations. While these host and iSeries databases manage the data, there is a great demand to integrate this data with applications running on Linux, UNIX, and Windows operating systems.

DB2 Connect servers enable local and remote client applications to create, update, control, and manage DB2 databases and host systems using:

- Structured Query Language (SQL)
- DB2 APIs (application programming interfaces)
- ODBC (Open Database Connectivity)
- JDBC (Java™ Database Connectivity)
- SQLJ (Structured Query Language for Java)
- DB2 CLI (call level interface)
- ADO .NET (Microsoft ActiveX Data Objects .NET)

Several enhancements and changes in DB2 Version 9.1 affect the functionality and capabilities of DB2 Connect. The following DB2 Version 9.1 enhancements and changes apply to DB2 Connect use.

Application development enhancements

- "BINARY, VARBINARY, and DECFLOAT data type support in .NET and CLI client applications" on page 43
- "Client support for trusted connections to DB2 for z/OS databases" on page 45
- "Command line processor (CLP) 64 KB limit for SQL statements has been removed" on page 45
- "DB2 .NET Data Provider enhancements and support for the .NET Framework 2.0" on page 45
- "Developer Workbench replaces the Development Center" on page 47
- "IBM Database Add-Ins for Microsoft Visual Studio 2005 enhancements" on page 49
- "JDBC and SQLJ enhancements" on page 52
- "New and changed development software support" on page 53
- "XML support in SQL Assist" on page 36
- "XML support in Developer Workbench" on page 34

Changes in existing functionality

- “32-bit DB2 Connect instance support changes” on page 22
- “Application ID format changed” on page 123
- “Changed parameters and output for the db2batch command” on page 133
- “Configuration parameters default value changes” on page 117
- “IBM Software Development Kit (SDK) for Java 5.x support added (AIX, Linux, and Windows)” on page 135
- “The -w option of the db2icrt, db2ilist, and db2iupdt commands is ignored” on page 145

Client and connectivity enhancements

- “A single client (DB2 Client) is provided for application development and administration” on page 64
- “Connection timeout support for database applications added” on page 63
- “DB2 Runtime Client enhancements (Windows)” on page 64
- “Internet Protocol Version 6 (IPv6) support added” on page 65

Discontinued and deprecated functionality

- “CLISchema CLI keyword no longer supported” on page 140
- “DB2 Administration Tools are no longer supported on some platforms” on page 152
- “DB2 JDBC Type 2 driver is deprecated” on page 145
- “db2profc and db2profp utilities are discontinued” on page 150
- “db2secv82 command is deprecated” on page 141
- “Desktop icon and folder making utilities are no longer supported (Linux)” on page 151
- “NetBIOS and SNA communication protocols are no longer supported” on page 152
- “Type 3 JDBC support is discontinued” on page 153
- “VSE and VM objects are no longer supported in the DB2 Control Center” on page 154

Federated enhancements

- “Statement level isolation for nicknames” on page 68
- “Two-phase commit for multivendor data sources” on page 68
- “User mapping retrieval from an external repository is supported” on page 67

Installation, migration, and fix pack considerations

- “Application and routine migration changes” on page 106
- “Coexistence of multiple DB2 versions and fix packs enhancements (Linux and UNIX)” on page 69
- “Coexistence of multiple DB2 versions and fix packs now supported (Windows)” on page 70
- “DB2 Client easier to upgrade to DB2 Connect Personal Edition” on page 21
- “DB2 install images package format changes (Linux and UNIX)” on page 106
- “db2_deinstall command parameter changes (Linux and UNIX)” on page 106
- “Documentation updates are available between product releases” on page 107
- “Installation CD changes for national language versions of DB2 products (Linux and UNIX)” on page 108

- “Installation of DB2 products without using an Administrator user ID now supported (Windows)” on page 72
- “Licensing support changes” on page 108
- “Locally installed Information Center installation restrictions” on page 108
- “Manual installation enhancements (Linux and UNIX)” on page 73
- “Migration support changes during installation (Windows)” on page 109
- “Migration support changes for DB2 database systems” on page 110
- “Multiple CD installation now required for some products (Linux and UNIX)” on page 111
- “New DB2 product uninstall features using the db2unins command (Windows)” on page 74
- “Response file enhancements” on page 74

Manageability enhancements

- “Automated evaluation of object maintenance policies by the DB2 UDB for z/OS health monitor” on page 78
- “EXEC SQL utility support for loading data into DB2 UDB for z/OS tables” on page 82
- “First Steps interface enhancements” on page 83

Product and packaging changes

- “DB2 product-line changes” on page 27
- “DB2 Version 9.1 component replacements and name changes” on page 26

Troubleshooting and problem determination enhancements

- “Data collection improvements for database system failures” on page 101
- “Trace mask support added to the db2trc command” on page 102

Related concepts:

- “Highlights of Version 9.1 summary” on page 3

DB2 Client easier to upgrade to DB2 Connect Personal Edition

DB2 Connect Personal Edition provides access to host and iSeries DB2 data servers in addition to DB2 data servers on Linux, UNIX, and Windows. DB2 Connect Personal Edition is available for Windows and Linux operating systems.

If you want to upgrade from the DB2 Client to DB2 Connect Personal Edition, you can now do so by simply supplying the appropriate Electronic Certificate File. Providing the Electronic Certificate File enables DB2 Connect Personal Edition functionality. You no longer need to perform an entire installation process for this upgrade.

Related concepts:

- “DB2 Connect” in *DB2 Connect User’s Guide*
- “Typical steps required to install and configure DB2 Connect Personal Edition” in *Quick Beginnings for DB2 Connect Personal Edition*
- “A single client (DB2 Client) is provided for application development and administration” on page 64
- “DB2 Connect enhancements summary” on page 19

Related reference:

- “DB2 Connect product offerings” in *DB2 Connect User’s Guide*

32-bit DB2 Connect instance support changes

Growing business data processing needs for faster, more scalable applications is driving requirements for increased processing power from hardware. 64-bit server support paves the way towards delivering higher processing power, more memory, and better application performance. The AIX, HP, Solaris, Linux on AMD64 and Intel EM64T, Linux on IA64, Linux on IBM System p™, Linux on IBM System z™, Windows x64, and Windows on IA64 operating systems all support 64-bit processors with a default 64-bit kernel and a default 64-bit user space.

Responding to market demand while also building on the 64-bit DB2 server support introduced in DB2 UDB Version 8, IBM is reducing the number of supported 32-bit platforms and making DB2 server support on 64-bit hardware and operating systems a higher priority. IBM will, however, continue to support those 32-bit Windows and Linux platforms that are often preferred for building or running small and medium business applications.

DB2 Connect Version 9.1 products are supported primarily on 64-bit hardware and operating systems only. In general, migrating to DB2 Connect Version 9.1 requires that you migrate existing 32-bit DB2 Connect server instances to 64-bit DB2 Connect server instances. There are three possible DB2 Connect Version 8 to DB2 Connect Version 9 migrations:

- 32-bit Version 8 to 32-bit Version 9
- 32-bit Version 8 to 64-bit Version 9
- 64-bit Version 8 to 64-bit Version 9

Migrations from 32-bit DB2 Connect Version 8 database servers to 32-bit DB2 Connect Version 9 servers are only supported on the following operating systems:

- The Linux for x86 operating system
- The Windows for x86 operating system
- The Windows for x64 operating system (where the 32-bit DB2 Connect server for the Windows x86 operating system is supported)

To migrate existing 32-bit compiled applications who use DB2 Connect and to build 64-bit applications successfully, refer to the following topics for information.

Connection changes

- DB2 client to DB2 server connection support

Client application support changes

- Support for 32-bit database applications created in DB2 UDB or DB2 Connect Version 8 to continue working in DB2 Connect Version 9
- Changed environment variable value settings to facilitate 32-bit and 64-bit development and deployment
- Updated sample build scripts are available for building new applications

JVM support changes

- A 32-bit JVM is provided with DB2 Connect instances for Linux x86 and Windows on x86
- A 64-bit JVM is provided for all other DB2 Connect instances.

- The 64-bit JVMs are no longer provided separately on a CD.

Related concepts:

- “Migration overview for database applications and routines” in *Migration Guide*
- “Support changes for 32-bit and 64-bit DB2 servers” in *Migration Guide*

Chapter 3. Product, packaging, and terminology changes

Product, packaging, and terminology changes summary

Product, packaging and terminology changes introduced in version 9.1 include component replacements, name changes, and product-line changes. One of the terminology changes for version 9.1 affects the coding of some SQL statements.

For more information about product, packaging, and terminology changes introduced in version 9.1, follow the links in the “Related concepts” section.

Related concepts:

- “DB2 Version 9.1 component replacements and name changes” on page 26
- “DB2 product-line changes” on page 27
- “Distribution key terminology change” on page 27
- “New name for DB2 Universal Database for Linux, UNIX, and Windows” on page 25
- “New names for DB2 Information Integrator products” on page 25

New name for DB2 Universal Database for Linux, UNIX, and Windows

Beginning with the version 9.1 release, the DB2 Universal Database™ for Linux, UNIX, and Windows product name has been simplified by removing “Universal Database” and “UDB”. This change has been implemented on user interfaces, in documentation, and in packaging materials. Previous versions of DB2 database products and documentation retain “Universal Database” and “UDB” in the product naming.

Also starting in version 9.1, the term *data server* is introduced to describe the product. A data server provides software services for the secure and efficient management of structured information. DB2 Version 9.1 is a hybrid relational and XML data server.

Related concepts:

- “Product, packaging, and terminology changes summary” on page 25
- “DB2 Version 9.1 component replacements and name changes” on page 26

New names for DB2 Information Integrator products

In Version 9.1, the brand name, packaging and product names for DB2 information integration products have changed. The information integration documentation has been updated to reflect these changes.

In Version 8.2, DB2 Information Integrator was also referred to as WebSphere® Information Integrator, especially in marketing materials. The following table outlines the renaming of the information integration products and packaging changes using the V8.2 WebSphere product names.

Table 1. New names for DB2 information integration products

| Version 9.1 product | Closest Version 8.2 product | Packaging change |
|---------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------|
| WebSphere Federation Server | WebSphere Information Integrator WebSphere Information Integrator Standard Edition WebSphere Information Integrator Advanced Edition WebSphere Information Integrator Advanced Edition Unlimited | WebSphere Federation Server does not include the replication or event publishing capabilities. |
| WebSphere Replication Server | WebSphere Information Integrator Replication Edition | WebSphere Replication Server includes the Q replication but not the event publishing capability. |
| WebSphere Event Publisher | WebSphere Information Integration Event Publisher Edition | None |
| WebSphere Replication Server for z/OS | WebSphere Information Integrator Replication for z/OS | None |
| WebSphere Event Publisher for z/OS | WebSphere Information Integration Event Publisher for DB2 Universal Database for z/OS | None |

Note: The names WebSphere Data Event Publisher and WebSphere Event Publisher represent the same product. You will see WebSphere Data Event Publisher in marketing material and on the Web, and you might see either name used in the product and in the product documentation.

Related concepts:

- “Product, packaging, and terminology changes summary” on page 25

DB2 Version 9.1 component replacements and name changes

As DB2 database systems continue to evolve, the related components and component names also change. Table 2 outlines renamed product features in DB2 Version 9.1 for Linux, UNIX, and Windows:

Table 2. New names for DB2 product features

| Previous feature name | Version 9.1 feature name |
|------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------|
| IBM DB2 Universal JDBC Driver | IBM DB2 Driver for JDBC and SQLJ |
| DB2 Application Development Client and DB2 Administration Client | DB2 Client The DB2 Application Development Client and DB2 Administration Client have been combined and renamed. |
| DB2 Run-Time Client and DB2 Run-Time Client Lite | DB2 Runtime Client The DB2 Run-Time Client and DB2 Run-Time Client Lite have been combined and renamed. |

Table 2. New names for DB2 product features (continued)

| Previous feature name | Version 9.1 feature name |
|-----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------|
| Development Center | Developer Workbench |
| | The Development Center has been replaced with the Developer Workbench, which is a new separately installable feature of DB2 Version 9.1. |
| DB2 Add-in for Visual Studio .NET | IBM Database Add-Ins for Visual Studio 2005 |
| | For version 9.1, this feature is no longer included with the DB2 client and server products, and must be installed separately. |
| DB2 Geodetic Extender | DB2 Geodetic Data Management Feature |

Related concepts:

- “A single client (DB2 Client) is provided for application development and administration” on page 64
- “DB2 Runtime Client enhancements (Windows)” on page 64
- “Developer Workbench replaces the Development Center” on page 47
- “JDBC and SQLJ enhancements” on page 52
- “New name for DB2 Universal Database for Linux, UNIX, and Windows” on page 25
- “Product, packaging, and terminology changes summary” on page 25
- “IBM Database Add-Ins for Microsoft Visual Studio 2005 enhancements” on page 49
- “IBM DB2 Development Add-In overview” in *Developing ADO.NET and OLE DB Applications*

DB2 product-line changes

In version 9.1, IBM has updated the list of DB2 database products available and added several new features. For descriptions of these products and to view the related licensing and marketing information, see the DB2 Database for Linux, UNIX, and Windows home page at www.ibm.com/db2/udb.

DB2 Intelligent Miner™ Modeling, Scoring and Visualization are now part of DB2 Data Warehouse Enterprise Edition. For more information about DB2 Data Warehouse Edition, see the DB2 Data Warehouse Edition for Linux, UNIX and Windows home page at www.ibm.com/software/data/db2/udb/dwe.

Related concepts:

- “Licensing support changes” on page 108
- “New names for DB2 Information Integrator products” on page 25
- “Product, packaging, and terminology changes summary” on page 25

Distribution key terminology change

In DB2 V9.1, the term *partitioning key* is changed to *distribution key*. A distribution key is a column (or group of columns) that is used to determine the database partition in which a particular row of data is stored.

A new DISTRIBUTE BY clause replaces the PARTITIONING KEY clause used in previous releases. The old PARTITIONING KEY clause is deprecated but is supported for backwards compatibility. There is no restriction on using this old clause with the new PARTITION BY RANGE clause. Changes to the ALTER TABLE statement are as follows:

- To add a distribution key, use ADD DISTRIBUTE BY HASH instead of ADD PARTITIONING KEY.
- To drop a distribution key, use DROP DISTRIBUTION instead of DROP PARTITIONING KEY.

In DB2 V9.1, the term *table partitioning key* refers to an ordered set of one or more columns in a table. The values in the table partitioning key columns are used to determine the data partition in which each table row belongs. You define a table partitioning key by using the CREATE TABLE statement with the PARTITION BY clause. This syntax change aligns all of the data organization scheme clauses of the CREATE TABLE statement with the {DISTRIBUTE | PARTITION | ORGANIZE} BY <algorithm> pattern.

Related concepts:

- “Distribution keys” in *Administration Guide: Planning*
- “Table partitioning” in *Administration Guide: Planning*
- “Table partitioning keys” in *Administration Guide: Planning*
- “ADD PARTITIONING KEY clause of the ALTER TABLE statement is deprecated” on page 143
- “DROP PARTITIONING KEY clause of the ALTER TABLE statement is deprecated” on page 144

Related tasks:

- “Changing distribution keys” in *Administration Guide: Implementation*

Related reference:

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

Chapter 4. Native XML data store support

Native XML data store support summary

The native XML data store enables well-formed XML documents to be stored in their hierarchical form within columns of a table. XML columns are defined with the XML data type. By storing XML data in XML columns, the data is kept in its native hierarchical form, rather than stored as text or mapped to a different data model.

Because the native XML data store is fully integrated into the DB2 database system, the stored XML data can be accessed and managed by leveraging DB2 functionality.

The storage of XML data in its native hierarchical form enables efficient search and retrieval of XML. XQuery, SQL, or a combination of both can be used to query XML data. SQL functions that return XML data or take XML arguments (referred to as SQL/XML functions) also enable XML data to be constructed or published from values retrieved from the database.

XML data can only be stored in single-partition databases defined with the UTF-8 code set. Note that using XML features prevents future use of the Database Partitioning Feature available with DB2 Enterprise Server Edition for Linux, UNIX, and Windows.

Use of the new XML data type and related native XML data store support is available as a separate feature of DB2 Version 9.1. You must acquire the same licensing terms and conditions as the underlying DB2 data server. For details regarding product packaging and licensing, see the DB2 Database for Linux, UNIX, and Windows home page at www.ibm.com/db2/udb.

Related concepts:

- “Administration tools support summary - native XML data store” on page 39
- “Application development support summary - native XML data store” on page 32
- “Command line processor (CLP) and command line tool support summary - native XML data store” on page 40
- “XML schema, DTD, and external entity management using the XML schema repository (XSR)” on page 30
- “XML support in SQL statements and SQL/XML functions” on page 35
- “DB2 Net Search Extender support for the XML data type” on page 32
- “Performance enhancements summary - native XML data store” on page 37
- “XQuery builder” on page 31
- “XQuery language support” on page 30
- “Native XML data store overview” in *XML Guide*
- “XML input and output overview” in *XML Guide*

XML schema, DTD, and external entity management using the XML schema repository (XSR)

The XML schema repository (XSR) is a repository for all XML artifacts required to validate and process XML instance documents stored in XML columns. It stores copies of XML schemas, DTDs, and external entities referenced in your XML documents.

The XSR allows you to manage the dependencies XML documents have transparently, without requiring changes to the XML document content.

Related concepts:

- “XML schema, DTD, and external entity management using the XML schema repository (XSR)” in *XML Guide*
- “Native XML data store support summary” on page 29

XQuery language support

XQuery is a generalized language for querying XML data. DB2 allows XQuery to be invoked directly, obtaining data by calling functions that extract XML data from DB2 tables and views. XQuery can also be invoked from an SQL query. In this case, the SQL query can pass XML data to XQuery in the form of bound variables. XQuery supports various expressions for processing XML data and for constructing new XML objects such as elements and attributes. The programming interface to XQuery provides facilities similar to those of SQL to execute queries and retrieve query results.

XQuery is a functional programming language that was designed by the World Wide Web Consortium (W3C) to meet specific requirements for querying XML data. Unlike relational data, which is predictable and has a regular structure, XML data is highly variable. Because the structure of XML data is unpredictable, the queries that you need to perform on XML data often differ from typical relational queries. The XQuery language provides the flexibility required to perform these kinds of operations. For example, you might need to create XML queries that search XML data for objects that are at unknown levels of the hierarchy or that perform structural transformations on the data. XQuery is a strongly-typed language in which the operands of various expressions, operators, and functions must conform to expected types. The type system for XQuery is based on XML Schema.

A query consists of an optional prolog that is followed by a query body. The prolog contains a series of declarations that define the processing environment for the query. The query body consists of an expression that defines the result of the query. Expressions are the basic building blocks of a query. Expressions can be used alone or in combination with other expressions to form complex queries. DB2 supports several kinds of expressions for working with XML data, including path expressions for locating nodes within a document tree, constructors for creating XML structures within a query, and FLWOR expressions for iteration and for binding of variables to intermediate query results.

XQuery uses the XQuery and XPath data model (XDM), which represents an XML document as a hierarchy (tree) of nodes that represent XML elements and attributes. The XDM allows XQuery to operate on the abstract, logical structure of an XML document or fragment, rather than its surface syntax. The inputs (if any)

of an XQuery expression are instances of the XDM, and the result of an expression is also an instance of the XDM. XML documents are converted into the XDM when they are stored in an XML column.

DB2 supports XQuery built-in functions for working with XML data. The library includes the following types of functions: string functions, numeric functions, functions that operate on boolean values, functions that operate on QNames, functions that operate on nodes, functions on sequences, and functions that operate on durations, dates, and times.

A query that invokes XQuery directly begins with the keyword XQUERY. This keyword indicates that XQuery is being used and that the DB2 server must therefore use case sensitivity rules that apply to the XQuery language. After establishing the processing environment for the query, the query must retrieve input data. DB2 provides the following functions to retrieve input data from an XML column: `db2-fn:xmlcolumn` and `db2-fn:sqlquery`.

The Developer Workbench provides an XQuery builder for creating queries without having to understand the details of XQuery semantics.

Related concepts:

- “XQuery builder” on page 31
- “Native XML data store support summary” on page 29
- “XML support in SQL statements and SQL/XML functions” on page 35
- “XQuery” in *IBM DB2 XQuery Reference*

XQuery builder

The XQuery builder is an Eclipse-based tool to help you create queries against XML data that is in DB2 databases. The XQuery builder is part of the DB2 Developer Workbench.

With the XQuery builder, you can create complete queries without needing to understand XQuery semantics. You can build an XML query visually by selecting sample resultant nodes from a tree representation of a schema or XML document and dragging the nodes onto a return grid. After a node is listed on the return grid, you can drill down into the query to add predicates and sorting preferences. You can drill down multiple levels in a query to specify nested predicates, clauses, and expressions. For example, you might select a node and then drill down to add a predicate. Within that predicate, you might drill down again to add another predicate.

After you build your query, you can run it directly from Developer Workbench to test the query.

Related concepts:

- “Developer Workbench replaces the Development Center” on page 47
- “Native XML data store support summary” on page 29
- “XQuery language support” on page 30

DB2 Net Search Extender support for the XML data type

DB2 Net Search Extender V9.1 is now part of a separately orderable feature and fully supports the XML data type.

All Net Search Extender text search functions can be used on XML documents that are stored natively in the database. To search in specific parts of an XML document, issue an SQL search query by using the SECTION keyword. Additionally, you can combine an SQL text search query with XQuery by using the input function db2-fn:sqlquery().

Net Search Extender is now also available on Linux and Windows X64, and Linux on zSeries 64-bit.

For more information, see the DB2 Net Search Extender home page at www.ibm.com/software/data/db2/extenders/netsearch.

Related concepts:

- “Native XML data store support summary” on page 29
- “Full-text search in XML documents” in *XML Guide*

Application development support - native XML data store

Application development support summary - native XML data store

For V9.1 application development enhancements introduced to support the native XML data store, follow the links in the ‘Related concepts’ section.

Related concepts:

- “Annotated XML schema decomposition” on page 37
- “Application programming language support for XML” on page 32
- “SAMPLE database enhancements” on page 56
- “XML data type support in SQL and external procedures” on page 33
- “XML support in Developer Workbench” on page 34
- “IBM Database Add-Ins for Microsoft Visual Studio 2005 enhancements” on page 49
- “New samples” on page 55
- “XML support in SQL Assist” on page 36
- “XML support in SQL statements and SQL/XML functions” on page 35
- “XQuery builder” on page 31

Application programming language support for XML

Application development support of the new native XML data store enables applications to combine XML and relational data access and storage. Both internally encoded and externally encoded XML data are supported.

The following programming languages support the new XML data type:

- C or C++ (embedded SQL or DB2 CLI)

- COBOL
- Java (JDBC or SQLJ)
- C# and Visual Basic (DB2 .NET Data Provider)
- PHP

Java, DB2 CLI, or DB2 .NET Data Provider applications can use XML, binary, or character application data types to store XML data in or fetch XML data from XML columns. Embedded SQL applications can use XML, LOB or LOB_FILE application data types.

Stored procedures and user-defined functions can pass XML values in input or output parameters.

Related concepts:

- “Application development support summary - native XML data store” on page 32
- “XML data type support in SQL and external procedures” on page 33
- “Application programming language support for XML” in *XML Guide*
- “XML data encoding” in *XML Guide*
- “XML support in Developer Workbench” on page 34

XML data type support in SQL and external procedures

XML data can be passed to SQL procedures and external procedures by including parameters of data type XML in CREATE PROCEDURE parameter signatures. Existing procedure features support the implementation of procedural logic flow around SQL statements that produce or make use of XML values as well as the temporary storage of XML data values in variables.

Parameters of type XML are supported in:

- SQL procedures
- External procedures and external functions implemented in the following programming languages: C, C++, COBOL, Java, and .NET CLR supported languages including C# and Visual Basic.

Variables of type XML are supported in:

- SQL procedures
- External procedures and external functions implemented in the following programming languages: C, C++, COBOL, Java, and .NET CLR supported languages including C# and Visual Basic.

XML parameter and XML variable within procedures can be:

- Referenced in contexts including SQL statements where XML values are allowed
- Assigned to other variables using the following statements:
 - SELECT...INTO statement
 - VALUES...INTO statement
 - FETCH...INTO statement
 - CALL statement
 - EXECUTE ...INTO statement
 - SET statement

Related concepts:

- “Parameters in C and C++ routines” in *SQL Guide*
- “Parameters in SQL procedures” in *SQL Guide*
- “Parameters to .NET CLR routines” in *SQL Guide*
- “Effect of commits and rollbacks on XML parameter and variable values in SQL procedures” in *SQL Guide*
- “Parameters and variables of data type XML in SQL functions” in *SQL Guide*
- “XML and XQuery support in SQL procedures” in *SQL Guide*
- “Performance of routines” in *SQL Guide*
- “XML data type support in external routines” in *SQL Guide*
- “Cursors for XQuery expressions in SQL procedures” in *SQL Guide*
- “Application development support summary - native XML data store” on page 32

Related tasks:

- “Creating .NET CLR routines” in *SQL Guide*
- “Creating C and C++ routines” in *SQL Guide*
- “Creating Java routines from the command line” in *SQL Guide*
- “Creating SQL procedures” in *SQL Guide*

Related reference:

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

XML support in Developer Workbench

Developer Workbench contains the following types of XML support:

- Support for the XML data type
- Support for XML schemas
- XML document validation
- XQuery builder

Stored procedure support

- You can create stored procedures that contain XML data type parameters or return XML data types.
- You can run stored procedures that contain XML data types as input or output parameters.
- You can import XML queries that were generated by the XQuery builder into the procedure body when you are creating a stored procedure.

Data Output view support

- You can view XML data type columns on the Results page.
- For any column that can contain XML documents, you can view the content as a tree or the document text.

SQL builder support

- The XML data type is displayed anywhere that other data types are displayed.
- You can select XML functions in the Expression builder.
- You can run SQL statements that contain host variables where the column associated with the host variable is an XML data type.

- You can insert or update column values when the column value is an XML data type.

XML schema support

- From the Database Explorer in Developer Workbench, you can load existing XML schemas and XML schema documents from the XML schema repository in the database and view properties such as target namespace or schema location.
- You can register a new XML schema with its corresponding XML schema documents from the file system.
- You can drop XML schemas and XML schema documents from the XML schema repository in the database.
- You can view and edit the source for XML schema documents that make up an XML schema.
- You can annotate XML schemas for decomposition.

XML document validation

- You can edit and update an XML data type column.
- You can perform XML value validation for the XML document in the column against a registered XML schema.

XQuery builder

For details about the XQuery builder, see “XQuery builder” on page 31.

Related concepts:

- “Application development support summary - native XML data store” on page 32
- “Developer Workbench replaces the Development Center” on page 47

XML support in SQL statements and SQL/XML functions

Many SQL statements support the new XML data type. This enables you to perform many common database operations with XML data, such as creating tables with XML columns, adding XML columns to existing tables, creating indexes over XML columns, creating triggers on tables with XML columns, and inserting, updating, or deleting XML documents.

The set of SQL/XML functions, expressions, and specifications supported by the DB2 database system has been enhanced to take full advantage of the new XML data type.

You can query XML data using XQuery expressions with the XMLQUERY and XMLTABLE functions. These functions enable you to execute XQuery expressions from within an SQL context. The XMLCAST specification is useful in such cases where an XML value must be cast to an SQL value for further processing in the SQL context. (XMLCAST also supports casting from SQL types to XML types.) The XMLEXISTS predicate determines if an XQuery expression returns an empty sequence or a sequence that contains one or more items.

For validation of XML documents, the XMLVALIDATE function is available. XMLVALIDATE references XML schemas that are registered with and stored in the new XML Schema Repository (XSR). The VALIDATED predicate is used to determine if an XML document has already been validated using the XMLVALIDATE function.

XML data can be converted from XML data to character or BLOB data with the XMLSERIALIZE function. The converse operation of parsing character or BLOB data to yield XML data can be performed with the XMLPARSE function. While XMLSERIALIZE and XMLPARSE explicitly serialize and parse XML data, XML data can also be implicitly serialized and parsed by binding XML values to character, binary, and XML application data types.

New and changed SQL/XML functions allow you to construct or publish XML using the new XML data type. These functions are: XMLAGG, XMLATTRIBUTES, XMLCOMMENT, XMLCONCAT, XMLDOCUMENT, XMLELEMENT, XMLFOREST, XMLNAMESPACES, XMLPI, XMLTEXT, and XMLTABLE.

Related concepts:

- “Application development support summary - native XML data store” on page 32
- “Native XML data store support summary” on page 29
- “XML schema, DTD, and external entity management using the XML schema repository (XSR)” on page 30
- “XQuery language support” on page 30
- “Addition of XML columns to existing tables” in *XML Guide*
- “Creation of tables with XML columns” in *XML Guide*
- “Deletion of XML data from tables” in *XML Guide*
- “Insertion into XML columns” in *XML Guide*
- “Introduction to querying XML data with SQL” in *XML Guide*
- “Publishing XML values with SQL/XML” in *XML Guide*
- “Updates of XML columns” in *XML Guide*
- “XML support in triggers” in *XML Guide*

Related reference:

- “ALTER TABLE statement” in *SQL Reference, Volume 2*
- “CREATE TABLE statement” in *SQL Reference, Volume 2*
- “CREATE TRIGGER statement” in *SQL Reference, Volume 2*
- “DELETE statement” in *SQL Reference, Volume 2*
- “INSERT statement” in *SQL Reference, Volume 2*
- “Supported functions and administrative SQL routines and views” in *SQL Reference, Volume 1*
- “UPDATE statement” in *SQL Reference, Volume 2*
- “XMLCAST specifications” in *SQL Reference, Volume 1*
- “XMLEXISTS predicate” in *SQL Reference, Volume 1*
- “VALIDATED predicate” in *SQL Reference, Volume 1*

XML support in SQL Assist

The XML data type and XML functions are now supported in SQL Assist. XML functions are listed in the Expression builder. The XML data type is displayed in SQL Assist anywhere that other data types are displayed.

Related concepts:

- “Application development support summary - native XML data store” on page 32

- “Application programming language support for XML” on page 32
- “XML data type support in SQL and external procedures” on page 33

Annotated XML schema decomposition

While the new native XML data store enables you to store and access XML data as XML, in its hierarchical form, there may be cases where accessing XML data as relational data is required. An example of such a requirement is an existing application that expects and treats XML in a relational form. For such cases, annotated XML schema decomposition can be used to store content from XML documents in columns of relational tables.

Annotated XML schema decomposition is a new feature that decomposes documents based on annotations specified in an XML schema. The annotations added to XML schema documents specify details such as the name of the target table and column the XML data is to be stored in, the default SQL schema for when a target table’s SQL schema is not identified, as well as any transformation of the content before it is stored.

Related concepts:

- “Application development support summary - native XML data store” on page 32
- “XML schema, DTD, and external entity management using the XML schema repository (XSR)” on page 30
- “Annotated XML schema decomposition” in *XML Guide*

Performance enhancements - native XML data store

Performance enhancements summary - native XML data store

For V9.1 performance considerations using the native XML data store, follow the links in the ‘Related concepts’ section.

Related concepts:

- “db2batch command changes for native XML data store” on page 41
- “Indexes over XML data” on page 37
- “Optimizer support for the native XML data store” on page 39
- “Performance guidelines for XML indexing” on page 38
- “RUNSTATS command support for the native XML data store” on page 38

Indexes over XML data

Indexing support is available for data stored in XML columns. The use of indexes over XML data can improve the efficiency of queries issued against XML documents. Similar to a relational index, an index over XML data indexes a column. They differ, however, in that a relational index indexes an entire column, while an index over XML data indexes a part of a column. You indicate which parts of an XML column are indexed by specifying an XML pattern, which is a limited XPath expression.

Related concepts:

- “Native XML data store support summary” on page 29

- “Performance enhancements summary - native XML data store” on page 37
- “Performance guidelines for XML indexing” on page 38
- “XML data indexing overview” in *Performance Guide*

Related reference:

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

Performance guidelines for XML indexing

Native XML data store allows you to index XML data stored in XML columns. The optimizer supports these indexes over XML data when evaluating SQL and XQuery statements.

Unlike relational indexes, indexes over XML data provide access to nodes within the document by creating index keys based on XML patterns. Some of the considerations for efficient query evaluation that apply to relational indexes also apply to indexes over XML data. For example, up-to-date statistics are required for the most efficient access to XML data. Other performance considerations apply only to indexes over XML data. For example, how restrictive an index over XML data is determines whether it can be used by a query statement.

The DB2 Information Center has been updated with a core set of guidelines that allow you to work with SQL and XQuery statements efficiently.

Related concepts:

- “Indexes over XML data” on page 37
- “Native XML data store support summary” on page 29
- “Performance enhancements summary - native XML data store” on page 37
- “Guidelines for matching indexes with queries overview” in *Performance Guide*
- “Query tuning guidelines” in *Performance Guide*

Related tasks:

- “Creating an index” in *Administration Guide: Implementation*

RUNSTATS command support for the native XML data store

The RUNSTATS command has been updated to support the collection of statistics on tables containing XML columns and on indexes over XML data.

These statistics are used by the optimizer to determine the optimal access path to XML data stored in XML columns. Up-to-date statistics are required for the most efficient access.

Related concepts:

- “Native XML data store support summary” on page 29
- “Performance enhancements summary - native XML data store” on page 37

Related reference:

- “RUNSTATS command” in *Command Reference*

Optimizer support for the native XML data store

The optimizer has been updated to support the evaluation of SQL, XQuery, and SQL/XML functions that embed XQuery, against XML and relational data. The optimizer exploits statistics gathered over XML data, as well as data from indexes over XML data, to produce efficient query execution plans.

Related concepts:

- “Native XML data store support summary” on page 29
- “Performance enhancements summary - native XML data store” on page 37
- “Guidelines for matching indexes with queries overview” in *Performance Guide*

Administration tools support - native XML data store

Administration tools support summary - native XML data store

For V9.1 administration tools enhancements introduced to support the native XML data store, follow the links in the ‘Related concepts’ section.

Related concepts:

- “Control Center support for native XML data store” on page 39
- “Explain and Visual Explain support for SQL/XML and XQuery statements” on page 40

Control Center support for native XML data store

The Control Center has been updated to support the native XML data type for many of its administrative functions. This allows database administrators to work with XML data alongside relational data from within a single GUI tool.

Examples of supported administrative tasks are:

- Creating tables with XML columns
- Creating indexes over XML columns using the new Create Index wizard
- Viewing the contents of XML documents stored in XML columns
- Working with the XML schemas, DTDs, and external entities required to validate and process XML documents
- Collecting statistics on tables containing XML columns

Related concepts:

- “Administration tools support summary - native XML data store” on page 39
- “Command line processor (CLP) and command line tool support summary - native XML data store” on page 40
- “DB2 command line processor support for the native XML data store” on page 40
- “Native XML data store support summary” on page 29

Related reference:

- “db2cc - Start control center command” in *Command Reference*

Explain and Visual Explain support for SQL/XML and XQuery statements

The Explain facility and the Visual Explain GUI tool have been updated to support SQL enhancements for querying XML data and to support XQuery statements. These updates to the Explain facility and to the Visual Explain GUI tool allow you to see quickly how the DB2 database evaluates query statements against XML data.

Several operators are provided to explain statements issued against XML data stored in XML columns. Query cost estimates are provided, along with optimizer output that shows how statements issued against XML data are evaluated, including optimizer use of indexes over XML data.

Related concepts:

- “Visual Explain overview” in *Administration Guide: Implementation*
- “Administration tools support summary - native XML data store” on page 39
- “Native XML data store support summary” on page 29

Related reference:

- “Explain tables” in *SQL Reference, Volume 1*

Command line processor (CLP) and command line tool support - native XML data store

Command line processor (CLP) and command line tool support summary - native XML data store

For V9.1 enhancements to the DB2 command line processor (CLP) and system command utilities introduced to support the native XML data store in Version 9.1, follow the links in the ‘Related concepts’ section.

Related concepts:

- “db2batch command changes for native XML data store” on page 41
- “db2look command changes for native XML data store” on page 41
- “DB2 command line processor support for the native XML data store” on page 40
- “Import and export utility support for the native XML data store” on page 42

DB2 command line processor support for the native XML data store

Several DB2 commands have been updated or added to support the native storage of XML data. These updates allow you to work with XML data alongside relational data from the DB2 command line processor (CLP).

Examples of tasks that you can perform on XML data from the CLP include:

- Issuing XQuery statements by prefixing them with the XQUERY keyword.
- Importing and exporting XML data.
- Collecting statistics on XML columns.

- Calling stored procedures with IN, OUT, or INOUT parameters of XML data type.
- Working with the XML schemas, DTDs, and external entities required to validate and process XML documents.
- Reorganizing indexes over XML data and tables containing XML columns.
- Decomposing XML documents.

Related concepts:

- “Command line processor (CLP) and command line tool support summary - native XML data store” on page 40
- “db2batch command changes for native XML data store” on page 41
- “db2look command changes for native XML data store” on page 41

db2batch command changes for native XML data store

The **db2batch** command has been updated to process both SQL and XQuery statements. Users can issue XQuery statements by prefixing them with the XQUERY keyword.

In addition to native XML data store support, the **db2batch** command has undergone other changes that affect some of its options, option parameters, and output. See the links in “Related concepts” for more information.

Related concepts:

- “Command line processor (CLP) and command line tool support summary - native XML data store” on page 40
- “Changed parameters and output for the db2batch command” on page 133

Related reference:

- “db2batch - Benchmark tool command” in *Command Reference*

db2look command changes for native XML data store

The **db2look** command has been updated to allow you to reproduce the database objects required to validate and process XML documents. These include the XML schemas, DTDs, and external entities registered with the XML schema repository (XSR).

The **db2look** command can export all the XSR objects required to validate and process XML documents, along with the DDL statements needed to register them at the target database.

Related concepts:

- “Command line processor (CLP) and command line tool support summary - native XML data store” on page 40
- “Statistics for modeling production databases” in *Performance Guide*

Related reference:

- “db2look - DB2 statistics and DDL extraction tool command” in *Command Reference*

Import and export utility support for the native XML data store

The import and export utilities have been updated to support the native XML data type. These utilities treat XML data like LOB data: both types of data are stored outside the actual table. Application development support for importing and exporting XML data is also provided by updated db2Import and db2Export APIs.

These updated utilities permit data movement of XML documents stored in XML columns that is similar to the data movement support for relational data.

Related concepts:

- “Export Overview” in *Data Movement Utilities Guide and Reference*
- “Administration tools support summary - native XML data store” on page 39
- “Import Overview” in *Data Movement Utilities Guide and Reference*
- “Control Center support for native XML data store” on page 39

Related tasks:

- “Exporting data” in *Data Movement Utilities Guide and Reference*
- “Importing data” in *Data Movement Utilities Guide and Reference*

Related reference:

- “EXPORT command” in *Command Reference*
- “IMPORT Command” in *Command Reference*

Chapter 5. Application development enhancements

Application development enhancements summary

DB2 Version 9.1 provides numerous enhancements that simplify database application development. Enhancements include a new Developer Workbench, enhanced functionality for Visual Studio 2005 and DB2 .NET Data Provider, and new application samples. DB2 Version 9.1 also introduces integrated debugging capabilities and many other features that allow you to reduce development time.

For application development enhancements introduced in DB2 Version 9.1, follow the links in the "Related concepts" section.

Related concepts:

- "Application development support summary - native XML data store" on page 32
- "BINARY, VARBINARY, and DECFLOAT data type support in .NET and CLI client applications" on page 43
- "JDBC and SQLJ enhancements" on page 52
- "New and changed development software support" on page 53
- "New samples" on page 55
- "Client support for trusted connections to DB2 for z/OS databases" on page 45
- "Command line processor (CLP) 64 KB limit for SQL statements has been removed" on page 45
- "DB2 .NET Data Provider enhancements and support for the .NET Framework 2.0" on page 45
- "Developer Workbench replaces the Development Center" on page 47
- "External table function support across database partitions" on page 49
- "IBM Database Add-Ins for Microsoft Visual Studio 2005 enhancements" on page 49
- "Java routine class loader enhancements" on page 50
- "SAMPLE database enhancements" on page 56
- "TRIM and STRIP scalar functions added" on page 57
- "XML support in Developer Workbench" on page 34
- "XQuery builder" on page 31

BINARY, VARBINARY, and DECFLOAT data type support in .NET and CLI client applications

The data types BINARY and VARBINARY have been added to DB2 for z/OS Version 9 and support for the types has been added to DB2 CLI and DB2 .NET Data Provider.

DB2 CLI support of BINARY and VARBINARY:

The symbolic SQL data type SQL_BINARY, which is currently used to represent the data type CHAR FOR BIT DATA, is now also used to represent the BINARY

data type. If a function reports the data type of a column as being SQL_BINARY you will not be sure whether the type on the server is BINARY or CHAR FOR BIT DATA. The two types are only different in that the CHAR FOR BIT DATA data type pads short values with the byte value 0x20 (spaces) but the BINARY data type pads short values with the byte value 0x00.

The symbolic SQL data type SQL_VARBINARY, which is currently used to represent the data type VARCHAR FOR BIT DATA, is now also used to represent the VARBINARY data type. There is no difference in behavior between the VARBINARY and VARCHAR FOR BIT DATA data types.

The existing C symbolic data type SQL_C_BINARY can be used to hold values from both the BINARY and VARBINARY data types.

DB2 .NET Data Provider support of BINARY and VARBINARY:

A new data type named DB2Binary has been added that will accept CHAR FOR BIT DATA, VARCHAR FOR BIT DATA, BINARY, and VARBINARY data types. The enumeration value DB2Type.Binary has been added to represent the SQL data type BINARY. The enumeration value DB2Type.VarBinary has been added to represent the SQL data type VARBINARY.

Support is provided in both DB2 CLI and DB2 .NET Data Provider for the new DECFLOAT data type supported on DB2 for z/OS.

DB2 CLI support of DECFLOAT:

The C data types, SQLDECIMAL64 and SQLDECIMAL128, have been added for storing decimal float values. A wide range of conversions are supported for the types. A new connection setting named SQL_ATTR_DECFLOAT_ROUNDING_MODE allows the client to specify what type of rounding should occur if any operations on the server side require rounding of a decimal float value.

DB2 .NET Data Provider support of DECFLOAT:

A new data type named DB2DecimalFloat has been added for storing decimal float values. The DB2DecimalFloat data type will hold both DECIMAL(16) and DECIMAL(34) values. No math operations are currently supported on the DB2DecimalFloat data type but the type can be converted to and from both the Decimal and Double data types, which do support math operations.

The DB2Type.DecimalFloat16 and DB2Type.DecimalFloat34 enumeration values have been added to represent the DECFLOAT(16) and DECFLOAT(34) SQL data types.

Related concepts:

- “Application development enhancements summary” on page 43

Related reference:

- “SQL symbolic and default data types for CLI applications” in *Call Level Interface Guide and Reference, Volume 1*

Client support for trusted connections to DB2 for z/OS databases

DB2 CLI and IBM DB2 Driver for JDBC and SQLJ now supports making trusted connections to DB2 database servers that support trusted contexts. Trusted connections can acquire special set of privileges that are not available to it outside the trusted context. This allows end to end identity control and better auditing to help in compliance and government regulation issues.

If the database server is configured to allow it to do so, a client can create trusted connections using ODBC, XA, or new Java methods. The user name associated with the trusted connection can then be switched without the database server having to fully authenticate the new name.

Trusted contexts are currently supported only on DB2 Version 9.1 for z/OS.

Related concepts:

- “Application development enhancements summary” on page 43
- “DB2 Connect enhancements summary” on page 19
- “IBM DB2 Driver for JDBC and SQLJ trusted context support” in *Developing Java Applications*
- “Trusted connections through DB2 Connect” in *DB2 Connect User’s Guide*

Related tasks:

- “Creating and terminating a trusted connection through CLI” in *DB2 Connect User’s Guide*
- “Switching users on a trusted connection through CLI” in *DB2 Connect User’s Guide*

Command line processor (CLP) 64 KB limit for SQL statements has been removed

A Command line processor (CLP) limit of 64 KB for SQL statements and for CLP commands that contain SQL statement components has now been removed. In previous releases, SQL statements generated by other DB2 tools were not valid for use within the CLP when they exceeded the old CLP 64 KB limit. The new CLP limit of approximately 2 MB is comparable with the limits on the other DB2 tools.

Related reference:

- “Using command line SQL statements and XQuery statements” in *Command Reference*

DB2 .NET Data Provider enhancements and support for the .NET Framework 2.0

The DB2 .NET Data Provider now supports the Microsoft .NET Framework, Version 2.0, and has additional new features to help you develop more powerful .NET applications.

Support for the System.Data.Common base classes:

The DB2 .NET Data Provider supports the use of the common base classes declared in the System.Data.Common namespace. This enables you to develop a generic .NET database application without referencing any data provider-specific classes. For example, you can use the generic DBConnection class, instead of the DB2Connection class from the DB2 .NET Data Provider. Where applicable, the DB2 .NET Data Provider classes are inherited from their corresponding classes in the System.Data.Common namespace.

Instances of these common base classes are created from a Factory class. You can use the generic Factory class, which is created as follows:

```
DbProviderFactory myFactory = DbProviderFactories.GetFactory("IBM.Data.DB2");
```

You can also use the DB2 .NET Data Provider factory class, DB2Factory, which is created as follows:

```
DB2Factory myFactory = DB2Factory.Instance;
```

DB2Types classes:

The DB2Types classes provide a means to represent DB2 database column values as individual nullable objects. Instances of DB2Types classes are also useful as parameters for CLR stored procedures or user-defined functions. Following is a list of the DB2Types classes:

- DB2Binary
- DB2Blob
- DB2Clob
- DB2Date
- DB2Decimal
- DB2DecimalFloat
- DB2Double
- DB2Int16
- DB2Int32
- DB2Int64
- DB2Real
- DB2Real370
- DB2RowId
- DB2String
- DB2Time
- DB2TimeStamp
- DB2Xml

All of these DB2Types classes belong to the IBM.Data.DB2Types namespace.

Scrollable and updatable result sets:

The DB2 .NET Data Provider now contains a DB2ResultSet class, which contains functions that enable your applications to update individual column values, or to scroll through result sets backwards or forwards.

Data paging capability:

The DB2Command class now has an ExecutePageReader method, which enables applications to fetch a specific set of rows from the database. The

ExecutePageReader method accepts values for the starting row number, and the number of rows to be fetched, and returns a DataReader object. This feature is useful when looking for a particular set of rows, and is much faster and simpler than finding the desired rows by scrolling through the entire result set.

Bulk data copy:

With the bulk data copy facility in the DB2 .NET Data Provider, you can copy data from an available data source into a DB2 database table. To perform a bulk data copy operation, you first define the column mappings from the data source to the DB2 table where the data is to be copied, and then perform the actual copy of the data. The bulk data copy facility is accessible through the DB2BulkCopy, DB2CopyColumnMapping and DB2CopyColumnMappingCollection classes.

Update batch size:

The DB2DataAdapter property, UpdateBatchSize enables applications to determine the number of SQL statements to collect before issuing them as a batch to the DB2 database server. This can provide a performance boost, as it will result in fewer individual transmissions of data between the client application and the database server.

Related concepts:

- “DB2 .NET Data Provider” in *Developing ADO.NET and OLE DB Applications*
- “Application development enhancements summary” on page 43

Developer Workbench replaces the Development Center

The Development Center from DB2 UDB for Linux, UNIX, and Windows Version 8 is replaced in DB2 V9.1 by an Eclipse-based tool called Developer Workbench. The Developer Workbench includes functionality that is comparable to Development Center. In addition to existing Development Center functionality, there are some additional new features as described in the following section.

Developer Workbench information center and tutorials

Developer Workbench help and tutorials are available in an information center that is installed with Developer Workbench. This information is for Developer Workbench only, and it is not installed with the DB2 information center CD. To access the Developer Workbench help and tutorials, click **Help > Help Contents** from the main menu in the product. You can also link directly to important getting started information from the Welcome page in Developer Workbench by clicking **Help > Welcome**.

Migrate existing Development Center projects

You can use a wizard to migrate existing Development Center projects into Developer Workbench.

Compare routines

You can compare and make changes between two routines that are contained within a data development project in Developer Workbench. You can also compare routine attributes for routines that are stored on a server.

Deploy routines to unlike servers

You can deploy routines that were created for one DB2 database to a DB2 database on a different platform. For example, you can create a routine for a DB2 for Linux, UNIX, and Windows database and then deploy it to a DB2 for z/OS database. Not all server combinations are supported.

Binary deploy

For SQL or Java stored procedures targeting DB2 UDB for z/OS Version 8 or higher, you can deploy without going through a full rebuild. The binaries for a SQL procedure or the JAR for a Java procedure are copied from the source to the target system.

Launch Visual Explain

You can launch Visual Explain for DB2 for z/OS or DB2 for Linux, UNIX, and Windows SQL statements from either the routine editor for SQL routines or from the wizard that is used to create a routine.

Develop SQLJ applications

You can develop SQLJ applications by using the following features:

- Generate an SQLJ template file by using a wizard
- Translate and compile automatically
- Customize by using a wizard
- Print the profile file
- Edit SQLJ applications by using code assist and templates
- Debug SQLJ files

Team support

You can share your Developer Workbench data development project by using either CVS or ClearCase®. After you share your project, you can manage all changes and update history, and you can synchronize your files with the repository.

Table data editing

You can use an editor to edit the data that is contained in a table. You can edit existing values, delete an existing row, or insert a new row.

Data extract and load

You can extract the data from a table or view into a file on the local file system. You can use this file to load the data into a table.

Stored procedure debugger

Developer Workbench includes integrated stored procedure debugging capabilities. You can debug SQL or Java stored procedures that target supported DB2 servers, or Java stored procedures that target supported Derby servers.

XML support

Developer Workbench contains support for XML functions, the XML data type, and XML schema registration. You can also create XQueries with the XQuery builder. See the links under the "Related Concepts" section for more details.

Developer Workbench will also provide the following additional features when DB2 V9.1 for z/OS becomes available:

Multiple JAR support

You can create Java stored procedures that depend on code that is contained in multiple JAR files that are installed on the server. You can also package multiple Java stored procedures within the same JAR file on the server.

SQL procedure versioning for z/OS

There is support for native SQL stored procedures and versioning of native SQL stored procedures targeting DB2 for z/OS servers.

Package variation support for SQL and SQLJ Java stored procedures

You can create package variations from the Database explorer for SQL and SQLJ stored procedures targeting DB2 for z/OS. These package variations are used for creating copies of existing packages with different bind options.

Related concepts:

- “Application development enhancements summary” on page 43
- “XML support in Developer Workbench” on page 34
- “XQuery builder” on page 31

External table function support across database partitions

In partitioned database environments, user-defined table functions that do not execute SQL statements can be invoked in parallel across all DB2 database partitions. User-defined table functions generally enable the extension and customization of the SQL language. In DB2 Version 9.1, table functions can be defined such that for a single table function reference, the table function is invoked on all database partitions. The union of the results of the table-function execution on each of the database partitions is returned as a single table result set.

Related concepts:

- “Application development enhancements summary” on page 43
- “External routine implementation” in *SQL Guide*
- “Routines: Table functions” in *SQL Guide*

Related reference:

- “SNAPAPPL administrative view and SNAP_GET_APPL table function – Retrieve appl logical data group snapshot information” in *Administrative SQL Routines and Views*
- “CREATE FUNCTION (External Table) statement” in *SQL Reference, Volume 2*

IBM Database Add-Ins for Microsoft Visual Studio 2005 enhancements

The IBM Database Add-Ins for Microsoft Visual Studio 2005 provides tools for rapid application development, database schema development, and debugging.

- Database activity is now performed in the Microsoft Server Explorer. The IBM Server Explorer has been removed from the IBM Database Add-Ins for Microsoft Visual Studio 2005 but still remains for IBM DB2 Development Add-In for Microsoft Visual Studio .NET 2003.
- You can build Windows applications and Web sites for the IBM DB2 Database without writing any code.
- There is support to generate and deploy IBM Web Services on DB2 Embedded Application Server and for Microsoft Web Services using Microsoft web service projects. Web services can be created and deployed without writing a single line of code. Support includes the ability to deploy and alter Web services, test-run Web services, and browse previously deployed web services.

All DB2 database products are supported: DB2 Database for Linux, UNIX, and Windows, DB2 Universal Database for iSeries, and DB2 Universal Database for z/OS.

- New IBM designers provide an easier way for you to work with database objects. With the designers, you can perform the following actions:
 - Create and alter tables, views, and procedures.
 - Create and alter roles and define access privileges for tables, views and procedures.
 - Clone tables and procedures.
 - View or create scripts for all database objects.
 - Show data from tables and views, and filter the data based on an SQL condition.
 - Import and export data to a table or view.
- You can run scripts before and after you run procedures and functions, and you can save input or in-out parameter values across Visual Studio sessions. You can commit or roll back transactions.
- New user interface gives you the ability to view single or multiple result sets for a procedure. Where possible, you can discover result sets automatically, and you can manually define or customize a result set definition.
- You can seamlessly debug SQL procedures on Linux, UNIX, and Windows or zSeries servers.

There is full support for the DB2 native XML store, including the following functionality:

- Use an XML data type for columns and procedures
- Provide an XML index for an XML column
- Have the capability to visualize XML data
- Update, import, and export XML data
- Validate an XML database against a registered XML schema
- Register and unregister XML schemas
- Generate sample data based on an XML schema
- Create and register annotated XML schemas
- Execute and visualize XQuery and SQL/XML scripts
- Apply XSLT to XML data for customized visualization

Related concepts:

- “Application development enhancements summary” on page 43
- “IBM DB2 Development Add-In overview” in *Developing ADO.NET and OLE DB Applications*

Java routine class loader enhancements

Application developers who develop Java routines can now utilize multiple Java classes that have the same name by including them in separate JAR files and by explicitly identifying them in routine create statements.

When creating a Java procedure, the specification of a *jar-id* to identify a jar file:

- Ensures that the correct Java class is located and loaded for the procedure.
- Improves the performance with which a Java class is located and loaded at run time.

Previously, upon a Java procedure call, the DB2 database manager searched through a list of cached Java classes and the set of classes defined within the

CLASSPATH path value for the first Java class with a class name matching the one specified in the *EXTERNAL NAME* clause of the *CREATE PROCEDURE* or *CREATE FUNCTION* statement used to create the routine without considering the name of the JAR file in which the class might reside.

The specification of a JAR file name within the *EXTERNAL NAME* clause allows DB2 to more quickly resolve and load the Java class associated with the routine.

To make use of this new functionality with new and existing routines:

For new Java procedures and functions:

To make use of this new functionality when creating new Java procedures, be sure to specify the *jar_id* clause option in the *CREATE PROCEDURE* (external) or the *CREATE FUNCTION* (external) statement which specifies the jar identifier of the JAR collection on the database server that contains the Java class that implements the routine. This is a recommended best practice.

For existing Java procedures and functions:

If a JAR file for a Java routine was installed on the database server using the system-defined procedure *install_jar*, but no *jar-id* was explicitly specified in the *EXTERNAL NAME* clause of the *CREATE* statement for the routine, it is recommended that these routine definitions be modified so that the *EXTERNAL NAME* clause explicitly specifies the *jar-id* value. This can be done using the *ALTER PROCEDURE* or *ALTER FUNCTION* statement.

For example, consider an existing Java procedure named *myproc* that references a method *mymethod* in the class *myclass* within the jar file associated with the *jar-id* *myjar* that was created using the following SQL statement:

```
CREATE PROCEDURE myproc
LANGUAGE JAVA
PARAMETER STYLE JAVA
EXTERNAL NAME 'myclass!mymethod'
FENCED
```

The following SQL statement can be used to update this definition such that it explicitly specifies the *jar-id*:

```
ALTER PROCEDURE myproc() EXTERNAL NAME 'myjar:myclass!mymethod'
```

The definition for the procedure that results in the database is as follows:

```
CREATE PROCEDURE myproc
LANGUAGE JAVA
PARAMETER STYLE JAVA
EXTERNAL NAME 'myjar:myclass!mymethod'
FENCED
```

When the procedure *myproc* is called, the class *myclass* is directly loaded from the jar file associated with *jar-id* *myjar*.

Specifying the *jar_id* clause option when creating Java routines is a recommended best practice.

Related concepts:

- “Application development enhancements summary” on page 43
- “External routine implementation” in *SQL Guide*

Related tasks:

- “Creating external routines” in *SQL Guide*
- “Creating Java routines from the command line” in *SQL Guide*

Related reference:

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

JDBC and SQLJ enhancements

The IBM DB2 Driver for JDBC and SQLJ contains the following major enhancements for DB2 V9.1:

- Compliance with the JDBC 3.0 specification.
- Support for the new XML column data type. This support includes:
 - New DB2-only methods and a new DB2-only data type to support update and retrieval of data in XML columns, and stored procedure calls with XML parameters.
 - New DB2-only methods for performing XML schema registration.
- New DB2-only methods to support trusted connections

These new methods enable an application server to use a system authorization ID to establish a connection to a DB2 for z/OS database server and then reuse that connection on behalf of a new user.

- Support for heterogeneous pooling and connection reuse

Connection pooling is a framework for caching physical data source connections, which are equivalent to DB2 threads. When JDBC reuses physical data source connections, the expensive operations that are required for the creation and subsequent closing of `java.sql.Connection` objects are minimized. The IBM DB2 Driver for JDBC and SQLJ provides a factory of pooled connections that are used by WebSphere Application Server or other application servers. The application server actually does the pooling.

Previously, the IBM DB2 Driver for JDBC and SQLJ supported homogeneous connection pooling, in which all `Connection` objects that come from a connection pool should have the same properties. With DB2 V9.1, the IBM DB2 Driver for JDBC and SQLJ also supports heterogeneous pooling, in which `Connection` objects with different properties can share the same connection pool.

- SSL support

The IBM DB2 Driver for JDBC and SQLJ now includes SSL (Secure Sockets Layer) support to database servers that also have SSL support. When an application sets the JDBC driver `sslConnection` property during connection to a database server, the driver obtains the connection using an SSL socket.

- Tolerable errors

This enhancement changes the behavior of the IBM DB2 Driver for JDBC and SQLJ when a query returns a +100 SQLCODE. Previously, when a call to `ResultSet.next` returned false (no more rows), the IBM DB2 Driver for JDBC and SQLJ did not generate an `SQLWarning`. With DB2 V9.1, the driver accumulates an `SQLWarning` with return code +4202 when errors are encountered and tolerated as specified by the `RETURN DATA UNTIL` clause.

- `sendDataAsIs` property

Previously, the IBM DB2 Driver for JDBC and SQLJ always did conversion from application data types to column data types if DESCRIBE information was available. Some anomalies occurred as a result of this automatic conversion. With DB2 V9.1, the IBM DB2 Driver for JDBC and SQLJ includes the `sendDataAsIs` connection property to let the application specify whether the driver should do data type conversion. If the `sendDataAsIs` property is set to true, the driver assumes the data type based on the `setXXX` method that is used.

- Enhancements for connections to DB2 for z/OS database servers
 - Support for progressive streaming for retrieval of LOBs and XML data.

When the IBM DB2 Driver for JDBC and SQLJ streams through large LOB or XML data progressively, it can perform efficient and fast data retrieval when the database server supports the new progressive streaming architecture. DB2 Version 9.1 for z/OS supports progressive streaming.

With progressive streaming, the database server dynamically determines the most efficient mode in which to return LOB or XML data, based on the size of the LOBs or XML objects.
 - Support for the new DB2 for z/OS BINARY, VARBINARY, and DECFLOAT column data types.

Applications can now use existing Java data types to retrieve data from or update data in BINARY, VARBINARY, and DECFLOAT columns in tables on DB2 for z/OS database servers.

Related concepts:

- “Application development enhancements summary” on page 43
- “IBM DB2 Driver for JDBC and SQLJ trusted context support” in *Developing Java Applications*
- “JDBC and SQLJ connection pooling support” in *Developing Java Applications*
- “LOBs in JDBC applications with the IBM DB2 Driver for JDBC and SQLJ” in *Developing Java Applications*
- “Java support for XML schema registration and removal” in *Developing Java Applications*
- “XML data in JDBC applications” in *Developing Java Applications*
- “XML data in SQLJ applications” in *Developing Java Applications*

Related reference:

- “Properties for the IBM DB2 Driver for JDBC and SQLJ” in *Developing Java Applications*

New and changed development software support

There are a number of changes to the list of development software titles supported for the development of DB2 applications. Following are the changes, listed by operating system.

AIX:

Java IBM Developer Kit for AIX, Java Technology Edition, Version 5

C and C++

IBM XL C/C++ Enterprise Edition Version 7.0 for AIX, IBM XL C/C++ Enterprise Edition Version 8.0 for AIX

Windows (32-bit and 64-bit):

C and C++

Microsoft Visual C++ 2005, Intel Proton Compiler for Windows 32-bit applications, Version 9.0.021 or later, Intel Proton Compiler for Windows x64, Version 9.0.024 or later

C# Microsoft Visual C# 2005 (with the 32-bit .NET Framework, Version 2.0 only)

Java IBM Developer Kit and Runtime Environment for Windows, Java Technology Edition, Version 5

Visual Basic .NET

Microsoft Visual Basic 2005 (with the 32-bit .NET Framework, Version 2.0 only)

Linux (for Intel x86, AMD 32-bit):**C and C++**

GNU/Linux gcc version 3.3.3
Intel C Compiler version 9.0

Java IBM Developer Kit and Runtime Environment for Linux, Java Technology Edition, Version 5

Linux (for IA64):**C and C++**

GNU/Linux gcc version 3.3.3
Intel C Compiler version 9.0

Java IBM Developer Kit and Runtime Environment for Linux, Java Technology Edition, Version 1.4.2 SR3

Linux (for zSeries, x86_64):**C and C++**

GNU/Linux gcc version 3.3.3

Java IBM Developer Kit and Runtime Environment for Linux, Java Technology Edition, Version 5

Linux (for POWER):**C and C++**

GNU/Linux gcc version 3.3.3
IBM XL C for Linux version 7.0

Java IBM Developer Kit and Runtime Environment for Linux, Java Technology Edition, Version 5

Solaris:**C and C++**

Forte C Versions 6.2
Sun ONE Studio versions 9, 10 Compiler Collection

HP-UX on PA-RISC:

C HP ANSI C Compiler Version B.11.11.14

C++ HP aC++ Compiler Version A.03.63

HP-UX on IPF (IA64):

C HP ANSI C Compiler Version A.06.05

C++ HP aC++ Compiler Version A.06.05

Related concepts:

- “Application development enhancements summary” on page 43

Related reference:

- “Supported operating systems for database application development” in *Getting Started with Database Application Development*
- “Supported programming languages and compilers for database application development” in *Getting Started with Database Application Development*

New samples

DB2 sample programs can be used as templates to create your own application programs and as learning tools to understand DB2 product functionality. For DB2 V9.1, there are many new sample programs available that demonstrate key features and enhancements. These include:

XQuery support

Axes, FLWOR expression. SQL/XML and nested XQueries

XML Schema support

Schema registration and validation of XML document as per the schema

XML DDL support

Alter table to add XML column

XML support for insert, update, and delete operations

Inserting the new XML value in XML type column, updating and deleting the existing values

XML support for parsing, validating, and serializing

Implicit and Explicit parsing of compatible data types.

Hybrid use of SQL and XQuery

SQL/XML functions like XMLTable, XMLQuery and XMLExists

Support parameter type of XML for PSM/STP

Stored procedure with XML data type parameter

XML Values Index support

Indexes on different node types of XML value

SQL/XML constructors

Creating a new XML value from relational data

XML Decomposition

Decomposition of XML document as per annotated schema

Utility support for XML

XML data type support for import, export, runstats, db2look, and db2batch utilities

Table Partitioning

Table partitioning, roll-in, roll-out, alter partition features

Online SET INTEGRITY

Set integrity after altering a partitioned table

Automatic Storage enhancements

Single point of management with multiple paths, auto-resize for tablespaces

Label-Based Access Control

LBAC features to control access basing on security policy, security label

Administration through SQL interfaces including ADMIN_CMD support

ADMIN_CMD support for backup, UPDATE CONFIGURATION, export, and other tasks

Row compression

Compressing the table at row level

Support for larger indexes

Creating large index key parts

Autonomic features

Autonomic feature default support with the new extended limit

Large RID support

Increased RID size, migrating from regular to large DMS tablespace

These and other samples can be found in the following location:

- On Windows: %DB2PATH%\sql1lib\samples (where %DB2PATH% is a variable that determines where DB2 is installed)
- On UNIX: \$HOME/sql1lib/samples (where \$HOME is the home directory of the instance owner)

Related concepts:

- “Sample files” in *Samples Topics*
- “Application development enhancements summary” on page 43
- “SAMPLE database enhancements” on page 56

SAMPLE database enhancements

Sample databases that demonstrate new and previously not demonstrated DB2 database features can be created from First Steps or by using a new *db2sampl* command. Sample databases can be used as a learning tool or as a database for trying out new DB2 features.

Sample databases can be created that include the following database features:

- Tables with XML data type columns that contain XML data
- Indexes over XML data type columns
- XML schema repository containing XML schema documents
- Primary keys
- Foreign keys and referential integrity constraints
- Indexes over relational data type columns
- Check constraints
- Triggers (includes before-triggers and after-triggers)
- Views
- SQL user-defined function
- SQL stored procedure
- Updated CLOB data source files

- Multi-dimensional clustered tables

The *db2sampl* command provides options and support for specifying:

- the directory in which to create the database files
- the name of the sample database
- the schema with which to qualify database object names
- the optional inclusion of relational database objects and data
- the optional inclusion of XML database objects and data
- the forced removal of databases with the same name prior to the sample database creation
- the command run in verbose mode and output status messages to standard output

Some new features of the *db2sampl* command behavior include:

- Improved error handling
- Improved error output redirected to the administration log file
- Improved performance

When the default options are specified in the *db2sampl* command the sample database that is created can be used in conjunction with sample applications to demonstrate DB2 application and administration features.

The database objects and data of the DB2 Version 8 SAMPLE database can still be created by specifying the *-v8* option in the *db2sampl* command.

Related concepts:

- “Application development enhancements summary” on page 43

Related reference:

- “The SAMPLE database” in *Samples Topics*
- “db2sampl - Create sample database command” in *Command Reference*

TRIM and STRIP scalar functions added

Support for the TRIM scalar function is added in Version 9.1. This function is used to remove blanks or occurrences of another specified character from the end or the beginning of a string expression.

Support for the STRIP function, which is identical to the TRIM function, is also supported for compatibility with legacy applications on other platforms.

Related reference:

- “STRIP scalar function” in *SQL Reference, Volume 1*
- “TRIM scalar function” in *SQL Reference, Volume 1*

Chapter 6. Backup, logging, and recovery enhancements

Backup, logging, and recovery enhancements summary

DB2 Version 9.1 backup, logging, and recovery enhancements include new features such as the ability to restart an interrupted recovery operation and automatically generated scripts for performing redirected restore operations. These enhancements and others will help you ensure the availability of your data.

For backup, logging, and recovery enhancements introduced in DB2 Version 9.1, follow the links in the “Related concepts” section.

Related concepts:

- “Continue a recover operation that ended during the rollforward phase” on page 59
- “Rebuild database function provides new restore options” on page 60
- “Redirected restore operation using an automatically generated script” on page 59
- “Tivoli Storage Manager support added (Windows x64)” on page 61

Continue a recover operation that ended during the rollforward phase

If you issue the RECOVER DATABASE command following an incomplete recover operation that ended during the rollforward phase, the recover utility will now attempt to continue the previous recover operation, without redoing the restore phase. If you want to force the recover utility to redo the restore phase, you need to issue the RECOVER DATABASE command with the RESTART option to force the recover utility to ignore any prior recover operation that failed to complete.

A RECOVER will not complete if, for example, it was interrupted by the user (if Ctrl+C is pressed or the FORCE APPLICATION command is used) or if DB2 was unable to retrieve log files.

Related concepts:

- “Backup, logging, and recovery enhancements summary” on page 59

Related tasks:

- “Using recover” in *Data Recovery and High Availability Guide and Reference*

Related reference:

- “FORCE APPLICATION command” in *Command Reference*
- “RECOVER DATABASE command” in *Command Reference*

Redirected restore operation using an automatically generated script

You can now perform a redirected restore operation by generating a redirected restore script from an existing backup image, modifying any of the paths or container sizes in the script, then running the script to recreate the database with the new set of containers.

A redirected restore is a restore operation in which you redefine table space containers. Container redirection provides considerable flexibility for managing table space containers. For example, even though adding containers to SMS table spaces is not supported, you could accomplish this by specifying an additional container when invoking a redirected restore operation.

Related concepts:

- “Backup, logging, and recovery enhancements summary” on page 59
- “Redefine table space containers by restoring a database using an automatically generated script” in *Data Recovery and High Availability Guide and Reference*

Related tasks:

- “Performing a redirected restore using an automatically generated script” in *Data Recovery and High Availability Guide and Reference*

Rebuild database function provides new restore options

The ability to rebuild a database is new functionality that involves restoring a database or a subset of its table spaces using a set of restore operations. This functionality makes DB2 recovery more robust and versatile and provides you with a more complete recovery solution.

Because you can rebuild a database from table space backup images, it means that you no longer have to take as many full database backups. As databases grow in size, opportunities for taking a full database backup are becoming limited. With table space backup as an alternative, you no longer need to take full database backups as frequently. Instead, you can take more frequent table space backups and plan to use them, along with log files, in case of a disaster.

In a recovery situation, if you need to bring a subset of table space online faster than others, you can use rebuild to accomplish this. The ability to bring only a subset of table spaces online is especially useful in a test and production environment.

Rebuilding a database involves a series of potentially many restore operations. A rebuild operation can use a database image, or table space images, or both. It can use full backups, or incremental backups, or both. Rebuilding allows you to build a database that is connectable and that contains the subset of table spaces that you need to have online, while keeping table spaces that can be recovered at a later time offline.

Related concepts:

- “Backup, logging, and recovery enhancements summary” on page 59
- “Choosing a target image for database rebuild” in *Data Recovery and High Availability Guide and Reference*
- “Database rebuild” in *Data Recovery and High Availability Guide and Reference*

Related tasks:

- “Rebuilding a database using selected table space images” in *Data Recovery and High Availability Guide and Reference*
- “Rebuilding selected table spaces” in *Data Recovery and High Availability Guide and Reference*

Tivoli Storage Manager support added (Windows x64)

Tivoli® Storage Manager (TSM) provides storage management and data access services to protect your organization's data from failures and other errors. TSM supports various communication methods, provides administrative facilities to manage the backup and storage of files, and provides facilities for scheduling backups.

TSM support has been added for all Windows x64 systems. The minimum required level of TSM client API is Version 5.3.2.

Related concepts:

- "Backup, logging, and recovery enhancements summary" on page 59

Related reference:

- "Tivoli Storage Manager" in *Data Recovery and High Availability Guide and Reference*

Chapter 7. Client and connectivity enhancements

Client and connectivity enhancements summary

DB2 Version 9.1 client and connectivity enhancements include new features such as Internet Protocol Version 6 (IPv6) communication protocol support, new support for running ODBC and CLI applications without a DB2 Client, and connection timeout support. These enhancements and others ensure that you have flexible and effective methods of accessing data from client systems and applications.

For client and connectivity enhancements introduced in DB2 Version 9.1, follow the links in the “Related concepts” section.

Related concepts:

- “Connection timeout support for database applications added” on page 63
- “A single client (DB2 Client) is provided for application development and administration” on page 64
- “DB2 Runtime Client enhancements (Windows)” on page 64
- “Internet Protocol Version 6 (IPv6) support added” on page 65
- “ODBC and CLI applications can now run without a DB2 client” on page 65

Connection timeout support for database applications added

You can now set a connection timeout value for DB2 database connections. A connection timeout value is a limit to the amount of time that an application should wait for a connection. Setting a connection timeout value is useful in case the database server is inaccessible. In this situation, it can take a long time for connection requests to fail and return.

For .NET, CLI, ODBC, and OLE DB applications, you need to use the `ConnectTimeout` configuration keyword to enable connection timeout. For other types of application interfaces, such as the command line processor (CLP), you need to set the `DB2TCP_CLIENT_CONTIMEOUT` registry variable. For more information, follow the related links.

Related concepts:

- “Client and connectivity enhancements summary” on page 63

Related reference:

- “ConnectTimeout CLI/ODBC configuration keyword” in *Call Level Interface Guide and Reference, Volume 1*
- “Interaction between client connection timeout and client reroute” in *Administration Guide: Implementation*
- “Communications variables” in *Performance Guide*

A single client (DB2 Client) is provided for application development and administration

Prior to Version 9, there were two similar clients, the DB2 Application Development Client and the DB2 Administration Client. The Version 9 DB2 Client merges these two clients.

On the Windows operating systems, the following installation enhancements supporting the DB2 Client are included:

- You can install the DB2 Client from a server installation image.
- If you have both a DB2 Client and DB2 server installed on the same computer, you need to download and apply only a single fix pack image to upgrade both the server and client. If you have only a DB2 Client installed, you will still need to download and apply a client-specific fix pack.

Other enhancements affecting the DB2 Client are as follows:

- You can install multiple copies of a client on the same computer. These copies can be the same or different versions: for example, V8.2, V9.1.100, and V9.1.300.
- When migrating a client from Version 8, you now have the choice of replacing the existing client or adding the client while keeping the existing version.

Related concepts:

- “Types of clients - DB2 Runtime Client and DB2 Client” in *Quick Beginnings for DB2 Clients*
- “Client and connectivity enhancements summary” on page 63
- “Coexistence of multiple DB2 versions and fix packs enhancements (Linux and UNIX)” on page 69
- “Coexistence of multiple DB2 versions and fix packs now supported (Windows)” on page 70

DB2 Runtime Client enhancements (Windows)

Prior to Version 9, there were two similar runtime clients on the Windows operating systems, the DB2 Run-Time Client and the DB2 Run-Time Client Lite. The Version 9 Runtime Client evolved from the Version 8 Run-Time Client Lite and is now the only runtime client available for the Windows operating systems. On the Linux and UNIX platforms, the Version 9 Runtime Client has changed little from the Version 8 DB2 Run-Time Client.

The Windows version of the DB2 Runtime Client includes the following enhancements:

- It is licensed so that you can freely redistribute it with applications that you sell to other companies.
- It is available in a 64-bit version.
- It can coexist with other DB2 products on the same computer.

Other enhancements affecting the DB2 Runtime Client are as follows:

- You can install multiple copies of a client on the same computer. These copies can be the same or different versions, for example, V8.2, V9.1.100 and V9.1.300.

- Version 9 introduces a new application driver for ODBC or CLI applications. If you used a runtime client in the past, you might prefer using this driver, as it uses an even smaller application deployment footprint.

Related concepts:

- “Types of clients - DB2 Runtime Client and DB2 Client” in *Quick Beginnings for DB2 Clients*
- “Client and connectivity enhancements summary” on page 63
- “Coexistence of multiple DB2 versions and fix packs now supported (Windows)” on page 70
- “ODBC and CLI applications can now run without a DB2 client” on page 65

Internet Protocol Version 6 (IPv6) support added

DB2 Version 9.1 adds support for Internet Protocol Version 6 (IPv6). This means that DB2 Version 9.1 can connect to servers using IPv4 or IPv6 addresses. Some commands have been added or enhanced to provide IPv6 support. For example, the existing CATALOG TCPIP NODE command has been supplemented with the additional commands, CATALOG TCPIP4 NODE and CATALOG TCPIP6 NODE, to enable you to request a particular type of connection.

TCP/IP nodes migrated from Version 8 can return both IPv4 and IPv6 addresses. For details about the migration, see “Migration essentials for DB2 clients” in *Migration Guide*.

Related concepts:

- “Client and connectivity enhancements summary” on page 63

Related reference:

- “CATALOG TCPIP/TCPIP4/TCPIP6 NODE command” in *Command Reference*
- “REGISTER command” in *Command Reference*

ODBC and CLI applications can now run without a DB2 client

In addition to the DB2 CLI and ODBC driver that is installed as part of the DB2 Client and the DB2 Runtime Client, there is a new driver, called the IBM DB2 Driver for ODBC and CLI that can be installed without a DB2 client. The IBM DB2 Driver for ODBC and CLI provides runtime support for ODBC and CLI applications, and it also provides connectivity for those applications.

Being able to install the IBM DB2 Driver for ODBC and CLI without a DB2 client makes deploying database applications easier:

- You can include the driver in your database application installation package
- Distribution size, installation footprint, and memory footprint are reduced

There are various install options available for the IBM DB2 Driver for ODBC and CLI:

- You can install the driver on a machine that already has a DB2 client installed
- You can have multiple installations of this driver on a single machine

Related concepts:

- “Client and connectivity enhancements summary” on page 63

- “DB2 CLI and ODBC drivers” in *Call Level Interface Guide and Reference, Volume 1*
- “IBM DB2 Driver for ODBC and CLI overview” in *Call Level Interface Guide and Reference, Volume 1*
- “Introduction to DB2 CLI and ODBC” in *Call Level Interface Guide and Reference, Volume 1*

Chapter 8. Federation enhancements

Federation enhancements summary

A federated system is a special type of distributed database management system (DBMS). With federation, you can send distributed requests to multiple data sources using a single SQL statement.

For information about federation enhancements introduced in DB2 Version 9.1, follow the links in the “Related concepts” section.

Related concepts:

- “User mapping retrieval from an external repository is supported” on page 67
- “Statement level isolation for nicknames” on page 68
- “Two-phase commit for multivendor data sources” on page 68

User mapping retrieval from an external repository is supported

A federated server uses a user mapping to connect to many data sources. By default, the user mappings are stored in a DB2 database, which has certain limitations. You can now develop your own plug-in to retrieve user mappings from other external repositories. A sample plug-in is provided that allows you to retrieve the user mappings from a Lightweight Directory Access Protocol (LDAP) server. User mappings that are stored in an LDAP server offer the following benefits:

- Lower maintenance. Your user mappings are stored in a centralized external repository that can be used by multiple federated servers. Therefore, you need to update your user mappings in only one location, instead of on each federated server.
- Increased security. You have greater control over the security measures that are used to store your user mappings. If you store your user mappings in a DB2 database, the remote passwords are stored in the SYSUSEROPTIONS table. The passwords are weakly encrypted and use a fixed encryption key. However, if you use an external repository such as an LDAP server, you can implement your own encryption algorithm and secret key.

Related concepts:

- “Federation enhancements summary” on page 67
- “Advantages of using an external repository” in *WebSphere Data Source Configuration Guide*
- “Developing a plugin for retrieving user mappings from an external repository Overview” in *WebSphere Data Source Configuration Guide*
- “LDAP sample plugin” in *WebSphere Data Source Configuration Guide*
- “Overview of the user mapping plugin for external repositories” in *WebSphere Data Source Configuration Guide*

Statement level isolation for nicknames

For DB2 data sources and Microsoft SQL Server data sources, you can use statement level isolation. You must use the WITH isolation clause in statements that use nicknames if you want to use statement level isolation.

For all other uses of nicknames to access relational data sources, the federated server maps its current isolation level to a corresponding isolation level at the data source, at each connection to the data source. After a connection is made to a data source, the isolation level for duration of the connection cannot be changed.

Related concepts:

- “Federation enhancements summary” on page 67
- “Isolation levels and performance” in *Performance Guide*

Two-phase commit for multivendor data sources

Federated two-phase commit is now available for the supported data sources of the WebSphere Information Integrator federated server. Two-phase commit combines updates to multiple sources in one transaction so that all sources involved are updated or none are updated. This strategy ensures that the sources remain synchronized.

DB2 Version 9.1 and DB2 Connect users may be interested in two-phase commit capabilities for multivendor data sources offered by the WebSphere Information Integrator federated server. The data sources supported by the WebSphere Information Integrator federated server include:

- DB2 family data sources through the Distributed Relational Database Architecture (DRDA[®]) protocol:
 - DB2 UDB for Linux, UNIX, and Windows, Version 8.1 or later
 - DB2 UDB for z/OS, Version 7.1 or later
 - DB2 UDB for iSeries, Version 5.3 or later
- Informix IDS Version 7.31 or later, Version 9.40 or later, Version 10.0 or later
- Informix XPS Version 8.40 or later
- Microsoft SQL Server 2000 and Microsoft SQL Server 2005 for a federated server only on Windows
- Oracle, Version 8.1.7 or later, with the XA library
- Sybase Adaptive Server Enterprise, Version 12 or later, with the XA library for a federated server only on Windows

Related concepts:

- “Federation enhancements summary” on page 67
- “Two-phase commit” in *Administration Guide: Planning*

Chapter 9. Installation, migration, and fix pack enhancements

Installation, migration, and fix pack enhancements summary

With new features and enhancements such as coexistence of multiple DB2 versions and fix packs on Windows, manual installation enhancements, and support for concurrent copies of DB2 database systems on Linux and UNIX, DB2 Version 9.1 allows your IT staff to spend more time supporting business goals and less time installing and configuring database systems.

For information about installation, migration, and fix pack enhancements introduced in DB2 Version 9.1, follow the links in the “Related concepts” section.

Related concepts:

- “Coexistence of multiple DB2 versions and fix packs enhancements (Linux and UNIX)” on page 69
- “Coexistence of multiple DB2 versions and fix packs now supported (Windows)” on page 70
- “Installation of DB2 products without using an Administrator user ID now supported (Windows)” on page 72
- “Manual installation enhancements (Linux and UNIX)” on page 73
- “Migration resources for DB2 database systems added” on page 73
- “New DB2 product uninstall features using the db2unins command (Windows)” on page 74
- “Response file enhancements” on page 74
- “Update capabilities added to the DB2 Information Center” on page 75

Coexistence of multiple DB2 versions and fix packs enhancements (Linux and UNIX)

You can install multiple copies of DB2 database systems on Linux or UNIX operating systems without the need for alternate FixPaks. DB2 fix packs for the Linux or UNIX operating systems can now be used to install a new DB2 copy at the corresponding fix pack level. Other key benefits of this feature include:

- **Install anywhere:** You can install DB2 database systems using any valid path that you choose.
- **Install any number of times:** You can install two or more copies of the same database system on one computer. The code levels can be different as well.
- **Service each copy independently:** You can update one copy without affecting any of the other copies.

You can use this feature in a number of ways:

- Database administrators can deploy one version of a DB2 database system for production, and a newer version to test new fix packs. Then, the production instance can be rolled over to the new installation path on an instance-by-instance basis, without having to apply the fix pack to the earlier version (although in-place fix packs are available). You can simply uninstall the earlier version when you no longer need it.

- Database administrators can have independent copies of DB2 database systems for different purposes. This independence allows different databases on the same computer to run at different fix pack levels. For example, one group, such as human resources, can apply fixes without affecting other groups, such as finance.
- If you are an independent software vendor, you can embed your own DB2 copy into your product.

On supported Linux and UNIX operating systems, a new command, **db2ls**, provides information about DB2 database systems and features installed on your system. You can use this command to first list where DB2 database systems are installed and which level of the DB2 database system is installed. After identifying the installation paths where DB2 database systems are installed, you can use this command to list all, or specific, products and features installed in a particular installation path.

In Version 9, the **db2ls** command is the only method for querying a DB2 product. You can no longer query DB2 products using Linux or UNIX operating system native utilities such as `pkgadd`, `rpm`, `SMIT`, or `swinstall`. You must change any scripts containing a native installation utility that you use to interface with and query DB2 installations.

Related concepts:

- “Installation, migration, and fix pack enhancements summary” on page 69
- “Multiple DB2 copies on the same computer (Windows)” in *Administration Guide: Implementation*

Related tasks:

- “Installing DB2 servers (Linux and UNIX)” in *Quick Beginnings for DB2 Servers*

Related reference:

- “db2ls - List installed DB2 products and features command” in *Command Reference*
- “Multiple DB2 copies roadmap” in *Administration Guide: Implementation*

Coexistence of multiple DB2 versions and fix packs now supported (Windows)

DB2 Version 9 on the Windows operating system introduces the ability to install multiple DB2 server and client copies on the same system. Each DB2 installation copy can either be at the same level or at a different level of DB2.

When installing a DB2 product on a machine that already has a DB2 copy installed, you can install a new DB2 copy or update existing DB2 copies.

DB2 licensing:

In previous DB2 versions on the Windows operating system, a DB2 database system was licensed as an entire physical machine. DB2 Version 9 is licensed using an installation path model, based on the products installed in a given install directory.

Default DB2 copy:

As there can be multiple DB2 copies on the same system, you must specify a DB2 copy as the default DB2 installation for remote connections. You can update the default DB2 installation with the Default DB2 Selection wizard.

If you have DB2 UDB Version 8 and DB2 Version 9 coexisting on your system, DB2 UDB Version 8 is always the default DB2 copy on that system.

New DB2_COPY_NAME registry variable:

The DB2_COPY_NAME variable stores the name of the copy of DB2 currently in use. To switch to a different DB2 copy when multiple DB2 copies are installed, you cannot use the DB2_COPY_NAME variable, but you can choose one of the following methods:

- you can use the DB2 command window from the Start → Programs → IBM DB2 → *<DB2 Copy Name>* → Command Line Tools → DB2 Command Window: the command window is already set up with the correct environment variables for the particular DB2 copy chosen.
- you can run the db2envar.bat file from a command window using the fully qualified path (*<DB2 Copy install directory>\bin\db2envar.bat*) for the DB2 copy that you want to use.

DB2 Version 9 and DB2 UDB Version 8 coexistence restrictions:

Coexistence of DB2 Version 9 and DB2 UDB Version 8 copies has some restrictions:

- A DB2 UDB Version 8 DAS cannot administer a DB2 Version 9 installation. A DB2 Version 9 DAS can administer both DB2 Version 9 and DB2 UDB Version 8 instances.
- The DB2 UDB Version 8 copy is assigned as the default DB2 copy on the physical machine. To select a new copy, you must migrate the Version 8 copy to Version 9, or uninstall the Version 8 copy and use the Default DB2 Selection wizard.
- 32-bit and 64-bit versions of the DB2 database system cannot coexist on the Windows for AMD64 and Intel EM64T operating system. If you have a 32-bit DB2 database system and install a 64-bit DB2 database system, the DB2 32-bit database system is automatically migrated. If you want to replace multiple 32-bit DB2 database systems with a single 64-bit database system, you must first migrate all of the 32-bit copies to a single 32-bit copy. Next, you must uninstall all of the 32-bit copies except for the one that contains all of the 32-bit instances. Finally, you can install a DB2 Version 9 64-bit database system. During the installation, all of the 32-bit copies are automatically migrated.
- DB2 Net Search Extender does not support multiple copies of different versions or releases on the same physical machine.

Client connectivity:

You can only use one DB2 copy in the same process at the same time.

Fix packs and migration:

If you are using a refresh image to update your DB2 database system, you must specify whether you want to update an existing DB2 copy or install a new DB2 copy. You cannot update more than one DB2 copy at the same time. To update a DB2 copy installed on your system, you must rerun the installation.

The DB2 installation process gives you the option to migrate from DB2 UDB Version 8 (in the same path) or to install a new DB2 Version 9 copy without modifying the existing DB2 UDB Version 8 installation.

Application development impact:

Having multiple DB2 copies on one system has a significant impact on how you develop your applications. The biggest change is how you make your database applications work with a specific DB2 instance that is not the default DB2 instance.

Related concepts:

- “Environment variables and the profile registry” in *Administration Guide: Implementation*
- “Multiple DB2 copies on the same computer (Windows)” in *Administration Guide: Implementation*
- “Installation, migration, and fix pack enhancements summary” on page 69

Related tasks:

- “Installing DB2 servers (Linux and UNIX)” in *Quick Beginnings for DB2 Servers*
- “Installing DB2 servers (Windows)” in *Quick Beginnings for DB2 Servers*
- “Using the Default DB2 Selection wizard (Windows)” in *Quick Beginnings for DB2 Servers*
- “DB2 Version 9 co-existence and migration with DB2 UDB Version 8 (Windows)” in *Quick Beginnings for DB2 Servers*

Related reference:

- “System environment variables” in *Performance Guide*
- “Response file keywords (Windows and UNIX)” in *Installation and Configuration Supplement*
- “Multiple DB2 copies roadmap” in *Administration Guide: Implementation*

Installation of DB2 products without using an Administrator user ID now supported (Windows)

If you do not want to use an Administrator user ID when installing a DB2 product on a Windows operating system, you now have another option. You can use the Windows operating system elevated privileges feature to perform the installation with a Windows Power User ID or a Restricted User ID.

Related concepts:

- “Installation, migration, and fix pack enhancements summary” on page 69

Related reference:

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

Manual installation enhancements (Linux and UNIX)

The **db2_install** command has been simplified and enhanced to support the ability to install multiple DB2 copies on the same computer and DB2 products that have installation images that span multiple CDs. Additional options are provided to help you complete the following tasks:

- Install non-English language support using the -c and -L options
- Capture troubleshooting information with the trace facility using the -t option

The **db2_deinstall** command, previously only provided on the DB2 product CD, is now installed as part of the base installation image. It is located at *DB2DIR/install*, where *DB2DIR* is the location where the current version of the DB2 database product is installed. The **db2_deinstall** command is not provided on the product CDs any more.

Related concepts:

- “Installation, migration, and fix pack enhancements summary” on page 69
- “Coexistence of multiple DB2 versions and fix packs enhancements (Linux and UNIX)” on page 69

Related tasks:

- “Listing DB2 products installed on your system (Linux and UNIX)” in *Quick Beginnings for DB2 Servers*
- “Installing a DB2 product manually” in *Installation and Configuration Supplement*

Related reference:

- “db2_deinstall - Uninstall DB2 products or features command” in *Command Reference*
- “db2_install - Install DB2 product command” in *Command Reference*

Migration resources for DB2 database systems added

The most important resources available to assist you with the migration of DB2 database systems are:

- The topics under “Migration to DB2 Version 9” in the DB2 Information Center describe the entire migration process and the migration components. The migration components are DB2 database systems, DB2 clients, and database applications and routines.
- The DB2 Migration Guide in PDF format provides information about the entire migration process and components, including planning your migration. This guide is available for download from the DB2 Version 9 manuals Web page at www.ibm.com/software/data/db2/udb/support/manualsv9.html.
- The DB2 database system migration portal at <http://www.ibm.com/support/docview.wss?rs=73&uid=swg21200005> provides you with one place to access all additional resources and up-to-date information about the migration process as they become available. These resources include white papers and sample scripts for migration.

Related concepts:

- “Installation, migration, and fix pack enhancements summary” on page 69
- “Migration support changes for DB2 database systems” on page 110

New DB2 product uninstall features using the `db2unins` command (Windows)

In DB2 Version 9.1, you can use the `db2unins` command to:

- Conveniently uninstall multiple DB2 products within a DB2 instance in a single operation. Previously, you could uninstall DB2 products only one at a time. In DB2 Version 9.1, you can specify the `db2unins` command with the `-p` option.
- Silently uninstall DB2 products within a DB2 instance, as defined in the uninstall response file. Previously, you had to use the Add/Remove Program function in the Windows Control Panel to uninstall a DB2 product. In DB2 Version 9.1, you can specify the `db2unins` command with the `-u` option.
- Automatically force the uninstall of all DB2 products on the system. Use this feature for systems that are so badly damaged that the only option is to manually clean them. Previously, you had to manually perform this task. In DB2 Version 9.1, you can specify the `db2unins` command with the `-f` option.

There are other options for the `db2unins` command, such as, the `-d` option which allows you to view the complete list of DB2 products that are installed in the current DB2 copy, and the `-l` option which allows you to generate a log file for the `db2unins` command.

Related concepts:

- “Installation, migration, and fix pack enhancements summary” on page 69

Related reference:

- “`db2unins` - Uninstall DB2 database product command” in *Command Reference*

Response file enhancements

A response file is an ASCII text file that contains setup and configuration information. Unlike installing with the DB2 Setup wizard, installing with a DB2 response file lets you install DB2 products or features without user interaction.

In Version 9.1, new response file keywords are added to support new function and enhancements. Ready-to-use sample response files with default entries are included on the DB2 CD in `db2/platform/samples` (where *platform* refers to the hardware platform).

Related concepts:

- “Response file installation basics” in *Installation and Configuration Supplement*
- “Installation, migration, and fix pack enhancements summary” on page 69

Related reference:

- “Available sample response files” in *Installation and Configuration Supplement*
- “Response file keywords (Windows and UNIX)” in *Installation and Configuration Supplement*

Update capabilities added to the DB2 Information Center

If you choose to install the DB2 Version 9.1 Information Center locally, rather than use the online version available at <http://publib.boulder.ibm.com/infocenter/db2help>, you can take advantage of the update capability built into the Information Center itself. With this feature, you can search for and download updates for your Information Center simply by clicking the Update button in the Information Center window.

Related concepts:

- “Installation, migration, and fix pack enhancements summary” on page 69

Related tasks:

- “Updating the DB2 Information Center installed on your computer or intranet server” on page 161

Chapter 10. Manageability enhancements

Manageability enhancements summary

Manageability is improved in DB2 Version 9.1 with several new enhancements including default enablement of autonomic computing features, ALTER TABLE statement enhancements, automatic storage support for multi-partition databases, and support for copying schemas between databases. Other improvements include enhanced access to DB2 administration commands through SQL, automatic table and index reorganization enhancements, and load support for partitioned tables. These features and others will allow you to spend less time managing your databases and more time managing your business.

For manageability enhancements introduced in DB2 Version 9.1, follow the links in the “Related concepts” section.

Related concepts:

- “Adaptive, self-tuning memory allocation” on page 77
- “ALTER TABLE statement enhancements” on page 78
- “Automated evaluation of object maintenance policies by the DB2 UDB for z/OS health monitor” on page 78
- “Automatic storage enhancements” on page 79
- “Automatic table and index reorganization enhancements” on page 80
- “Buffer pool identifiers displayed in database memory output” on page 81
- “Copy schema support” on page 81
- “Enhanced access to DB2 administration commands through SQL” on page 81
- “EXEC SQL utility support for loading data into DB2 UDB for z/OS tables” on page 82
- “Fast communications manager (FCM) enhancements” on page 82
- “First Steps interface enhancements” on page 83
- “Identifier length limits increased to 128 bytes for additional objects and names” on page 84
- “Load support for partitioned tables added” on page 84
- “New SQL Monitor Area (SQLMA) for monitoring information flow in partitioned database environments” on page 85

Adaptive, self-tuning memory allocation

Adaptive, self-tuning memory simplifies the task of memory configuration by automatically setting values for memory configuration parameters and sizing buffer pools. When enabled, the memory tuner dynamically distributes available memory resources between several memory consumers including sort, package cache, and lock list, as well as buffer pools.

On Windows and AIX platforms, the self-tuning memory feature can also determine the overall database memory requirements and dynamically tune total database shared memory usage. This allows the database manager to consume

more physical memory if required by the workload, and free up that memory to the operating system when database memory requirements are low.

Use the `database_memory` configuration parameter to set a limit to the amount of memory that each database will use. See the related links for more information.

In addition to simplifying the task of memory configuration, this new adaptive self-tuning memory feature improves performance by providing a superior configuration that is dynamic and responsive to significant changes in workload characteristics.

Related concepts:

- “Self tuning memory” in *Performance Guide*
- “Self tuning memory manager roadmap” in *Performance Guide*
- “Automatic self-tuning memory enabled by default during database creation” on page 113
- “Configuration parameters default value changes” on page 117
- “`database_memory` configuration parameter change” on page 119
- “Manageability enhancements summary” on page 77

ALTER TABLE statement enhancements

You can now use the ALTER TABLE statement to:

- Drop a column, using the new DROP COLUMN clause
- Change a column type to a larger variety of data types, using the ALTER COLUMN SET DATA TYPE clause
- Change the nullability attribute of a column, using the SET NOT NULL or the DROP NOT NULL clause

When changing these table attributes using SQL, it is no longer necessary to drop the table and then recreate it, a time consuming process that can be complex when object dependencies exist.

A table alteration that affects the row format of the data is known as a *REORG-recommended operation*, and requires table reorganization before most subsequent operations on the table are allowed.

Related concepts:

- “Manageability enhancements summary” on page 77

Related reference:

- “ALTER TABLE statement” in *SQL Reference, Volume 2*
- “REORG INDEXES/TABLE command” in *Command Reference*

Automated evaluation of object maintenance policies by the DB2 UDB for z/OS health monitor

This topic refers collectively to the following products as DB2 UDB for z/OS:

- DB2 UDB for z/OS Version 8
- DB2 UDB for OS/390 and z/OS Version 7

On the z/OS system, a DB2 UDB for z/OS health monitor is started as a task for each DB2 subsystem to be monitored. The Control Center's Create and Change Object Maintenance Policy wizards now allow you to automate the evaluation of object maintenance policies by the DB2 UDB for z/OS health monitor.

The DB2 UDB for z/OS health monitor triggers policy evaluations at scheduled times and intervals, as defined in the policy. During each policy evaluation, the criteria for recommending maintenance is checked against the thresholds set in the object maintenance policy to determine the need for object maintenance, that is, whether COPY, REORG, RUNSTATS, STOSPACE, ALTER TABLESPACE, or ALTER INDEX are required, and to identify restricted states, such as CHKP, on table space and index objects. When objects are identified to be in the alert state during policy evaluation, the policy health alert contacts are notified at their e-mail addresses or pager numbers. The list of health alert contacts for each DB2 subsystem is defined in and managed from the Control Center.

To enable the Control Center to support these new features and to activate the DB2 UDB for z/OS health monitor and modify the object maintenance tables, you need to apply APAR PK20053 for:

- z/OS Enablement (FMID JDB881D), if connecting to a DB2 UDB for z/OS Version 8 server
- OS/390 Enablement (FMID JDB771D), if connecting to a DB2 UDB for OS/390 and z/OS Version 7 server

Related concepts:

- "Control Center overview" in *Administration Guide: Implementation*
- "Introduction to the health monitor" in *System Monitor Guide and Reference*
- "Manageability enhancements summary" on page 77
- "DB2 UDB for z/OS health monitor overview" in *Administration Guide: Implementation*

Automatic storage enhancements

In Version 9.1, automatic storage support is added for multi-partition databases. In addition, the Control Center has been enhanced to let you create databases that use automatic storage, and to enable you to add storage paths to existing databases.

Support added for multi-partition databases:

If you have the Enterprise Server Edition with the Database Partitioning Feature (DPF), you may now use multi-partition databases that support automatic storage. An automatic storage database is one in which table spaces can be created and whose container and space management characteristics are completely determined by the DB2 database manager. Databases that are enabled for automatic storage have a set of one or more storage paths associated with them. A table space can be defined as "managed by automatic storage" and have its containers assigned and allocated by the DB2 database manager based on those storage paths.

Also, you can specify paths in addition to drive letters when creating a database in a Windows operating environment.

Control Center enhancements:

You can create an automatic storage database or add a storage path to an existing automatic storage database using the Control Center.

To create an automatic storage database using the Control Center:

1. Expand the object tree until you find the **Databases** folder.
2. Right-click the **Databases** folder, and select **Create** → **Standard** or **Create** → **With Automatic Maintenance** from the pop-up menu.
3. Follow the steps to complete this task.

To add a storage path to an existing database using the Control Center:

1. Open the Add Storage window: Expand the object tree until you see the **Table Spaces** folder of the database to which you want to add a storage path. Right-click the **Table Spaces** folder and select **Manage Storage** → **Add Automatic Storage** from the pop-up menu. The Add Storage window opens.
2. Click **Add**. The Add Storage Path window opens.
3. Specify the storage path.

Related concepts:

- “Automatic storage databases” in *Administration Guide: Implementation*
- “Database Partitioning Feature (DPF)” in *Administration Guide: Planning*
- “Configuration parameters default value changes” on page 117
- “Manageability enhancements summary” on page 77

Automatic table and index reorganization enhancements

New policy options allow you to

- Specify the use of a system temporary table space of an appropriate page size for an offline table reorganization. This table space is used to store a temporary copy of the table being automatically reorganized.
- Apply a size limit to tables considered for offline reorganization. A new option is added to the scope specification for table reorganization to set a size limit for offline reorganization operations. Any tables larger than the size limit will not have offline reorganizations applied automatically. Online reorganizations will still be applied automatically. If a table spans multiple database partitions, the size limit applies to the size of the table portion contained on each database partition, calculated as an average.
- Specify that automatic index reorganizations should be run online and in the online maintenance window.
- Specify whether to keep or rebuild compression dictionaries during reorganization.

Related concepts:

- “Configuration parameters default value changes” on page 117
- “Manageability enhancements summary” on page 77

Related tasks:

- “Enabling automatic table and index reorganization” in *Performance Guide*

Buffer pool identifiers displayed in database memory output

Version 9.1 provides users with the capability to distinguish multiple buffer pools in the database memory output of the **db2mtrk** command. Buffer pool identifiers are now displayed in parentheses next to the buffer pools (including system buffer pools).

Related concepts:

- “Buffer pool management” in *Performance Guide*

Related reference:

- “db2mtrk - Memory tracker command” in *Command Reference*

Copy schema support

The db2move utility and the ADMIN_COPY_SCHEMA procedure allow you to conveniently make copies of a database schema and its associated database objects. Once a model schema is established, you can use it as a template for creating new versions.

You can use the ADMIN_COPY_SCHEMA procedure to copy a single schema within the same database.

You can also use the db2move utility with the **-co** option and the COPY action to copy a single schema or multiple schemas from a source database to a target database. If you use the db2move utility, the source and the target database must be different. If you want to copy a schema within the same database, use the ADMIN_COPY_SCHEMA procedure.

Related concepts:

- “Manageability enhancements summary” on page 77

Related tasks:

- “Copying a schema” in *Administration Guide: Implementation*

Related reference:

- “db2move - Database movement tool command” in *Command Reference*
- “ADMIN_COPY_SCHEMA procedure – Copy a specific schema and its objects” in *Administrative SQL Routines and Views*
- “ADMIN_DROP_SCHEMA procedure – Drop a specific schema and its objects” in *Administrative SQL Routines and Views*

Enhanced access to DB2 administration commands through SQL

SQL administrative routines were introduced in DB2 UDB for Linux, UNIX, and Windows Version 8 and have been expanded in DB2 Version 9 to include more administrative tasks. New administrative views have also been added in DB2 Version 9.

The SQL administrative routines and views provide a primary, easy-to-use programmatic interface to administer DB2 through SQL. They encompass a collection of built-in views, table functions, procedures, and scalar functions for

performing a variety of DB2 administrative tasks such as: reorganizing a table, capturing and retrieving monitor data, or retrieving the application ID of the current connection.

These routines and views can be invoked from an SQL-based application, a DB2 command line or a command script.

In addition to the new administrative views, routines, and procedures, DB2 Version 9 includes:

- Expanded support for monitoring your database
- Expanded support for DB2 command execution through the ADMIN_CMD procedure
- Multi-partition database support for administrative views and routines

In order to provide expanded support for the existing administrative routines, some of the DB2 UDB for Linux, UNIX, and Windows Version 8 routines have been replaced with new, more comprehensive routines or views in DB2 Version 9.

For a list of all supported administrative SQL routines and view, including the new ones, see "Supported administrative SQL routines and views" in *Administrative SQL Routines and Views*.

Related concepts:

- "Manageability enhancements summary" on page 77
- "Some SQL administrative routines have been replaced" on page 143

Related reference:

- "Supported administrative SQL routines and views" in *Administrative SQL Routines and Views*

EXEC SQL utility support for loading data into DB2 UDB for z/OS tables

This topic refers collectively to the following products as DB2 UDB for z/OS:

- DB2 UDB for z/OS Version 8
- DB2 UDB for OS/390 and z/OS Version 7

The Control Center's Load Table and Load on Table Space notebooks now support the INCURSOR option of the LOAD utility. The Control Center uses the EXEC SQL utility control statement to declare a cursor, and the result table from the declared cursor is used as input to the LOAD utility.

Related concepts:

- "Manageability enhancements summary" on page 77

Related reference:

- "LOAD command" in *Command Reference*

Fast communications manager (FCM) enhancements

The re-architecture of the DB2 communications subsystem has resulted in several enhancements involving Database Partitioning Feature (DPF) servers:

- Separate sender and receiver communications daemons are now used to improve communication speed.
- NUMA partitioning of FCM resources by logical partition is now supported on any platform that offers per-partition affinity.
- The number of FCM buffers and FCM channels can be dynamically changed during execution, eliminating the need for tuning. There are several other dynamic configuration improvements that automatically monitor resource usage without your involvement.

These enhancements affect several configuration parameters and monitor elements.

- The `fcm_num_buffers` parameter has a new `AUTOMATIC` setting that allows the DB2 database system to attempt to tune the configured parameter value if the database is not making full use of resources. This setting is turned on by default.
- The new `fcm_num_channels` configuration parameter specifies the number of FCM channels. It replaces the deprecated `fcm_num_rqb`, `fcm_num_anchors`, and `fcm_num_connect` parameters. This parameter is set to `AUTOMATIC` by default.
- Two new monitor elements replace a number of deprecated ones.
 - The new `ch_free` monitor element indicates the number of inter-node communication channels that are currently free. It replaces the deprecated `ma_free`, `ce_free`, and `rb_free` monitor elements.
 - The new `ch_free_bottom` monitor element indicates the lowest number of free inter-node communication channels reached during processing. It replaces the deprecated `ma_free_bottom`, `ce_free_bottom`, and `rb_free_bottom` monitor elements.

The deprecated monitor elements no longer appear in the dynamic snapshot streams in FCM snapshots.

Related concepts:

- “Fast communications manager (FCM) communications” in *Administration Guide: Implementation*
- “Configuration parameters default value changes” on page 117
- “Scalability enhancements summary” on page 95

Related reference:

- “`ch_free` - Channels Currently Free monitor element” in *System Monitor Guide and Reference*
- “`ch_free_bottom` - Minimum Channels Free monitor element” in *System Monitor Guide and Reference*
- “`fcm_num_buffers` - Number of FCM buffers configuration parameter” in *Performance Guide*
- “`fcm_num_channels` - Number of FCM channels configuration parameter” in *Performance Guide*

First Steps interface enhancements

The new First Steps is designed as a portal with easy access to basic database maintenance tasks instructions, application development resources, and technical resources available for your DB2 products.

The information in First Steps is tailored according to the DB2 product that you install. For example, if you install a DB2 database system, you can perform any of the following tasks with First Steps assistance:

- Create the sample database or your own database
- Back up data, load data, configure access, partition databases, and migrate database systems.
- Create an application using any of the supported development environments
- Review technical resources that provide more in-depth knowledge
- Check for product updates

First Steps provides links to the online DB2 Information Center and to relevant IBM Web sites. To use all of the First Steps features, you need an internet connection and access to the online DB2 Information Center.

Related concepts:

- “First Steps interface” in *Quick Beginnings for DB2 Servers*

Related reference:

- “db2fs - First steps command” in *Command Reference*

Identifier length limits increased to 128 bytes for additional objects and names

The list of objects or names for which you can specify identifiers that are up to 128 bytes long has been extended to include:

- Indexes
- Index specifications
- Methods
- Routine-specific names
- User-defined functions (UDFs)

Increased identifier length limits enable you to specify more descriptive names for database objects.

Increased limits might also affect existing behavior. You should modify any application that uses object names that were created outside of the application to accept longer returned names. An example of such an application is one that accepts index names from a catalog view.

Related concepts:

- “Manageability enhancements summary” on page 77

Related reference:

- “Road map to the catalog views” in *SQL Reference, Volume 1*

Load support for partitioned tables added

The load utility now supports loading data into a partitioned table. Tables using any combination of DISTRIBUTE BY, PARTITION BY, and ORGANIZE BY algorithms are supported.

The load utility inserts data records into the correct data partition without the need to use an external utility to partition the input data before loading. Maintenance of non-partitioned indexes is fully supported, and you are not required to pre-sort your table data.

You cannot load a subset of data partitions and leave the remaining data partitions in full access state. If the source data targets only a subset of data partitions, all data partitions remain in the load in progress state for the duration of the load operation.

Rows violating the range constraint are optionally inserted into the exception table.

Related concepts:

- “Load overview” in *Data Movement Utilities Guide and Reference*
- “Load considerations for partitioned tables” in *Data Movement Utilities Guide and Reference*
- “Improved large database management using table partitioning” on page 12
- “Manageability enhancements summary” on page 77

Related tasks:

- “Loading data” in *Data Movement Utilities Guide and Reference*
- “Loading data in a partitioned database environment” in *Data Movement Utilities Guide and Reference*

Related reference:

- “LOAD command” in *Command Reference*

New SQL Monitor Area (SQLMA) for monitoring information flow in partitioned database environments

There is a new SQL Monitor Area (SQLMA) for monitoring information flow in partitioned database environments. The SQLMA is an interface between a client and server that the db2GetSnapshot API uses to send database monitor snapshot requests to the database manager. The db2GetSnapshotSize API also uses the SQLMA to estimate the size of the snapshot output. Applications using these APIs will continue to work without being changed or recompiled.

Only the old SQLMA interface is supported when Version 9.1 client applications are attached to a DB2 Universal Database Version 8 instance.

Related concepts:

- “Manageability enhancements summary” on page 77

Related reference:

- “db2AddSnapshotRequest API - Add a snapshot request” in *Administrative API Reference*
- “db2GetSnapshot API - Get a snapshot of the database manager operational status” in *Administrative API Reference*
- “db2GetSnapshotSize API - Estimate the output buffer size required for the db2GetSnapshot API” in *Administrative API Reference*
- “sqlma data structure” in *Administrative API Reference*

Chapter 11. National language enhancements

National language enhancements summary

DB2 Version 9.1 enhancements for national languages include Unicode support for character-based string functions. This enhancement simplifies processing of multiple byte character data with the addition of several new scalar functions and enhancements to existing functions.

For national language enhancements introduced in DB2 Version 9.1, follow the links in the “Related concepts” section.

Related concepts:

- “Unicode support added for character-based string functions” on page 87

Unicode support added for character-based string functions

Processing of multiple byte character data has been simplified with the addition of several new scalar functions and modifications to existing functions. The new functions include:

- CHARACTER_LENGTH
- OCTET_LENGTH
- POSITION
- SUBSTRING

The modified existing functions include:

- LENGTH
- LOCATE

These functions process strings along character boundaries rather than along byte or double-byte boundaries. Each function (except OCTET_LENGTH) accepts an argument specifying the *code unit*, or string unit of the result:

- CODEUNITS16 specifies that the result is to be expressed in 16-bit UTF-16 code units
- CODEUNITS32 specifies that the result is to be expressed in 32-bit UTF-32 code units
- OCTETS specifies that the result is to be expressed in bytes

This argument is optional for the existing functions.

Related reference:

- “Character strings” in *SQL Reference, Volume 1*
- “CHARACTER_LENGTH scalar function” in *SQL Reference, Volume 1*
- “LENGTH scalar function” in *SQL Reference, Volume 1*
- “LOCATE scalar function” in *SQL Reference, Volume 1*
- “OCTET_LENGTH scalar function” in *SQL Reference, Volume 1*
- “POSITION scalar function” in *SQL Reference, Volume 1*
- “SUBSTRING scalar function” in *SQL Reference, Volume 1*

Chapter 12. Performance enhancements

Performance enhancements summary

DB2 Version 9.1 introduces a number of performance enhancements, including data row compression, adaptive self-tuning memory, enhanced query performance using statistical views, faster data loading capabilities, materialized query table (MQT) enhancements, and 64 KB memory page size support on POWER5+™ processors (on the AIX operating system). These enhancements and others ensure that you get the highest performance when accessing and updating your data.

For more information about performance enhancements introduced in DB2 Version 9.1, follow the links in the “Related concepts” section.

Related concepts:

- “64 KB page size support added for POWER5+ processors (AIX)” on page 93
- “Adaptive, self-tuning memory allocation” on page 77
- “Enhanced query performance using statistical views” on page 89
- “Faster data loading using SOURCEUSEREXIT customizable user exit” on page 90
- “Load from cursor with remote fetch” on page 90
- “Materialized query table (MQT) enhancements” on page 90
- “Online processing of the SET INTEGRITY statement” on page 91
- “Optimizer registry variables and keywords added” on page 91
- “Performance enhancements summary - native XML data store” on page 37
- “Row compression support added” on page 92

Enhanced query performance using statistical views

Statistics can now be collected for views. Views that are associated with statistics are called statistical views; they improve access plans for queries whose definition overlaps with the query definition, leading to improved query performance.

Statistical views provides the optimizer with accurate statistics for determining cardinality estimates for queries with complex sets of (possibly correlated) predicates involving one or more tables. Cardinality estimation is the process whereby the optimizer uses statistics to determine the size of partial query results after predicates are applied or aggregation is performed. The accuracy of cardinality estimates depends on the predicates and the available statistics. Statistics are available to represent the distribution of data values within a column, which can improve cardinality estimates when the data values are unevenly distributed.

Related concepts:

- “Statistical views” in *Performance Guide*
- “Performance enhancements summary” on page 89

Faster data loading using SOURCEUSEREXIT customizable user exit

The new SOURCEUSEREXIT option of the load utility allows you to feed data into the utility by using a customized script or executable. This facility supports multiple parallel user exits and is supported in partitioned database environments.

Related concepts:

- “Load overview” in *Data Movement Utilities Guide and Reference*
- “Moving data using a customized application (user exit)” in *Data Movement Utilities Guide and Reference*
- “Load in a partitioned database environment - overview” in *Data Movement Utilities Guide and Reference*
- “Performance enhancements summary” on page 89

Load from cursor with remote fetch

You can now load data from one database into another by referencing a nickname within an SQL query. Alternatively, you can use the new DATABASE option within the DECLARE CURSOR statement, or its equivalent, the new sqlu_remotefetch_entry media entry with the db2Load API.

Using the DATABASE option or the sqlu_remotefetch_entry media entry provides usability and performance benefits over using nicknames.

Related concepts:

- “Load overview” in *Data Movement Utilities Guide and Reference*
- “Moving data using the CURSOR file type” in *Data Movement Utilities Guide and Reference*
- “Performance enhancements summary” on page 89

Related tasks:

- “Loading data” in *Data Movement Utilities Guide and Reference*

Materialized query table (MQT) enhancements

Materialized query tables (MQTs) have been enhanced in several ways, resulting in:

- Better support for designing MQTs and for troubleshooting queries that involve MQTs. The Explain output now indicates which MQTs were considered (but not chosen) by the optimizer for a query access plan; and also provides information about the reason. For example, the information can help determine whether the reason related to cost or an MQT not being a close enough match to be used.
- Better query performance for specific types of queries, such as for:
 - Queries that insert, update, or delete data in partitioned database environments where the MQT is replicated across database partitions
 - Queries that have expressions that are equivalent, but not identical, to those in the MQT definition
- Improved MQT maintenance. This might shorten refresh times, especially where the design consists of a logical hierarchy of MQTs. In this case, MQTs at the base of the hierarchy are refreshed first, and those results are applied to MQTs higher in the hierarchy.

Related concepts:

- “Materialized query tables” in *Performance Guide*
- “Performance enhancements summary” on page 89

Online processing of the SET INTEGRITY statement

You can now use the SET INTEGRITY statement to perform online integrity processing. Tables that are being processed by the SET INTEGRITY statement are no longer required to be completely offline during the entire duration of statement execution. You now have the option to allow read access, read and write access, or no access to a table while it is being processed for integrity.

You can also now use the SET INTEGRITY statement to:

- Generate values for an identity column, if one has been defined for the table. The SET INTEGRITY statement gives you the option of requesting identity column values for attached rows only, or for all rows in the table, including attached rows, loaded rows, and existing rows. You also have the option of leaving the current identity column values for all rows in the table unchanged.
- Request that (when the CURRENT REFRESH AGE special register is set to 'ANY') a REFRESH DEFERRED materialized query table be used to optimize the maintenance of another REFRESH DEFERRED materialized query table. (REFRESH IMMEDIATE materialized query tables are always considered during query optimization.)

Related concepts:

- “Materialized query table (MQT) enhancements” on page 90
- “Performance enhancements summary” on page 89

Related reference:

- “ALTER TABLE statement” in *SQL Reference, Volume 2*
- “Exception tables” in *SQL Reference, Volume 1*
- “REFRESH TABLE statement” in *SQL Reference, Volume 2*
- “SET INTEGRITY statement” in *SQL Reference, Volume 2*

Optimizer registry variables and keywords added

A new DB2_OPT_MAX_TEMP_SIZE registry variable is added:

You can use the new DB2_OPT_MAX_TEMP_SIZE registry variable to limit the amount of space that queries can use in temporary table spaces. Setting DB2_OPT_MAX_TEMP_SIZE can cause the optimizer to choose a plan that is more expensive than it would otherwise have chosen but which uses less space in the temporary table spaces. If you set DB2_OPT_MAX_TEMP_SIZE, be sure to balance your need to limit the use of temporary table space against the efficiency of the plan that your setting causes to be chosen.

Two new keywords NO_SORT_MGJOIN and NO_SORT_NLJOIN are added to the DB2_REDUCED_OPTIMIZATION registry variable:

The new NO_SORT_MGJOIN keyword instructs the optimizer to generate query plans that do *not* force sorts for merge scan joins (MSJN). The new

NO_SORT_NLJOIN keyword instructs the optimizer to generate query plans that do *not* force sorts for nested loop joins (NLJN). Use both keywords carefully, as they can severely impact performance.

Related concepts:

- “Performance enhancements summary” on page 89

Related reference:

- “Query compiler variables” in *Performance Guide*

Row compression support added

DB2 Version 9.1 includes a new dictionary-based row compression feature that you can use to compress data objects. When you compress data, you achieve disk storage space savings by representing the same data using fewer database pages. Large tables with rows that contain repetitive patterns will benefit from this feature.

Query performance might improve for tables that use row compression. Fewer I/O operations might be required to access the compressed data and more data can be cached in the buffer pool when it is compressed. Because user data is compressed within log records, it is possible that log records become smaller. For UPDATE log records, this might not be true.

The cost associated with row compression is due to extra CPU cycles needed to compress and decompress data. Compression and decompression is done on a per-row basis when the data within a row is accessed. To estimate the storage savings to be gained by using row compression, you can use the ROWCOMPESTIMATE option of the DB2 online INSPECT utility.

Rows cannot be compressed until the COMPRESS attribute for the table is enabled and a compression dictionary is created. The table COMPRESS attribute can be set via either the CREATE or ALTER TABLE statements. Use the REORG TABLE command to create a compression dictionary. When the REORG command is processed, all existing table rows are subject to being compressed.

Data row compression does not apply to index, LOB, LF, or XML objects.

Data row compression (COMPRESS clause) can be used with the existing value compression (VALUE COMPRESS clause). Use the following information to help you determine which method you should use, or whether you should use both methods together:

- Value compression provides an alternate method of representing the internal storage format of a data row. The disk storage savings depend on the table column definition, for more information see the CREATE TABLE statement. There is little overhead added by the implementation and use of this row format. Value compression is supported when using the DATA CAPTURE CHANGES clause.
- Data row compression does not depend on the table column definition. It replaces common byte patterns in a data row with shorter symbol strings. The storage savings are greater than the savings provided with value compression. However, there is added cost in implementing row compression as well as

processing costs associated with every time a row must be compressed or decompressed. Data row compression is not supported when using the DATA CAPTURE CHANGES clause.

Related concepts:

- “Performance enhancements summary” on page 89
- “Space compression for tables” in *Administration Guide: Implementation*
- “Space value compression for existing tables” in *Administration Guide: Implementation*
- “Space value compression for new tables” in *Administration Guide: Implementation*

Related reference:

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

64 KB page size support added for POWER5+ processors (AIX)

The AIX 5.3E operating system running on POWER5+ hardware provides a new 64 KB memory page size. 64 KB memory pages are fully pageable by the operating system, and the number of available 64 KB memory pages provided by the operating system can grow or shrink as needed, depending on system usage. Version 9.1 of the DB2 database manager allocates 64 KB memory pages when appropriate to enhance its performance when running on the AIX 5.3E operating system on POWER5+ hardware.

Related concepts:

- “Performance enhancements summary” on page 89

Chapter 13. Scalability enhancements

Scalability enhancements summary

Scalability in DB2 Version 9.1 is improved with several enhancements, including larger row identifiers that allow more data pages per object and more records per page, larger system temporary tables, fast communications support for parallel communications daemons, and indexes that you can define on a greater number of columns. These enhancements and others help ensure that DB2 Database for Linux, UNIX and Windows provides you with the scalability you need for your growing business.

For more information about the scalability enhancements introduced in DB2 Version 9.1, follow the links in the “Related concepts” section.

Related concepts:

- “Fast communications manager (FCM) enhancements” on page 82
- “Table size limits increased to 1.1 trillion rows and 16 terabytes” on page 96
- “Larger system and user temporary tables” on page 95
- “Increased maximum number of index columns and maximum size of index keys” on page 95

Increased maximum number of index columns and maximum size of index keys

The maximum number of columns in an index has been increased from 16 to 64. In addition, the maximum size of an index key now varies depending on the index page size. The maximum is:

- 1 KB on a 4 KB page
- 2 KB on an 8 KB page
- 4 KB on a 16 KB page
- 8 KB on a 32 KB page

Previously, the maximum size of an index key was 1 KB, regardless of the index page size.

Related concepts:

- “Scalability enhancements summary” on page 95

Related reference:

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

Larger system and user temporary tables

System and user temporary table spaces hold temporary data required by the database manager while performing operations such as sorts or joins. A larger record identifier to uniquely identify rows within the database allows for larger system or user temporary tables to be held in the system or user temporary table spaces.

Related concepts:

- “Table and index management for standard tables” in *Performance Guide*
- “Scalability enhancements summary” on page 95
- “Table size limits increased to 1.1 trillion rows and 16 terabytes” on page 96

Table size limits increased to 1.1 trillion rows and 16 terabytes

Larger record identifiers allow more data pages per table object and the possibility of more records per page. Table spaces must not be “regular” to allow the use of large record identifiers; that is, both system temporary and user temporary table spaces support large record identifiers.

The default when you create a new DMS table space is now “large”, including when you specify `MANAGE BY AUTOMATIC STORAGE`. The `USERSPACE1` table space, when created as DMS as part of a `CREATE DATABASE` operation, is also “large” by default.

Large record identifiers do not apply to SMS table spaces.

Related concepts:

- “DMS table spaces” in *Administration Guide: Planning*
- “Larger system and user temporary tables” on page 95
- “Scalability enhancements summary” on page 95

Related reference:

- “`CREATE DATABASE` command” in *Command Reference*

Chapter 14. Security enhancements

Security enhancements summary

In DB2 Version 9.1, security enhancements such as label-based access control (LBAC), new SQL APIs, and a new security administrator authority level (SECADM) provide greater control over access to information assets and improved reporting capabilities. You can also benefit from enhanced security management capabilities with the RESTRICTIVE option on the CREATE DATABASE statement which allows greater control over database privileges. These security enhancements and other features will help you protect and manage your sensitive data.

For more information about security enhancements introduced in DB2 Version 9.1, follow the links in the “Related concepts” section.

Related concepts:

- “Data access security improved through label-based access control (LBAC)” on page 97
- “RESTRICT option added to CREATE DATABASE statement” on page 98
- “Security administrator (SECADM) authority added to centralize security privileges” on page 99
- “SETSESSIONUSER privilege added” on page 99
- “TRANSFER OWNERSHIP SQL statement added” on page 100

Data access security improved through label-based access control (LBAC)

Label-based access control (LBAC) greatly increases your control over who can access your data. LBAC lets you decide exactly who has write and has read access to individual rows and individual columns.

The LBAC capability is very configurable: you can tailor it to match your particular security environment. A security administrator, a user to whom the new SECADM authority has been granted, performs all LBAC configuration. The security administrator configures the LBAC system by creating *security policies*. A security policy describes the criteria that are used to decide who has access to specific data. Only one security policy can be used to protect any one table, but different tables can be protected by different security policies.

After creating a security policy, the security administrator creates *security labels* that are part of that policy. Exactly what makes up a security label is determined by the security policy and can be configured to represent the criteria that your organization uses to decide who should have access to particular data items. If you decide, for instance, that you want to look at a person’s position in the company or which projects that the person is part of to decide which data that person should see, you can configure your security labels so that each label includes that information. LBAC is flexible enough to let you set up anything from a system with very complicated criteria, to a system with very simple criteria, where each label represents either a “high” or a “low” level of trust.

Once created, a security label can be associated with individual columns and rows in a table to protect the data held there. Data that is protected by a security label is called *protected data*. The security administrator allows users access to protected data by granting them security labels. When a user tries to access protected data, that user's security label is compared to the security label protecting the data. The protecting security label will block some user security labels but not block others.

The security administrator can also grant exemptions to users. An *exemption* allows you to access protected data that your security labels might otherwise prevent you from accessing. Together your security labels and exemptions are called your *LBAC credentials*.

If you try to access a protected column that your LBAC credentials do not allow you to access then the access will fail and you will get an error message.

If you try to read protected rows that your LBAC credentials do not allow you to read then the database server acts as if those rows do not exist. Those rows cannot be selected as part of any SQL statement that you run, including SELECT, UPDATE, and DELETE. Even the aggregate functions ignore rows that your LBAC credentials do not allow you to read. The COUNT(*) function, for example, will return a count only of the rows to which you have read access.

Related concepts:

- “Label-based access control (LBAC) overview” in *Administration Guide: Implementation*
- “Security enhancements summary” on page 97

Related reference:

- “CREATE SECURITY LABEL COMPONENT statement” in *SQL Reference, Volume 2*
- “CREATE SECURITY LABEL statement” in *SQL Reference, Volume 2*
- “CREATE SECURITY POLICY statement” in *SQL Reference, Volume 2*
- “CREATE TABLE statement” in *SQL Reference, Volume 2*

RESTRICT option added to CREATE DATABASE statement

In DB2 UDB Version 8, when you create a new database, permissions to access objects in that database, such as SELECT access to the system catalog tables and views, are automatically granted to PUBLIC. The new RESTRICT option of the CREATE DATABASE statement changes this behavior. By including the RESTRICT option, you can grant only the permissions that you want. You do not have to worry about granting other permissions by default.

Related concepts:

- “Security enhancements summary” on page 97

Related tasks:

- “Creating a database” in *Administration Guide: Implementation*

Related reference:

- “CREATE DATABASE command” in *Command Reference*

Security administrator (SECADM) authority added to centralize security privileges

The new security administrator (SECADM) authority collects several security-related privileges under one authority. The abilities given by SECADM are not given by any other authority, not even SYSADM.

Having SECADM authority allows you to perform these actions:

- Create, drop, grant permission to access, or revoke the various objects that are part of label-based access control (LBAC)
- Use the TRANSFER OWNERSHIP statement on objects that you do not own
- Grant or revoke the SETSESSIONUSER privilege

Related concepts:

- “Database authorities” in *Administration Guide: Implementation*
- “Data access security improved through label-based access control (LBAC)” on page 97
- “Security enhancements summary” on page 97
- “SET SESSION AUTHORIZATION requires SETSESSIONUSER privilege” on page 129
- “SETSESSIONUSER privilege added” on page 99

Related reference:

- “TRANSFER OWNERSHIP statement” in *SQL Reference, Volume 2*

SETSESSIONUSER privilege added

You can grant the new SETSESSIONUSER privilege to a user or to a group. This privilege allows the holder to switch identities to any of the authorization IDs to which you granted the privilege. The identity switch is made by using the SQL statement SET SESSION AUTHORIZATION.

This new privilege provides more control over who has authority to switch session user identities. In DB2 UDB Version 8, users with DBADM or SYSADM authority could assume another user’s identity by using the SET SESSION AUTHORIZATION statement. In DB2 Version 9.1, the new SETSESSIONUSER privilege, which can only be granted by the security administrator authority (SECADM), is required to perform this task.

For backward compatibility and to avoid loss of existing user privileges upon migration to DB2 Version 9.1, any authorization ID that explicitly holds DBADM authority (as recorded in the SYSCAT.DBAUTH catalog view) is automatically granted the SETSESSIONUSER privilege on PUBLIC. A user who acquires DBADM authority after migration to DB2 Version 9.1 will not be able to change the session authorization ID unless they are explicitly granted the SETSESSIONUSER privilege.

Related concepts:

- “Security enhancements summary” on page 97
- “SET SESSION AUTHORIZATION requires SETSESSIONUSER privilege” on page 129

Related reference:

- “GRANT (SETSESSIONUSER Privilege) statement” in *SQL Reference, Volume 2*
- “REVOKE (SETSESSIONUSER Privilege) statement” in *SQL Reference, Volume 2*
- “SET SESSION AUTHORIZATION statement” in *SQL Reference, Volume 2*

TRANSFER OWNERSHIP SQL statement added

The new TRANSFER OWNERSHIP SQL statement provides the security administrator (with SECADM authority) or database object owner with the ability to change the ownership of a database object. For example, if an employee is leaving a company, all objects that he or she owns can be transferred to a different user. This task can be performed either by the leaving employee or a user with the SECADM authority. The TRANSFER OWNERSHIP SQL statement automatically grants the new owner the same privileges that the previous owner had when the object was created.

Related concepts:

- “Database authorities” in *Administration Guide: Implementation*
- “Authorization, privileges, and object ownership” in *Administration Guide: Implementation*
- “Security administrator (SECADM) authority added to centralize security privileges” on page 99
- “Security enhancements summary” on page 97

Related reference:

- “TRANSFER OWNERSHIP statement” in *SQL Reference, Volume 2*

Chapter 15. Troubleshooting and problem determination enhancements

Troubleshooting and problem determination enhancements summary

In DB2 Version 9.1, troubleshooting and problem determination enhancements provide increased control over the set of diagnostic information produced when you encounter problems. These enhancements provide valuable capabilities that will help address any problems you encounter.

For troubleshooting and problem determination enhancements introduced in DB2 Version 9.1, follow the links in the “Related concepts” section.

Related concepts:

- “Data collection improvements for database system failures” on page 101
- “High water mark option added for db2pd -fcm command” on page 101
- “Trace mask support added to the db2trc command” on page 102

Data collection improvements for database system failures

You now have increased control over the set of diagnostic information produced when the database manager encounters a panic, trap, exception, or segmentation violation. In such situations, the **db2cos** script is now automatically run. The **db2cos** script contains **db2pd** commands to collect information in an unlatched manner. You can edit the **db2cos** script to collect more or less information.

If you use the default **db2cos** script, a db2cos output file is created. You or DB2 customer support can use this file to troubleshoot the cause of the database manager problem.

Related concepts:

- “db2cos (callout script) output files” in *Troubleshooting Guide*
- “Troubleshooting and problem determination enhancements summary” on page 101

Related reference:

- “db2pd - Monitor and troubleshoot DB2 database command” in *Command Reference*

High water mark option added for db2pd -fcm command

In DB2 Version 9.1, the following changes to the **db2pd -fcm** command are introduced:

- You can use the new **hwm** option to see historical information about applications that consume large amounts of fast communication manager (FCM) resources. FCM usage statistics for applications are collected and retained so that even if the applications are not connected to the database, the details will be reported when you run the **db2pd** command with the **-fcm hwm** option.

- The output will now contain FCM channel usage statistics, including the high and low water mark values with respect to the number of channels used.

Related reference:

- “db2pd - Monitor and troubleshoot DB2 database command” in *Command Reference*

Trace mask support added to the db2trc command

The **db2trc** (trace) command now has the ability to set trace masks. Trace masks allow you to limit the operations recorded by the trace facility. They are provided by DB2 Support technical support as necessary.

You can add trace mask elements, delete trace mask elements, set the trace mask to a specific value, load a trace mask theme or load a list of trace mask actions by issuing the **db2trc** command with the **-M** option. You can also change values associated with trace mask that is already running by specifying the **change** option.

Related concepts:

- “Troubleshooting and problem determination enhancements summary” on page 101

Related reference:

- “db2trc - Trace command” in *Command Reference*

Part 2. What's changed from Version 8.2

Read this part if you are interested in changed, deprecated, or discontinued functionality. This information will help you safely migrate your work to Version 9.1.

This part contains the following topics:

Chapter 16, "Changes in existing functionality"

This chapter outlines the changes to existing DB2 database system functionality, including changes related to installation, migration, and fix packs; database setup; database administration; and application development.

Chapter 17, "Deprecated functionality"

This chapter lists the deprecated functionality, which refers to specific functions or features that are supported but are no longer recommended and might be removed in a future release.

Chapter 18, "Discontinued functionality"

This chapter lists functions and features that are no longer supported in DB2 Version 9.1.

Chapter 16. Changes in existing functionality

Changes in existing functionality summary

The following topics outline changes to existing functionality related to DB2 installation and setup, DB2 administration, and application development.

For additional information about changed functionality in DB2 Version 9.1, see "Version 9 incompatibilities with previous releases and changed behaviors" in *Administration Guide: Planning*.

Related concepts:

- "Administration changes summary" on page 123
- "Application development changes summary" on page 132
- "Database setup changes summary" on page 112
- "Installation, migration, and fix pack changes summary" on page 105
- "Product, packaging, and terminology changes summary" on page 25

Installation, migration, and fix pack changes

Installation, migration, and fix pack changes summary

For information about changes to existing DB2 database installation, migration, and fix pack functionality in DB2 Version 9.1, follow the links in the "Related concepts" section.

Related concepts:

- "Version 9 incompatibilities with previous releases and changed behaviors" in *Administration Guide: Planning*
- "Installation CD changes for national language versions of DB2 products (Linux and UNIX)" on page 108
- "Licensing support changes" on page 108
- "Locally installed Information Center installation restrictions" on page 108
- "Migration support changes during installation (Windows)" on page 109
- "Migration support changes for DB2 database systems" on page 110
- "Multiple CD installation now required for some products (Linux and UNIX)" on page 111
- "Application and routine migration changes" on page 106
- "db2_deinstall command parameter changes (Linux and UNIX)" on page 106
- "DB2 install images package format changes (Linux and UNIX)" on page 106
- "Documentation updates are available between product releases" on page 107

Application and routine migration changes

The steps required to successfully migrate existing applications and routines are different than for previous release migrations, particularly if the DB2 database server or DB2 database clients are migrated from a 32-bit environment to a 64-bit environment.

Although, depending on the application programming interface, few steps might be required to migrate applications and routines, in some cases due to application and routine support changes, application and routine migration might entail modifying and rebuilding source code. In anticipation of this impact, tooling and documentation support is provided to help you with this task.

Related concepts:

- “Application and routine feature support changes” on page 132
- “Hardware and operating system support” on page 17
- “Installation, migration, and fix pack changes summary” on page 105

Related tasks:

- “Migrating 32-bit database applications to run on 64-bit instances” in *Migration Guide*
- “Migrating 32-bit external routines to run on 64-bit instances” in *Migration Guide*

db2_deinstall command parameter changes (Linux and UNIX)

In previous versions of DB2, the **db2_deinstall** command was used to uninstall DB2. In Version 9.1, you must specify the **-a** option to remove all installed DB2 products in the current location.

Related concepts:

- “Changes in existing functionality summary” on page 105
- “Installation, migration, and fix pack changes summary” on page 105

Related tasks:

- “Removing DB2 products using the db2_deinstall or doce_deinstall command (Linux and UNIX)” in *Quick Beginnings for DB2 Servers*

Related reference:

- “db2_deinstall - Uninstall DB2 products or features command” in *Command Reference*

DB2 install images package format changes (Linux and UNIX)

Prior to V9.1, DB2 install images on Linux and UNIX are created in their respective operating system formats such as rpm on Linux or lpp on AIX. You could install these images using DB2 installation code or using operating system utilities such as pkgadd, rpm, SMIT, and swinstall. To query information about the DB2 products installed, you could use operating system utilities such as pkginfo, rpm, lspp and swinfo.

In Version 9.1, the DB2 install images no longer uses the operating system formats. To enable you to install multiple DB2 copies on the same system, all DB2 install images for Linux and UNIX are compressed in a tar.gz format. You should use the

DB2 installation programs to ensure that your DB2 products are deployed and set up correctly. If you have scripts that you used to install DB2 products using operating system commands, you must modify them to call DB2 installation programs (**db2setup** or **db2_install**) instead.

In Version 9.1, you can use only the **db2ls** command to query the installation of a DB2 product; you can no longer use Linux and UNIX operating system utilities such as **pkgadd**, **rpm**, **SMIT**, and **swinstall**. If you used scripts containing operating system commands to query DB2 installation packages, you must modify them to use the **db2ls** command.

Related concepts:

- “Multiple DB2 copies on the same computer (Linux and UNIX)” in *Installation and Configuration Supplement*
- “Database setup changes summary” on page 112
- “Installation, migration, and fix pack changes summary” on page 105

Related reference:

- “db2_install - Install DB2 product command” in *Command Reference*
- “db2ls - List installed DB2 products and features command” in *Command Reference*
- “db2setup - Install DB2 command” in *Command Reference*
- “Multiple DB2 copies roadmap” in *Administration Guide: Implementation*

Documentation updates are available between product releases

Some of the Information Center content may be updated after the initial release. The Information Center hosted on the IBM web site will always be up to date and therefore the following information only applies when viewing a locally installed Information Center.

To determine if there is an update available for a specific topic compare the ‘Last updated’ value in your locally installed topic to the same topic on the IBM hosted Information Center. The ‘Last updated’ value as well as the URL for the IBM hosted topic can be found at the bottom of most topics.

Not all topics will be refreshed in an update, however, so it is possible that the above comparison will show no change to a given topic even though there may be updates to other topics in the Information Center. To determine if there is an update available for the entire Information Center look for the ‘Last updated’ value on the Information Center home page. Compare the value in your locally installed Information Center home page to the latest value which is available on the IBM hosted Information Center home page at <http://publib.boulder.ibm.com/infocenter/db2help>.

Related concepts:

- “Installation, migration, and fix pack changes summary” on page 105
- “Navigating in the Information Center” in *Online DB2 Information Center*

Installation CD changes for national language versions of DB2 products (Linux and UNIX)

If you are installing a non-English version of DB2 Version 9.1, you must use one or more additional CDs called a National Language Package. This package contains national language support files (product files that are specific to a language).

Related concepts:

- “Installation, migration, and fix pack changes summary” on page 105

Related reference:

- “National language versions” in *Administration Guide: Planning*

Licensing support changes

Licensing DB2 database products is more flexible and efficient in Version 9.1.

You can manage licenses for DB2 products and features through the License Center or using the Licensed Management Tool command, **db2licm**. The License Center checks license information, statistics, and user access for each installed product and feature. The **db2licm** command performs basic functions such as adding, removing, listing, and modifying licenses and policies.

The License Center and the **db2licm** command have been updated in Version 9.1 to support DB2 product packaging changes.

For details regarding DB2 product packaging and licensing, see DB2 Database for Linux, UNIX, and Windows home page at www.ibm.com/db2/udb.

Related concepts:

- “License Center overview” in *Administration Guide: Implementation*
- “Installation, migration, and fix pack changes summary” on page 105

Related reference:

- “db2licm - License management tool command” in *Command Reference*
- “DB2 product license files” in *Installation and Configuration Supplement*

Locally installed Information Center installation restrictions

By default, DB2 products access DB2 documentation at 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*.

The following installation and coexistence restrictions exist:

- The Information Center is supported only on specific versions of Windows and Linux operating systems. For the list of operating systems where you can install the Information Center, refer to the installation topics in the related links section.
- You must install the Version 9 Information Center in a different location from other DB2 Version 9 products.

- You can install multiple Information Centers on a system, but you can only have one copy of each version. For example, you can have one Version 8.2 Information Center and one Version 9 Information Center on the same computer.
- To update the content of the Information Center, use the new documentation update feature of the Information Center; no documentation fixpaks will be provided.

In Version 9.1, the default installation paths for locally-installed Information Center have changed to:

- /opt/ibm/db2ic/V9/ (for Linux operating systems)
- C:\Program Files\IBM\DB2 Information Center\Version 9 (for Windows operating systems)

To remove a Version 9 Information Center, use the following methods:

- On the Linux platforms, use the **doce_deinstall** command. When you use this command, all files in the Information Center directories are removed.
- On the Windows platforms, use the Add/Remove Programs windows, accessible through the Windows Control Panel.

Related concepts:

- “Documentation updates are available between product releases” on page 107
- “Installation, migration, and fix pack changes summary” on page 105

Related tasks:

- “Installing the DB2 Information Center using the DB2 Setup wizard (Linux)” in *Quick Beginnings for DB2 Servers*
- “Installing the DB2 Information Center using the DB2 Setup wizard (Windows)” in *Quick Beginnings for DB2 Servers*

Related reference:

- “doce_deinstall - Uninstall DB2 Information Center command” in *Command Reference*
- “doce_install - Install DB2 Information Center command” in *Command Reference*

Migration support changes during installation (Windows)

In most cases, when you install DB2 Version 9 on the Windows operating system, you have the option to migrate an existing DB2 UDB Version 8 installation or to install a new copy of DB2 Version 9. If you choose to migrate your existing installation, your DB2 instances and DB2 Administration Server (DAS) are migrated automatically.

If you choose to install a new copy, you can manually migrate your instances by running the **db2imigr** command and manually migrate the DAS by running the **dasmigr** command after installation. You can also create new instances by using the **db2icrt** command.

If you install the Version 9 DB2 Runtime Client on Windows, the only option is to install as a new copy. After installation, you can manually migrate the client instance by running the **db2imigr** command.

Related concepts:

- “Migration to DB2 Version 9” in *Migration Guide*

- “Migration essentials for DB2 clients” in *Migration Guide*
- “Migration essentials for DB2 servers” in *Migration Guide*
- “Migration planning for your DB2 environment” in *Migration Guide*
- “Installation, migration, and fix pack changes summary” on page 105

Related tasks:

- “Installing DB2 clients (Windows)” in *Quick Beginnings for DB2 Clients*
- “An overview of installing your DB2 product (Windows)” in *Quick Beginnings for DB2 Servers*
- “Migrating a DB2 server (Windows)” in *Migration Guide*
- “Migrating a DB2 Client (Windows)” in *Migration Guide*
- “Migrating a DB2 Runtime Client (Windows)” in *Migration Guide*

Migration support changes for DB2 database systems

These are the new restrictions and changes in support for database system migration in DB2 Version 9:

Migration support changes in DB2 Version 9

- Migration to DB2 Version 9 from DB2 UDB Version 8 is supported. If you have DB2 UDB Version 7 or earlier, you cannot migrate directly to DB2 Version 9; you need to migrate to DB2 UDB Version 8 first.
- Migration is supported from a system with multiple DB2 database system copies.
- Instance migration is not supported on DB2 database systems where DB2 Data Links Manager Version 8 or DB2 Data Warehouse Manager Version 8 is installed. DB2 Data Links Manager and DB2 Data Warehouse Manager are not supported in DB2 Version 9. However, you could migrate to a DB2 Version 9 instance if you meet certain requirements.
- Instance profile registry variables are migrated when you migrate your instances to DB2 Version 9. However, the global profile registry variables set by the user are not migrated.
- Database migration fails if a database has:
 - User-defined distinct types (UDTs) with the names XML, BINARY, or VARBINARY are created in the database
 - User objects that use the system-defined DATALINK data type
 - High availability disaster recovery (HADR) database role is set to STANDBY

New migration restrictions in DB2 Version 9

- The NetBIOS protocol is no longer supported in DB2 Version 9. NetBIOS is not a valid keyword for the DB2COMM registry variable, and cataloged nodes and databases using this protocol do not work.
- Migration from certain versions of UNIX, Linux and Windows operating systems – such as AIX 4.3.3, Solaris 8, Windows 95, Windows 98, Windows NT[®] and Windows Me – is not supported. You need to upgrade to a supported version of the operating system before you migrate to DB2 Version 9. For a complete list of operating systems supported, visit the DB2 system requirements Web page at www.ibm.com/software/data/db2/udb/sysreqs.html.

- You must install 64-bit kernels prior to installing DB2 Version 9 on AIX, HP-UX, Solaris, and Linux (zSeries, POWER, or x86_64) operating systems.
- Instance bit size is determined by the platform where DB2 Version 9 is installed, and support for 32-bit kernels and 64-bit kernels has changed.

Support changes in DB2 Version 9 that impact routines

- 32-bit external unfenced stored procedures and user-defined functions (UDFs) cannot run on DB2 Version 9 64-bit instances.
- If you created SQL procedures in DB2 UDB Version 8.1 prior to FixPak 7 or Version 8.2 and migrated from a DB2 UDB Version 8 32-bit instance to a DB2 Version 9 64-bit instance, you need to re-create your SQL procedures using DB2 Version 9.
- If you migrate to a DB2 Version 9 64-bit instance and you have Java external routines, the **db2imigr** command sets the JDK_PATH database configuration parameter to INSTHOME/sql1lib/java/jdk64 on Linux and UNIX or DB2PATH\java\jdk on Windows. The JDK_PATH parameter must be set to a 64-bit JVM installation path to be able to run your Java external routines and DB2 tools.

For more details about changes in migration support, review the migration essentials for DB2 servers, DB2 clients, database applications, and routines.

For more details on migration concepts and what migration involves, review Migration to DB2 Version 9 in the *Migration Guide*.

Related concepts:

- “Functionality deprecated or discontinued in DB2 database products that impacts migration” in *Migration Guide*
- “Migration essentials for database applications” in *Migration Guide*
- “Migration essentials for DB2 clients” in *Migration Guide*
- “Migration essentials for routines” in *Migration Guide*
- “Migration overview for DB2 servers” in *Migration Guide*
- “Migration restrictions for DB2 servers” in *Migration Guide*
- “Migration essentials for DB2 servers” in *Migration Guide*
- “Support changes for 32-bit and 64-bit DB2 servers” in *Migration Guide*
- “Data Links Manager no longer supported” on page 148
- “Installation, migration, and fix pack changes summary” on page 105
- “The Data Warehouse Center and the Information Catalog Center are no longer included” on page 150

Related tasks:

- “Migrating SQL procedures” in *Migration Guide*
- “Recataloging nodes and databases that use NetBIOS and SNA protocols” in *Migration Guide*

Multiple CD installation now required for some products (Linux and UNIX)

Most DB2 product installation images are contained on a single CD. However, on Linux and UNIX operating systems, some DB2 database product installation images span more than a single CD, requiring multiple CDs to install the product.

When installing the DB2 database product, you can mount the CDs when prompted, or prior to installing the product, you can copy the installation image from the CDs to a file system. Then, during installation, you specify the location of the file system to which you copied the image. You can use either process when installing using the DB2 Setup wizard or a response file.

Multiple CDs are not required to install on Windows operating systems.

Related concepts:

- “Multiple CD installation (Linux and UNIX)” in *Quick Beginnings for DB2 Servers*
- “Installation, migration, and fix pack changes summary” on page 105

Related tasks:

- “Installing a DB2 product using a response file (Linux and UNIX)” in *Installation and Configuration Supplement*
- “Installing a DB2 product using a response file (Windows)” in *Installation and Configuration Supplement*

Database setup changes

Database setup changes summary

For changes to existing DB2 database setup functionality in DB2 Version 9.1, follow the links in the “Related concepts” section.

Related concepts:

- “Version 9 incompatibilities with previous releases and changed behaviors” in *Administration Guide: Planning*
- “Lock lists require additional space” on page 120
- “Product, packaging, and terminology changes summary” on page 25
- “Authority required to start the database manager has changed (Windows)” on page 113
- “32-bit DB2 database instance support changes” on page 121
- “Automatic configuration of prefetchers and page cleaners enabled by default” on page 113
- “Automatic self-tuning memory enabled by default during database creation” on page 113
- “Automatic statistics collection enabled by default during database creation” on page 114
- “Automatic storage enabled by default during database creation” on page 115
- “Configuration Advisor enabled by default during database creation” on page 116
- “Configuration parameters default value changes” on page 117
- “database_memory configuration parameter change” on page 119
- “Databases require additional space” on page 119
- “Increased log, table space, and memory requirements due to larger record identifiers (RIDs)” on page 120
- “Registry and environment variables default value changes” on page 121

Related reference:

- “Changes in DB2 registry variables, configuration parameters, and database physical design characteristics” in *Migration Guide*

Authority required to start the database manager has changed (Windows)

In previous versions of the DB2 database system, the database manager could be started by users belonging to the Administrators, Power Users, or System Operators groups.

In Version 9, the authority required to start the database manager on Windows depends on whether extended security is enabled.

- When extended security is enabled, the users must belong to the Administrators or DB2ADMNS group.
- When extended security is disabled, the users must belong to the Administrators, Power Users, or System Operators group.

Related concepts:

- “Extended Windows security using DB2ADMNS and DB2USERS groups” in *Administration Guide: Implementation*

Related reference:

- “db2start - Start DB2 command” in *Command Reference*
- “START DATABASE MANAGER command” in *Command Reference*

Automatic configuration of prefetchers and page cleaners enabled by default

Starting in IBM DB2 Version 9.1, the settings for the num_iocleaners and num_ioservers configuration parameters are set to AUTOMATIC by default. This means that the number of prefetchers and page cleaners started is based on environment characteristics such as the number of CPUs, the number of database partitions, and the parallelism settings of the table spaces in the database.

For existing databases, you can take advantage of this feature by setting the values of num_iocleaners and num_ioservers to AUTOMATIC.

Related concepts:

- “Configuration parameters default value changes” on page 117
- “Database setup changes summary” on page 112

Related reference:

- “num_iocleaners - Number of asynchronous page cleaners configuration parameter” in *Performance Guide*
- “num_ioservers - Number of I/O servers configuration parameter” in *Performance Guide*

Automatic self-tuning memory enabled by default during database creation

To simplify the configuration of several parameters that control memory resources, the self-tuning memory manager, which redistributes available memory resources

among memory consumers as workload requirements change, has been enabled by default for newly created database. You can enable self-tuning memory for migrated databases by setting the `self_tuning_mem` configuration parameter to ON and setting some or all of the following configuration parameters to AUTOMATIC:

- `pckcachesz`
- `locklist`
- `maxlocks`
- `sortheap`
- `sheapthres_shr`
- `database_memory` (You can only set `database_memory` to AUTOMATIC on Windows and AIX platforms.)

You can also have your buffer pools automatically tuned by setting their size to AUTOMATIC.

In Version 9.1, the `self_tuning_mem` database configuration parameter is automatically set to ON when you create a single-partition database. In addition, the configuration parameters listed above and the buffer pools are set to AUTOMATIC and therefore are automatically tuned. If you do not want to have the self-tuning memory manager enabled, you can turn it off by setting the `self_tuning_mem` configuration parameter to OFF after creating the database.

The self-tuning memory manager is not enabled by default on DPF systems.

Note: If you migrate a database from Version 8 to Version 9.1, this feature is not automatically enabled. To use this feature on a migrated database, you must enable it manually.

Related concepts:

- “Automatic features enabled by default” in *Administration Guide: Planning*
- “Self tuning memory” in *Performance Guide*
- “Automatic statistics collection enabled by default during database creation” on page 114
- “Automatic storage enabled by default during database creation” on page 115
- “Configuration Advisor enabled by default during database creation” on page 116
- “Configuration parameters default value changes” on page 117
- “`database_memory` configuration parameter change” on page 119
- “Database setup changes summary” on page 112

Automatic statistics collection enabled by default during database creation

Starting in IBM DB2 Version 9.1, automatic statistics collection (RUNSTATS) is enabled by default when a new database is created. This means that DB2 automatically determines which statistics are required by your workload and which statistics need to be updated. The RUNSTATS utility is then automatically run in the background, as needed, to ensure the correct statistics are collected and maintained.

You can disable automatic statistics collection by explicitly setting the `auto_runstats` database configuration parameter to OFF.

Note: If you migrate a database from Version 8 to Version 9.1, this feature is not automatically enabled. To use this feature on a migrated database, you must enable it manually.

Related concepts:

- “Automatic features enabled by default” in *Administration Guide: Planning*
- “Automatic statistics collection” in *Performance Guide*
- “Automatic statistics collection by table” in *Administration Guide: Planning*
- “Automatic statistics profiling using automatic statistics collection” in *Administration Guide: Planning*
- “Monitoring and notification for automatic features” in *Administration Guide: Planning*
- “Self tuning memory” in *Performance Guide*
- “Storage used by automatic statistics collection and profiling” in *Administration Guide: Planning*
- “Automatic self-tuning memory enabled by default during database creation” on page 113
- “Automatic storage enabled by default during database creation” on page 115
- “Configuration Advisor enabled by default during database creation” on page 116
- “Configuration parameters default value changes” on page 117
- “Database setup changes summary” on page 112

Related tasks:

- “Using automatic statistics collection” in *Performance Guide*

Automatic storage enabled by default during database creation

In Version 9.1, automatic storage is enabled by default when you create new databases. Automatic storage simplifies storage management by allowing you to specify *storage paths* where the database manager can place table space data and allocate space for various uses. If you do not want to use automatic storage, you must create your database by running the CREATE DATABASE command with the AUTOMATIC STORAGE option set to NO or by using the sqlcrea API with the SQLEDBDESCEXT parameter set to SQL_AUTOMATIC_STORAGE_NO.

The **db2look** command has changed due to the introduction of automatic storage databases to multiple partition configurations. You must now ensure that all database partitions are active before issuing the **db2look** command is issued. If any of the database partitions is not active, a warning message stating that DDL for a table space could not be generated is issued. This change to the **db2look** command affects all types of table spaces.

You might need to change the ALTER TABLESPACE statement and the RESTORE DATABASE command in DDL scripts when using automatic storage databases.

Related concepts:

- “Automatic self-tuning memory enabled by default during database creation” on page 113
- “Automatic statistics collection enabled by default during database creation” on page 114

- “Automatic storage enhancements” on page 79
- “Configuration Advisor enabled by default during database creation” on page 116
- “Database setup changes summary” on page 112
- “Automatic storage databases” in *Administration Guide: Implementation*
- “Automatic features enabled by default” in *Administration Guide: Planning*
- “About databases” in *Administration Guide: Planning*

Related tasks:

- “Adding an automatic storage path” in *Administration Guide: Implementation*

Related reference:

- “ALTER TABLESPACE statement” in *SQL Reference, Volume 2*
- “db2look - DB2 statistics and DDL extraction tool command” in *Command Reference*
- “RESTORE DATABASE command” in *Command Reference*

Configuration Advisor enabled by default during database creation

By default, new databases created in IBM DB2 Version 9.1 will benefit from the application of configuration recommendations provided by the Configuration Advisor. The Configuration Advisor generates recommendations for buffer pool size, database configuration parameter settings based on environment characteristics such as CPU speed and workload type.

For existing databases, you can take advantage of the Configuration Advisor enhancements by running the utility (the AUTOCONFIGURE command on the command line) against the database and selectively applying recommendations.

If you do not want to have the Configuration Advisor enabled by default at database creation, you can set the registry variable `DB2_ENABLE_AUTOCONFIG_DEFAULT` to NO before creating the database. If this registry variable is set to NO, the Configuration Advisor will not be invoked when creating a new database.

Note: If you migrate a database from Version 8 to Version 9.1, this feature will not be automatically enabled. To use this feature on a migrated database, you must enable it manually by running the Configuration Advisor through the Control Centre or via the AUTOCONFIGURE command.

Related concepts:

- “Automatic features enabled by default” in *Administration Guide: Planning*
- “Automatic self-tuning memory enabled by default during database creation” on page 113
- “Automatic statistics collection enabled by default during database creation” on page 114
- “Automatic storage enabled by default during database creation” on page 115
- “Configuration parameters default value changes” on page 117
- “Database setup changes summary” on page 112

Related tasks:

- “Defining the scope of configuration parameters using the Configuration Advisor” in *Administration Guide: Implementation*

Configuration parameters default value changes

The default values for the following configuration parameters have changed between V8.2 and V9.1 of the DB2 database.

Table 3. Configuration parameters with changed default values

| Parameter | V8.2 Default Value | V9.1 Default Value |
|-----------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| app_ctl_heap_sz - Application control heap size configuration parameter | <p>Database server with local and remote clients: 128</p> <p>Database server with local clients:</p> <ul style="list-style-type: none"> • 64 (on Windows and Linux operating systems) • 128 (on UNIX operating systems) <p>Partitioned database server with local and remote clients: 512</p> | <p>Database server with local and remote clients:</p> <ul style="list-style-type: none"> • 128 when INTRA_PARALLEL is not enabled • 512 when INTRA_PARALLEL is enabled <p>Database server with local clients:</p> <ul style="list-style-type: none"> • 64 (on Windows and Linux operating systems) when INTRA_PARALLEL is not enabled • 512 (on Windows and Linux operating systems) when INTRA_PARALLEL is enabled • 128 (on UNIX operating systems) when INTRA_PARALLEL is not enabled • 512 (on UNIX operating systems) when INTRA_PARALLEL is enabled <p>Partitioned database server with local and remote clients: 512</p> |
| auto_maint - Automatic maintenance configuration parameter | OFF | ON |
| auto_runstats - Automatic table runstats operations configuration parameter | OFF | ON |
| auto_tbl_maint - Automatic table maintenance configuration parameter | OFF | ON |
| avg_appls - Average number of active applications configuration parameter | 1 | AUTOMATIC |
| database_memory - Database shared memory size configuration parameter | AUTOMATIC | <ul style="list-style-type: none"> • AIX and Windows: AUTOMATIC • Linux, HP-UX, Solaris Operating System: Computed |
| java_heap_sz - Maximum Java interpreter heap size configuration parameter | 512 | <ul style="list-style-type: none"> • 32-bit platforms: 512 • 64-bit platforms: 1024 |

Table 3. Configuration parameters with changed default values (continued)

| Parameter | V8.2 Default Value | V9.1 Default Value |
|-----------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------|
| locklist - Maximum storage for lock list configuration parameter | <ul style="list-style-type: none"> • UNIX: 100 • Windows Database server with local and remote clients: 50 • Windows 64-bit Database server with local clients: 50 • Windows 32-bit Database server with local clients: 25 | AUTOMATIC |
| maxlocks - Maximum percent of lock list before escalation configuration parameter | <ul style="list-style-type: none"> • UNIX: 10 • Windows: 22 | AUTOMATIC |
| num_iocleaners - Number of asynchronous page cleaners configuration parameter | 1 | AUTOMATIC |
| num_ioservers - Number of I/O servers configuration parameter | 3 | AUTOMATIC |
| pckcachesz - Package cache size configuration parameter | -1 | AUTOMATIC |
| sheapthres - Sort heap threshold configuration parameter | <ul style="list-style-type: none"> • UNIX 32-bit platforms: 20 000 • Windows 32-bit platforms: 10 000 • 64-bit platforms: 20 000 | 0 |
| sheapthres_shr - Sort heap threshold for shared sorts configuration parameter | <i>sheapthres</i> | AUTOMATIC |
| sortheap - Sort heap size configuration parameter | 256 | AUTOMATIC |
| userexit - User exit enable configuration parameter | No | Off |

For a list of deprecated configuration parameters, see "Deprecated and discontinued features" in *Administration Guide: Planning*.

Related concepts:

- "Automatic configuration of prefetchers and page cleaners enabled by default" on page 113
- "Automatic self-tuning memory enabled by default during database creation" on page 113
- "Automatic statistics collection enabled by default during database creation" on page 114
- "Automatic storage enabled by default during database creation" on page 115
- "Automatic table and index reorganization enhancements" on page 80
- "Configuration Advisor enabled by default during database creation" on page 116
- "database_memory configuration parameter change" on page 119
- "sheapthres_shr parameter change" on page 130
- "Database setup changes summary" on page 112

Related reference:

- “Configuration parameters summary” in *Performance Guide*
- “Changes in DB2 registry variables, configuration parameters, and database physical design characteristics” in *Migration Guide*
- “RESET DATABASE CONFIGURATION command” in *Command Reference*

Databases require additional space

To accommodate the new Version 9.1 functionality, you must allocate more space for database objects than you would have for the same objects in Version 8.

Follow the related links for details on the Version 9.1 space requirements.

Related concepts:

- “Space requirements for database objects” in *Administration Guide: Planning*
- “Space requirements for indexes” in *Administration Guide: Planning*
- “Space requirements for log files” in *Administration Guide: Planning*
- “Space requirements for system catalog tables” in *Administration Guide: Planning*
- “Space requirements for user table data” in *Administration Guide: Planning*
- “Database setup changes summary” on page 112
- “Increased log, table space, and memory requirements due to larger record identifiers (RIDs)” on page 120

database_memory configuration parameter change

In DB2 Version 9.1, the COMPUTED setting of the database_memory configuration parameter is equivalent to the DB2 UDB Version 8 AUTOMATIC setting. If you want to maintain DB2 Version 8 behavior, set the database_memory parameter to COMPUTED. In DB2 Version 9.1, setting the database_memory parameter to AUTOMATIC enables the new self-tuning memory management feature that automatically tunes database memory usage.

The AUTOMATIC setting of database_memory is supported on the AIX and Windows operating systems only. On other platforms, the default value of database_memory is COMPUTED.

A new database configuration parameter, db_mem_thresh, has been added to control how much physical RAM is consumed by unused portions of the database_memory parameter. For more information, see “db_mem_thresh - Database memory threshold configuration parameter” in *Performance Guide*

Related concepts:

- “Adaptive, self-tuning memory allocation” on page 77
- “Configuration parameters default value changes” on page 117
- “Database setup changes summary” on page 112

Related reference:

- “database_memory - Database shared memory size configuration parameter” in *Performance Guide*
- “db_mem_thresh - Database memory threshold configuration parameter” in *Performance Guide*

Increased log, table space, and memory requirements due to larger record identifiers (RIDs)

Larger record identifiers (RIDs) which allow more data pages per table object and more records per page change the required amount of memory and space used by log files and system temporary table spaces.

Due to the increase in the RID sizes to support LARGE table space, the growth rate for log files and size of log records will increase. See "DB2 log records" in *Administrative API Reference* for current log record structures.

An increase in memory usage will result as each additional RID now requires 8 bytes of memory for single-partition environment and 16 bytes of memory for partitioned database environment.

The use of larger RIDs causes an increase of the row size of your result sets from queries or positioned updates. If the row size in your result sets is close to the maximum row length limit for your existing system temporary table space with the largest page size, you might need to create a system temporary table space with a larger page size. If your system temporary table space page size is 32 KB, you need to reduce the length of the information retrieved by your query or split your query.

Related concepts:

- "Space requirements for database objects" in *Administration Guide: Planning*
- "Database setup changes summary" on page 112
- "Databases require additional space" on page 119

Related reference:

- "REORGCHK command" in *Command Reference*
- "DB2 log records" in *Administrative API Reference*

Lock lists require additional space

Locking is the mechanism that the database manager uses to control concurrent access to data in a database by multiple applications. A lock list contains the locks held by all applications concurrently connected to a database.

In DB2 Version 9.1, the space required by each lock in a lock list has changed for some platforms, meaning that a lock list of a given size can no longer represent as many locks as it did before. Lock sizes have changed as follows:

- On 32-bit platforms, each lock requires 48 bytes of the lock list to record a lock on an object that has an existing lock on it. This value was 40 bytes in DB2 UDB Version 8.
- On 64-bit HP-UX/PA-RISC systems, each lock requires 80 bytes of the lock list to record a lock on an object that has an existing lock on it. This value was 64 bytes in DB2 UDB Version 8.

In addition to the increase in memory requirements, you will need to increase the lock list size by up to 1.7% for metadata overhead. On 64-bit HP-UX/PA-RISC systems, this overhead will be up to 2.5% of the total lock list.

Related concepts:

- “Configuration parameters default value changes” on page 117
- “Database setup changes summary” on page 112

Related reference:

- “locklist - Maximum storage for lock list configuration parameter” in *Performance Guide*

Registry and environment variables default value changes

The default values for the following variables have changed between V8.2 and V9.1 of the DB2 database.

Table 4. Registry and environment variables with changed default values

| Variable Name | V8.2 Default Value | V9.1 Default Value |
|-----------------------------------------------------------------------------------|--------------------|--------------------|
| DB2CHKSQLDA | OFF | ON |
| DB2_ALLOCATION_SIZE | 8 MB | 128 KB |
| DB2_COLLECT_TS_REC_INFO | OFF | ON |
| DB2_FORCE_FCM_BP ¹ | NO | YES |
| DB2_MDC_ROLLOUT | OFF | ON |
| DB2_SMS_TRUNC_TMPTABLE_THRESH | not set | 0 |
| Note: The DB2_FORCE_FCM_BP registry variable is deprecated in Version 9.1. | | |

Related concepts:

- “Database setup changes summary” on page 112

Related reference:

- “Changes in DB2 registry variables, configuration parameters, and database physical design characteristics” in *Migration Guide*
- “General registry variables” in *Administration Guide: Implementation*
- “Performance variables” in *Performance Guide*

32-bit DB2 database instance support changes

Growing business data processing needs for faster, more scalable applications is driving requirements for increased processing power from hardware. 64-bit server support paves the way towards delivering higher processing power, more memory, and better application performance. The AIX, HP, Solaris, Linux on AMD64 and Intel EM64T, Linux on IA64, Linux on IBM System p, Linux on IBM System z, Windows x64, and Windows on IA64 operating systems all support 64-bit processors with a default 64-bit kernel and a default 64-bit user space.

Responding to market demand while also building on the 64-bit DB2 database server support introduced in DB2 UDB Version 8, IBM is reducing the number of supported 32-bit platforms and making DB2 database server support on 64-bit hardware and operating systems a higher priority. IBM will, however, continue to support those 32-bit Windows and Linux platforms that are often preferred for building or running small and medium business applications.

DB2 Version 9.1 database servers are supported primarily on 64-bit hardware and operating systems only. In general migrating to DB2 Version 9.1 requires that you

migrate existing 32-bit DB2 database server instances to 64-bit DB2 database server instances. There are three possible DB2 UDB Version 8 to DB2 Version 9 database server migrations:

- 32-bit Version 8 to 32-bit Version 9
- 32-bit Version 8 to 64-bit Version 9
- 64-bit Version 8 to 64-bit Version 9

Migrations from 32-bit DB2 UDB Version 8 database servers to 32-bit DB2 Version 9 database servers are only supported on the following operating systems:

- The Linux for x86 operating system
- The Windows for x86 operating system
- The Windows for x64 operating system (where the 32-bit DB2 database server for the Windows x86 operating system is supported)

Migration of a 32-bit DB2 database server instance to a 64-bit DB2 database server instance can impact the functioning of DB2 database applications and routines. Refer to the following topics for information.

Connection changes

- DB2 client to DB2 server connection support

Client application support changes

- Support for 32-bit database applications created in DB2 UDB Version 8 to continue working in DB2 Version 9
- Changed environment variable value settings to facilitate 32-bit and 64-bit development and deployment
- Updated sample build scripts are available for building new applications

External routine support changes

- Limited support for 32-bit external routines created in DB2 UDB Version 8 to continue working in DB2 Version 9
- 32-bit routines defined using the NOT FENCED clause created in DB2 UDB Version 8 will no longer work on 64-bit DB2 Version 9 database servers in AIX, HP, SUN, Linux for AMD64 and Intel EM64T, Linux on POWER, Linux on zSeries environments
- Updated sample build scripts are available for building new routines

JVM support changes

- 64-bit JVM will only be provided with 64-bit DB2 database servers
- 32-bit JVM will only be provided for Linux x86 and Windows on x86
- The 64-bit JVMs are no longer provided separately on a CD

Related concepts:

- “External routine implementation” in *SQL Guide*
- “Database setup changes summary” on page 112
- “Migration overview for database applications and routines” in *Migration Guide*
- “Support changes for 32-bit and 64-bit DB2 servers” in *Migration Guide*

Related tasks:

- “Building .NET CLR routine code” in *SQL Guide*
- “Building C and C++ routine code” in *SQL Guide*
- “Building Java routine code” in *SQL Guide*

- “Building SQLJ routines” in *Developing Java Applications*

Administration changes

Administration changes summary

For changes to existing DB2 administration-related functionality in DB2 Version 9.1, follow the links in the “Related concepts” section.

Related concepts:

- “Application ID format changed” on page 123
- “Catalog table changes in support of table partitioning” on page 124
- “Handling of new database connections has changed” on page 127
- “Backup image naming convention has changed (Windows)” on page 124
- “Version 9 incompatibilities with previous releases and changed behaviors” in *Administration Guide: Planning*
- “Collating sequence IDENTITY replaces BINARY in command output” on page 125
- “db2mtrk command changes (Windows)” on page 126
- “Data movement utilities changes” on page 125
- “DMS table space type default value changed to LARGE” on page 126
- “Load from cursor incompatibilities with previous releases” on page 127
- “Location of diagnostic messages for automatic maintenance” on page 127
- “New DB2 agents may affect MAX_CONNECTIONS configuration and application monitor output” on page 128
- “REORGCHK command output changes” on page 129
- “ROLLFORWARD DATABASE output has changed” on page 129
- “sheapthres_shr parameter change” on page 130
- “SET SESSION AUTHORIZATION requires SETSESSIONUSER privilege” on page 129
- “Table space-level point-in-time rollforward operations restrictions” on page 131
- “The -w option of the db2icrt, db2ilist, and db2iupdt commands is ignored” on page 145
- “TSM management class no longer used for filtering when restoring or retrieving logs” on page 131
- “Write-to-table event monitor changes” on page 131

Application ID format changed

The format of the application ID has changed. If you have scripts that parse output that contains the application ID (such as the output from the **LIST APPLICATIONS** command), you might need to update them.

The new format presents the port number and IP address in a readable form and accommodates the longer IPv6 addresses. The format for the application ID remains: IP address.port #.timestamp. However, the IP address for an IPv4 address is now in the form a.b.c.d, for example, 9.11.22.314. For an IPv6 address, the form is a:b:c:d:e:f:g:h, where each of a through h is four hexadecimal digits: for example, 2001:0db8:4545:2::09ff:fef7:62dc. The

timestamp is of the form yymmddhhmmss where yy is the year, mm is the month, dd is the day, hh is the hour, mm is the minutes, and ss is the seconds.

Related concepts:

- “Administration changes summary” on page 123
- “Internet Protocol Version 6 (IPv6) support added” on page 65

Related reference:

- “LIST APPLICATIONS command” in *Command Reference*

Backup image naming convention has changed (Windows)

As of DB2 Version 9.1, the naming convention for backup images stored on Windows operating systems has changed to match the naming convention used for all other operating systems. File names for backup images created on disk will now consist of a concatenation of several elements, separated by periods:

```
DB_alias.Type.Inst_name.NODEnnnn.CATNnnnn.timestamp.Seq_num
```

For example:

```
STAFF.0.DB201.NODE0000.CATN0000.20050922120112.001
```

DB2 Universal Database Version 8 and earlier versions used a four-level subdirectory tree when storing backup images on Windows operating systems:

```
DB_alias.Type\Inst_name\NODEnnnn\CATNnnnn\yyymmdd\hhmmss.Seq_num
```

Backup images from earlier versions of the product that use the previous naming structure can still be restored on V9.1 DB2 database systems.

Related concepts:

- “Backup overview” in *Data Recovery and High Availability Guide and Reference*
- “Administration changes summary” on page 123

Related tasks:

- “Using backup” in *Data Recovery and High Availability Guide and Reference*

Catalog table changes in support of table partitioning

There are a number of changes to the catalogs in support of table partitioning that impact information for both partitioned and non-partitioned tables.

- The physical object ID and tablespace ID have been moved from SYSCAT.TABLES.TABLEID and SYSCAT.TABLES.TBSPACEID to SYSCAT.DATAPARTITIONS.PARTITIONOBJECTID and SYSCAT.DATAPARTITIONS.TBSPACEID. For compatibility TABLEID and TBSPACEID in SYSCAT.TABLES is identical to the physical PARTITIONOBJECTID and TBSPACEID in SYSCAT.DATAPARTITIONS.PARTITIONOBJECTID for non partitioned tables.
- The fields TBSPACE, INDEX_TBSPACE, and LONG_TBSPACE in SYSCAT.TABLES are always null for partitioned tables. You need to look in one or both of SYSCAT.INDEXES and SYSCAT.DATAPARTITIONS to determine where data, indexes and LOBs are stored for partitioned tables.

Related concepts:

- “Administration changes summary” on page 123

Related reference:

- “SYSCAT.DATAPARTITIONS catalog view” in *SQL Reference, Volume 1*
- “SYSCAT.TABLES catalog view” in *SQL Reference, Volume 1*

Collating sequence IDENTITY replaces BINARY in command output

The identity collating sequence specifies that multibyte characters should be sorted in the way that they appear in their code point table. Prior to DB2 Version 9.1, the value BINARY was displayed for this collating sequence in the output generated by the **db2look** command and the GET DATABASE CONFIGURATION command. In V9.1, the value IDENTITY is always displayed for this collating sequence in the output of these commands. The collating sequence itself has not changed.

Related concepts:

- “Collating sequences” in *SQL Guide*
- “db2look command changes for native XML data store” on page 41

Related reference:

- “db2look - DB2 statistics and DDL extraction tool command” in *Command Reference*
- “GET DATABASE CONFIGURATION command” in *Command Reference*

Data movement utilities changes

DB2 Version 9.1 contains the following changes to the DB2 load, import, and export utilities:

- Importing IXF files: Starting with Version 8 FixPak 9, the import utility now creates indexes using the definition in the IXF file to enable or disable reverse scans of indexes. Prior to Version 8 FixPak 9, when importing an IXF file from a Version 8 client, reverse scans of indexes are disabled by default on Version 8 servers, and they are enabled by default on Version 9 servers.
- Recreating tables using the IXF file format: In V8, you could issue the IMPORT command with the CREATE option for most IXF files. In V9.1, if a feature is not available to be recreated during the import process of IXF files using the CREATE option, a warning is returned during the export and an error during the import process. In some cases, you can force the creation of tables from IXF files by specifying the file type modifier FORCECREATE. This new behavior affects only files exported using DB2 V9.1.
- Naming conventions for exporting LOB files: In version V9.1, the exported lob file has a .lob extension, for example, filename.001.lob, filename.002.lob. The default name is named after the input data file name, for example, <datafile>.001.lob, <datafile>.002.lob. If the input data file is generated in DB2 UDB V8, the DB2 V9.1 import utility can read it correctly.
- Moving LOB data: The default paths and the order in which the load, import, and export utilities search for these paths have changed.
- Exporting and importing LOB data: The LOBSINFILE keyword is specified automatically if you specify the LOBS TO or LOBFILE options in the EXPORT command or the LOBS FROM option in the IMPORT command. In DB2 UDB V8, if the LOBSINFILE file type modifier is not specified, then the specified LOBS TO, LOBS FROM, and LOBFILE options are ignored. In DB2 V9.1, specifying the LOBS TO or LOBFILE options implies the LOBSINFILE file type

modifier in the EXPORT command and specifying the LOBS FROM option implies the LOBSINFILE file type modifier in the IMPORT command.

For more information about changes to the export utility, see the "Export Overview" in *Data Movement Utilities Guide and Reference*.

For more information about changes to the import utility, see the "Import Overview" in *Data Movement Utilities Guide and Reference*.

For more information about changes to the load utility, see the "Load overview" in *Data Movement Utilities Guide and Reference*.

Related concepts:

- "Reverse scans enabled by default for indexes, primary keys and unique keys" on page 136
- "Exporting large objects (LOBS)" in *Data Movement Utilities Guide and Reference*
- "Importing large objects (LOBS)" in *Data Movement Utilities Guide and Reference*

Related reference:

- "PC/IXF data types" in *Data Movement Utilities Guide and Reference*

db2mtrk command changes (Windows)

The **db2mtrk** (memory tracker) command provides complete reports of memory status for instances, databases and agents. The following changes to the **-d** and **-i** options of the command are introduced in Version 9.1:

- The **-d** option, which shows database level memory, is now supported on the Windows platforms.
- The **-i** option, which shows instance level memory, no longer shows the database level memory since this information is available through the **-d** option.

Related concepts:

- "Memory allocation in DB2" in *Performance Guide*
- "Administration changes summary" on page 123

Related reference:

- "db2mtrk - Memory tracker command" in *Command Reference*

DMS table space type default value changed to LARGE

Beginning in Version 9.1, when database-managed space (DMS) table spaces are created, the default type is LARGE. Prior to this release, the default type was REGULAR. LARGE table spaces allow more data pages per table object and more records per page. The record identifiers (RID) length has been increased to support LARGE table spaces.

When a table space is not specified during a CREATE TABLE operation, the database manager can choose to create the table in a large table space. If there are multiple table spaces (including regular and large table spaces) that qualify, then the table can be created in the large table space and not just regular table spaces.

Related concepts:

- "DMS table spaces" in *Administration Guide: Planning*

- “Administration changes summary” on page 123
- “Table size limits increased to 1.1 trillion rows and 16 terabytes” on page 96

Handling of new database connections has changed

In previous versions of the DB2 database, you could not specify how long an agent would wait for a client’s connect request before timing out and disconnecting from the client. In Version 9, if a new connection does not send its initial connect request within the connect timeout period, the server will terminate the connection. The connect timeout period is specified in seconds and can be adjusted using the DB2_SERVER_CONTIMEOUT registry variable. The default connect timeout period is 180 seconds.

Related concepts:

- “Administration changes summary” on page 123
- “Connection timeout support for database applications added” on page 63

Related reference:

- “Miscellaneous variables” in *Performance Guide*

Load from cursor incompatibilities with previous releases

If you perform a load operation on a DB2 UDB Version 8 or earlier version using the CURSOR file type and the PARTITION_ONLY partitioned-db-cfg load option is specified, a set of distributed data files is created. However, you cannot use those files to perform a load operation on a DB2 V9.1 server using the CURSOR file type and the LOAD_ONLY partitioned-db-cfg option as those files are not compatible with the new server. Likewise, the distributed data files that you create on a DB2 V9.1 server are incompatible with a DB2 UDB Version 8 or earlier version server.

Related concepts:

- “Loading data in a partitioned database environment - hints and tips” in *Data Movement Utilities Guide and Reference*
- “Moving data using the CURSOR file type” in *Data Movement Utilities Guide and Reference*
- “Load overview” in *Data Movement Utilities Guide and Reference*
- “Administration changes summary” on page 123

Related tasks:

- “Loading data” in *Data Movement Utilities Guide and Reference*
- “Loading data in a partitioned database environment” in *Data Movement Utilities Guide and Reference*

Location of diagnostic messages for automatic maintenance

The diagnostic level and location of messages related to automatic maintenance has changed.

In DB2 Universal Database Version 8, whenever automatic maintenance health indicators were evaluated (to determine if maintenance was required), a diagnostic record was written in db2diag.log file. Whenever a maintenance operation occurred as a result of these evaluations, another entry was written in the

db2diag.log file. These diagnostic records were classified as "event" records and would appear when the diagnostic level of the instance (as specified in the diaglevel database manager configuration parameter) was set to values of 3 or 4.

The following changes are introduced in DB2 Version 9.1:

- Whenever automatic maintenance health indicators are evaluated, a diagnostic record is written in the db2diag.log file. If a maintenance operation occurs as a result of these evaluations, a diagnostic record is written in both the db2diag.log and the notification log.
- The diagnostic records associated with automatic maintenance are classified as "info" records.
- These diagnostic records will only be written when the diagnostic level (diaglevel) or notification level (notifylevel) of the instance is set to a value of 4.

Related concepts:

- "About automatic maintenance" in *Administration Guide: Planning*

Related reference:

- "auto_maint - Automatic maintenance configuration parameter" in *Performance Guide*

New DB2 agents may affect MAX_CONNECTIONS configuration and application monitor output

Two new DB2 agents, db2stmm and db2taskd, have been introduced in DB2 Version 9.1. The db2stmm agent is part of the new self-tuning memory feature added in DB2 Version 9.1. The db2taskd agent is an internal daemon that distributes background database tasks. These new agents remain connected to the database at all times and are not active if the database is activated in exclusive mode. If MAX_CONNECTIONS, a parameter that controls the maximum number of applications that can be connected to the instance, is configured tightly to limit the number of application connections, the introduction of these two new agents may require that you reconfigure the MAX_CONNECTIONS parameter to ensure that the number of connections available remains constant after migration to DB2 Version 9.1

Both agents are database-initiated system applications. They appear in the database system monitor. As a result, if you have a script or tool that monitors the output of the LIST APPLICATIONS command, modifications may be required to account for the two new agents.

Related concepts:

- "Adaptive, self-tuning memory allocation" on page 77
- "Administration changes summary" on page 123
- "Database system monitor" in *System Monitor Guide and Reference*

Related reference:

- "max_connections - Maximum number of client connections configuration parameter" in *Performance Guide*
- "LIST APPLICATIONS command" in *Command Reference*

REORGCHK command output changes

The output generated by the REORGCHK command is changed in Version 9.1. The SCHEMA and NAME columns are concatenated into one column (SCHEMA.NAME). In addition, the SCHEMA.NAME for each table and index is broken into two rows, one of the actual fully qualified name of the table, and one for the fully qualified name of each index on that table. The actual data for the remaining columns follows each index name.

Related concepts:

- “Administration changes summary” on page 123

Related tasks:

- “Determining when to reorganize tables” in *Performance Guide*

Related reference:

- “REORGCHK command” in *Command Reference*

ROLLFORWARD DATABASE output has changed

The ROLLFORWARD DATABASE command could accept both Coordinated Universal Time (UTC) and local time as valid input formats since DB2 Version 8. In Version 9.1, the format of the output timestamp is the same as the format specified in the input.

This new functionality introduces consistency and removes ambiguity from the Last committed transaction output timestamp from the ROLLFORWARD DATABASE command. Consistency is introduced by having the ROLLFORWARD DATABASE command remember whether the USING LOCAL TIME option was specified. Any subsequent ROLLFORWARD DATABASE commands that are issued without specifying this option will use this remembered information to decide if their output should be in local time or UTC. Ambiguity is removed from the Last committed transaction timestamp output by explicitly printing the word UTC or Local beside the timestamp in the CLP output for all ROLLFORWARD DATABASE commands.

Related concepts:

- “Administration changes summary” on page 123

Related reference:

- “ROLLFORWARD DATABASE command” in *Command Reference*

SET SESSION AUTHORIZATION requires SETSESSIONUSER privilege

In DB2 UDB Version 8, users with DBADM or SYSADM authority could assume different authorization IDs on the same connection using the SET SESSION AUTHORIZATION statement. In DB2 Version 9.1, changing the session authorization ID to a new value using the SET SESSION AUTHORIZATION statement requires that the authorization ID of the SQL statement has the new SETSESSIONUSER privilege. A security administrator (with the new SECADM authority) can grant this privilege by using the new GRANT SETSESSIONUSER statement.

For backward compatibility and to avoid loss of existing user privileges, any authorization ID that explicitly holds DBADM authority (as recorded in the SYSCAT.DBAUTH catalog view) is automatically granted the SETSESSIONUSER privilege when you migrate to DB2 Version 9.1. Users who acquire DBADM authority after you migrate to DB2 Version 9.1 cannot change the session authorization ID unless they are explicitly granted the SETSESSIONUSER privilege.

Related concepts:

- “Security administrator (SECADM) authority added to centralize security privileges” on page 99
- “SETSESSIONUSER privilege added” on page 99
- “Administration changes summary” on page 123
- “Security enhancements summary” on page 97

Related reference:

- “GRANT (SETSESSIONUSER Privilege) statement” in *SQL Reference, Volume 2*
- “REVOKE (SETSESSIONUSER Privilege) statement” in *SQL Reference, Volume 2*
- “SET SESSION AUTHORIZATION statement” in *SQL Reference, Volume 2*

sheapthres_shr parameter change

The sheapthres_shr parameter represents a limit on the total amount of database shared memory that can be used by sort memory consumers at any one time. In DB2 Version 8, this limit was a hard limit. If the sort memory was close to the limit, an alert was generated. In DB2 Version 9.1, the sheapthres_shr parameter represents a soft limit. The sort memory heap can consume additional, unreserved, database shared memory if needed.

In DB2 Version 8, only sorts in SMP environments or sorts running the concentrator that could be accessed by more than one agent consumed shared memory, regardless of the value of the sheapthres parameter. In DB2 Version 9.1, if you set the sheapthres database manager configuration parameter to 0, all sorts will use shared memory. If you set the sheapthres database manager configuration parameter to a value greater than 0, then only sorts in SMP environments or sorts running the concentrator that could be accessed by more than one agent will consume shared memory. This behavior is the same as in Version 8.

Related concepts:

- “Adaptive, self-tuning memory allocation” on page 77
- “Administration changes summary” on page 123
- “Configuration parameters default value changes” on page 117

Related reference:

- “sheapthres - Sort heap threshold configuration parameter” in *Performance Guide*
- “sheapthres_shr - Sort heap threshold for shared sorts configuration parameter” in *Performance Guide*

Table space-level point-in-time rollforward operations restrictions

Point in time rollforward of a table space is available only from DB2 Version 9 clients. You should migrate any clients running an earlier version of the database product to Version 9 in order to roll a table space forward to a point in time.

Related concepts:

- “Rolling forward changes in a table space” in *Data Recovery and High Availability Guide and Reference*

Related tasks:

- “Using rollforward” in *Data Recovery and High Availability Guide and Reference*

TSM management class no longer used for filtering when restoring or retrieving logs

Management class is a Tivoli Storage Manager (TSM) concept that helps with the management of objects according to defined storage policies. When a backup image, a load copy image or a log file is written to TSM, a particular management class is associated with that object. After a log file is written to or a backup image is stored, the management class can be changed through TSM. Prior to DB2 Version 9.1, restore and log retrieval could search for objects based on a management class, if it was specified. Because the management class can change, filtering based on management class could produce incorrect results. Consequently, management class is no longer used as a basis for filtering.

Related concepts:

- “Administration changes summary” on page 123

Related reference:

- “Tivoli Storage Manager” in *Data Recovery and High Availability Guide and Reference*

Write-to-table event monitor changes

In a partitioned database environment, a write-to-table event monitor will only be active on database partitions where the table space containing the event monitor table exists. When the target table space for an active event monitor does not exist on a particular database partition, the event monitor will be deactivated on that database partition, and an error is written to the db2diag.log file.

In earlier versions of DB2, the event monitor would be active and would appear as an active event monitor process on these database partitions but would not write any data.

Related concepts:

- “Event monitor table management” in *System Monitor Guide and Reference*

Application development changes

Application development changes summary

For changes to existing application development-related functionality in DB2 Version 9.1, follow the links in the “Related concepts” section.

Related concepts:

- “Application and routine feature support changes” on page 132
- “Changed parameters and output for the db2batch command” on page 133
- “Cursor blocking is disabled in SQL procedures” on page 134
- “DB2 Embedded Application Server updated” on page 134
- “IBM Software Development Kit (SDK) for Java 5.x support added (AIX, Linux, and Windows)” on page 135
- “Removal of mutating table conflicts when calling procedures from triggers” on page 135
- “Reverse scans enabled by default for indexes, primary keys and unique keys” on page 136
- “Sample build scripts updated” on page 136

Application and routine feature support changes

The removal of support for most 32-bit database instances has resulted in changes in support for application and routines, as described below:

Client application connection support changes

- Client applications using DB2 Version 6 or Version 7 client instances cannot connect to DB2 Version 9 database servers.

Client application environment variable changes

- There are new environment variable values.

External routine support changes

- 32-bit unfenced routines (stored procedures and user-defined functions) created in DB2 Universal Database Version 8 will no longer work on 64-bit DB2 database servers in the AIX, HP, SUN, Linux on POWER, Linux for AMD64 and Intel EM64T, and Linux on zSeries environments. Migrating these routines to DB2 Version 9 requires that you rebuild them on the target 64-bit database server.

SQL procedure support changes

- SQL procedures that you created for 32-bit instances of DB2 Universal Database Version 8 prior to V8.2 (V8.1 FP7) will not run on 64-bit instances of DB2 Version 9. To successfully migrate these SQL procedures to DB2 Version 9, you must drop and recreate the SQL procedures using the target 64-bit database server.
- SQL procedures created for 32-bit instances of DB2 Universal Database Version 7 or Version 8 with any FixPak will continue to work on the supported 32-bit instances of DB2 Version 9. However, it is recommended that such procedures be recreated to take advantage of significant performance improvements introduced in later releases of DB2.

JVM support changes

- Only a 64-bit JVM is provided with 64-bit DB2 database servers.
- A 32-bit JVM is provided only for the Linux x86 and Windows on x86 operating systems.
- Java external routines require a 32-bit JVM for 32-bit DB2 database servers and a 64-bit JVM for 64-bit DB2 database servers.

Related concepts:

- “Application and routine migration changes” on page 106
- “Application development changes summary” on page 132
- “Hardware and operating system support” on page 17
- “Migration overview for database applications and routines” in *Migration Guide*

Related tasks:

- “Migrating database applications” in *Migration Guide*
- “Migrating routines” in *Migration Guide*

Changed parameters and output for the db2batch command

The db2batch command has undergone changes that affect some of its options, option parameters, and output. The db2batch command is a benchmark tool that reads SQL statements from either a flat file or standard input, dynamically prepares and describes the statements, and returns an answer set. The following changes have been made:

db2batch runs only in CLI mode

Embedded dynamic SQL mode, previously the default mode for db2batch, is no longer supported. Starting in DB2 Version 9.1, db2batch runs only in CLI mode. Specifying the -cli option (including the optional cache size argument) will not cause errors if specified, but this option is not required.

The db2batch.bnd file is not shipped

The db2batch command no longer requires the db2batch.bnd file and scripts should not attempt to issue a BIND or REBIND command using this file.

-iso isolation level option

The -iso option allows you to specify an isolation level. By default, db2batch will run at the RR isolation level. The TxnIsolation configuration keyword in the db2cli.ini file will have no effect on db2batch. If an isolation level other than RR is required, the -iso parameter must be specified.

Output changes

The output provided by the db2batch command has been improved and includes additional information, timestamps, better formatting, and clearer messages.

The -p option is unsupported

This option was used previously to allow only SELECT statements and is no longer supported.

New data types supported

In addition to data types supported in DB2 UDB Version 8, db2batch now supports the following data types: CLOB, GRAPHIC, VARGRAPHIC, LONGVARGRAPHIC, DBCLOB, BLOB, and XML.

Native XML data store support

Please refer to the related concepts below.

Other deprecated option parameters and new options

Some option parameters have been deprecated, and several other new options are available. Please refer to the Command Reference for more details.

Related concepts:

- “Application development changes summary” on page 132
- “db2batch command changes for native XML data store” on page 41

Related reference:

- “db2batch - Benchmark tool command” in *Command Reference*

Cursor blocking is disabled in SQL procedures

In Version 8, cursor blocking can be used in SQL procedures using the BLOCKING bind option. The BLOCKING bind option is specified by setting the DB2_SQLROUTINE_PREPOPTS registry variable, or by calling the SET_ROUTINE_OPTS system stored procedure before issuing the CREATE PROCEDURE statement.

In Version 9.1, cursor blocking is disabled in SQL procedures, regardless of the value that you specify for the BLOCKING bind option. The data will be retrieved one row at a time. This is a result of architecture changes to improve the performance of SQL procedures through a tighter integration between the SQL procedure interpreter and the SQL interpreter.

This change applies to FETCH statements as well as FETCH statements that are implicitly contained in FOR loops.

Related concepts:

- “Application development changes summary” on page 132
- “Binding” in *Administration Guide: Planning*

Related tasks:

- “Customizing precompile and bind options for SQL procedures” in *SQL Guide*

Related reference:

- “CLOSE statement” in *SQL Reference, Volume 2*
- “DECLARE CURSOR statement” in *SQL Reference, Volume 2*
- “FETCH statement” in *SQL Reference, Volume 2*
- “OPEN statement” in *SQL Reference, Volume 2*

DB2 Embedded Application Server updated

The DB2 Embedded Application Server (also referred to in DB2 UDB Version 8 as the *application server for DB2 UDB*) enables you to run the Web applications supplied with DB2 Version 9.1 without needing to purchase a separate application server.

The Web applications supplied with DB2 Version 9.1 are:

- DB2 Web Tools, for Web-based database administration
- DB2WebServices, an application that automates the deployment of .NET Web services from Microsoft Visual Studio to the DB2 Embedded Application Server

The XML Metadata Repository (XMR) application is no longer supplied with the DB2 Embedded Application Server. If you used the XMR application in V8, you must uninstall XMR and find a replacement product. WebSphere offers suitable replacement products.

If you used the DB2 Embedded Application Server in a prior release, you must upgrade it to the new version. Enhancements in DB2 Version 9.1 include:

- An option to install the application server on a separate server from the DB2 server. This enhancement applies to the DB2 Web Tools application.
- Simplified setup.

Related concepts:

- “DB2 Embedded Application Server overview and setup” in *Installation and Configuration Supplement*
- “DB2 Web Tools overview and setup” in *Installation and Configuration Supplement*
- “DB2WebServices application overview and setup” in *Installation and Configuration Supplement*
- “Application development changes summary” on page 132

IBM Software Development Kit (SDK) for Java 5.x support added (AIX, Linux, and Windows)

DB2 Version 9 now supports IBM Software Development Kit (SDK) for Java 5.x on the following operating system platforms: AIX 5, Linux on x86, Linux on AMD64/EM64T, Linux on zSeries, Linux on POWER, Windows x86 and Windows x64.

The IBM SDK is automatically installed on the server. If the client tools are installed, the IBM SDK is also installed on the client. If you are using the JDBC drivers with your own applications, you need to ensure the IBM SDK is installed.

Related concepts:

- “Application development changes summary” on page 132

Related reference:

- “IBM Software Development Kit for Java levels for DB2 products” in *Quick Beginnings for DB2 Servers*

Removal of mutating table conflicts when calling procedures from triggers

Before DB2 Version 9.1, CALL statement in triggers resulted in an error if an operation conflicted with other uses of the table by either the application, or a routine invoked directly or indirectly from that application. In Version 9.1 this restriction has been lifted. The new DB2_RESOLVE_CALL_CONFLICT registry variable is used to ensure that all modifications to tables are completed in compliance with the SQL Standard rules for triggers before executing the CALL statement. This variable is turned on by default.

In some cases, using this variable can cause triggers using the CALL statement to perform slightly slower. To revert to Version 8.2 behavior, set the DB2_RESOLVE_CALL_CONFLICT registry variable to N.

Related concepts:

- “Application development changes summary” on page 132

Related reference:

- “Miscellaneous variables” in *Performance Guide*

Reverse scans enabled by default for indexes, primary keys and unique keys

In DB2 Version 9, all new primary keys, unique keys and indexes (except extended index) will ALLOW REVERSE SCANS by default. Consequently, the access plan may change and query execution times may improve because the optimizer may be able to use the reverse index scan in some SQL statements. This feature also leads to improved index maintenance. In prior DB2 versions, some DB2 users created one forward scan index and one reverse scan index to speed up the application. Unfortunately, this requires the maintenance of two indexes. Now that reverse scans is enabled by default in DB2 9.1, the two indexes can be replaced with a single one that is enabled for reverse scans.

Attention: If you create two indexes on the same table, one specifying ASC and the other DESC, and if you do not specify the DISALLOW REVERSE SCANS option in the CREATE INDEX statement, the two indexes will default to ALLOW REVERSE SCANS. As a result, the latter index will not be created and DB2 will issue a duplicate index warning message.

Related concepts:

- “Options on the CREATE INDEX statement” in *Administration Guide: Implementation*
- “Application development changes summary” on page 132

Related reference:

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

Sample build scripts updated

On 64-bit platforms, the sample build files for building database applications and routines have been updated to build 64-bit applications and routines by default. The 32-bit options required for building 32-bit applications and routines are provided in the build files, but are commented out. To successfully build 32-bit applications and routines, the build files must be manually updated to make use of the 32-bit options.

On Windows and Linux x86 platforms, the build scripts will still build 32-bit applications and routines by default.

The sample build scripts for each application programming interface (API) or programming language are located with the sample programs for the API or programming language. These files are located in the subdirectories of the following product directory:

- sql1lib/samples (UNIX)
- sql1lib\samples (Windows).

Note: Only the build scripts specific to the current operating system are installed.
For other build scripts, refer to the DB2 Information Center.

Related concepts:

- “Sample files” in *Samples Topics*
- “Application development changes summary” on page 132
- “Application development enhancements summary” on page 43

Chapter 17. Deprecated functionality

Deprecated functionality summary

Deprecated functionality means that a specific function or feature is supported but is no longer recommended and might be removed in a future release.

For information about deprecated functionality in DB2 Version 9.1, follow the links in the "Related concepts" section. For additional information about deprecated administration functionality and incompatibilities, see "Deprecated and discontinued features" in *Administration Guide: Planning*. This topic provides recommendations for replacement functionality, where appropriate.

Related concepts:

- "ADD PARTITIONING KEY clause of the ALTER TABLE statement is deprecated" on page 143
- "Check pending table state is replaced and iCheckPending parameter is deprecated" on page 139
- "CLISchema CLI keyword no longer supported" on page 140
- "COLNAMES column in SYSCAT.INDEXES is deprecated" on page 140
- "db2secv82 command is deprecated" on page 141
- "db2ilist command options are deprecated (Linux and UNIX)" on page 140
- "Database logging using raw devices is deprecated" on page 142
- "DB2 JDBC Type 2 driver is deprecated" on page 145
- "DB2LINUXAIO registry variable is deprecated" on page 141
- "Distribution key terminology change" on page 27
- "DROP PARTITIONING KEY clause of the ALTER TABLE statement is deprecated" on page 144
- "External routines now require an explicit entry point specification" on page 141
- "Some SQL administrative routines have been replaced" on page 143
- "Type 1 indexes are deprecated" on page 144

Check pending table state is replaced and iCheckPending parameter is deprecated

In Version 8, a table with constraints that have not yet been verified is in the check pending state. In Version 9.1, the set integrity pending state replaces the check pending state. These are equivalent states.

In Version 8, the iCheckPending parameter is used as an input parameter of the db2Load API to specify whether a table should be put into the check pending state. In Version 9.1, this parameter is deprecated; you should use the iSetIntegrityPending parameter as a replacement. Table 5 on page 140 shows a comparison between the iCheckPending and the iSetIntegrityPending parameter values:

Table 5. Comparison between the *iCheckPending* and *iSetIntegrityPending* parameter values

| <i>iCheckPending</i> parameter value | <i>iSetIntegrityPending</i> parameter value |
|--------------------------------------|---------------------------------------------|
| SQLU_CHECK_PENDING_CASCADE_IMMEDIATE | SQLU_SI_PENDING_CASCADE_IMMEDIATE |
| SQLU_CHECK_PENDING_CASCADE_DEFERRED | SQLU_SI_PENDING_CASCADE_DEFERRED |

Related concepts:

- “Online processing of the SET INTEGRITY statement” on page 91

Related reference:

- “SET INTEGRITY statement” in *SQL Reference, Volume 2*

CLISchema CLI keyword no longer supported

For DB2 Version 9 clients connecting to DB2 for Linux, UNIX, and Windows DB2 database servers, the CLISchema keyword is deprecated. For DB2 Version 9 clients connecting to DB2 for z/OS database servers, the CLISchema keyword is dropped.

Setting the CLISchema keyword improved performance, primarily for client applications connecting to DB2 for z/OS. For information about configurations that improve the performance of client applications connecting to DB2 for z/OS, see “CLI/ODBC application performance tuning” in *DB2 Connect User’s Guide*. One keyword that is similar to CLISchema is SysSchema. For information about the SysSchema keyword, see “SysSchema CLI/ODBC Configuration Keyword” in *Call Level Interface Guide and Reference, Volume 1*.

Related concepts:

- “Deprecated functionality summary” on page 139

COLNAMES column in SYSCAT.INDEXES is deprecated

In DB2 Version 9.1, the COLNAMES column in SYSCAT.INDEXES is deprecated and will be removed in a future release. This column contains valid information only if each column name is less than 30 bytes and if there are fewer than 16 columns in the index. Either a blank or a NULL value is returned if any column name is greater than 30 bytes or if there are more than 16 columns.

Related concepts:

- “Deprecated functionality summary” on page 139

Related reference:

- “SYSCAT.INDEXCOLUSE catalog view” in *SQL Reference, Volume 1*
- “SYSCAT.INDEXES catalog view” in *SQL Reference, Volume 1*

db2ilist command options are deprecated (Linux and UNIX)

In DB2 Version 8, you could use the **db2ilist** command to list all of the DB2 instances that were available on a system, including GA and Fixpack instances. This command now only lists the instances related to the current installation path. In addition, DB2 V9.1 only supports one type of DB2 instance on each UNIX or Linux platform. Therefore, the following options are deprecated:

- w Lists the 31-, 32-, or 64-bit instances.
- a Lists information including the DB2 install path associated with an instance, as well as its bit width (32 or 64).
- p Lists the DB2 install path that an instance is running from.

These options were only valid on the AIX 5L™, HP-UX, Linux and Solaris operating systems.

Related concepts:

- “Coexistence of multiple DB2 versions and fix packs enhancements (Linux and UNIX)” on page 69

Related reference:

- “db2ilist - List instances command” in *Command Reference*

DB2LINUXAIO registry variable is deprecated

In DB2 Version 9.1, the DB2LINUXAIO variable is deprecated. It might become obsolete in a future release, as it might no longer be necessary.

Related concepts:

- “Deprecated functionality summary” on page 139

db2secv82 command is deprecated

The **db2secv82** command, which you can use to set the permissions for DB2 objects (for example, files, directories, network shares, registry keys, and services) is deprecated. You can use the **db2extsec** command as a replacement.

Related concepts:

- “Deprecated functionality summary” on page 139

Related reference:

- “db2extsec - Set permissions for DB2 objects command” in *Command Reference*

External routines now require an explicit entry point specification

Support for default function entry points in external routine libraries is deprecated. This support was only available for 32-bit AIX and Windows database servers. You should no longer rely on the DB2 database manager to resolve and load the function specified by a default entry point; instead, you should specify an explicit routine library entry point for each routine.

In the CREATE statement for external routines, the EXTERNAL NAME clause allows you to specify the name of the file containing the external routine library, class, or assembly and the entry point for the function associated with the routine. The entry point specification is *!proc-id* for a procedure and *!func-id* for a function.

If you are creating a new external routine, you should specify a *!proc-id* value to ensure that the database manager always locates and loads the correct library for the routine.

If you have an existing external routine definition that specifies the EXTERNAL NAME clause without a value, you should modify the definition to provide an explicit entry point value. You can do this by using the ALTER PROCEDURE or ALTER FUNCTION statement. For example, consider an external procedure named myproc that references a sub-routine within library mylib that resolves to a default entry point, which is a function named func1:

```
CREATE PROCEDURE myproc
LANGUAGE C
PARAMETER STYLE SQL
EXTERNAL NAME 'mylib'
FENCED
```

You can use the following SQL statement to specify the entry point explicitly:

```
ALTER PROCEDURE myproc() EXTERNAL NAME 'mylib!func1'
```

The definition for the routine that results in the database is as follows:

```
CREATE PROCEDURE myproc
LANGUAGE C
PARAMETER STYLE SQL
EXTERNAL NAME 'mylib!func1'
FENCED
```

When the procedure myproc is called, the library mylib is loaded, and the function at explicit entry point func1 is resolved and loaded.

If you are not sure what the explicit entry point for a routine should be because you no longer have the source code from which the library was built, you no longer have the export file, or the routine was provided by an independent vendor, you can use some AIX and Windows operating system commands and tools to determine the entry points defined within a library. With a list of entry points for the library, it should be easier to determine which entry point to specify for the routine.

Related concepts:

- “Deprecated functionality summary” on page 139
- “External routines” in *SQL Guide*

Related tasks:

- “Creating external routines” in *SQL Guide*

Related reference:

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

Database logging using raw devices is deprecated

As of DB2 Version 9.1, the use of raw devices for database logging is deprecated. As an alternative to using raw logs, you can use either direct input/output (DIO) or concurrent input/output (CIO).

Related concepts:

- “Backup, logging, and recovery enhancements summary” on page 59
- “Deprecated functionality summary” on page 139

Related tasks:

- “Configuring database logging options” in *Data Recovery and High Availability Guide and Reference*

Related reference:

- “Configuration parameters for database logging” in *Data Recovery and High Availability Guide and Reference*

Some SQL administrative routines have been replaced

To provide expanded support, some of the DB2 UDB for Linux, UNIX, and Windows administrative routines have been replaced with more comprehensive administrative routines or administrative views in DB2 Version 9. In most cases, these new table functions and administrative views return additional information. The administrative views will always be based on the most current version of the table functions and, therefore, allow for more application portability.

You should modify applications that use the Version 8 table functions to use the new functions or administrative views. The new table functions have the same base names as the original functions but are suffixed with `_Vxx` for the version of the product in which they were added (for example, `_V91`). Since the columns used in the administrative views might vary from one release to the next (that is, some might be added or deleted), you should select specific columns from the administrative views, or describe the result set if a `SELECT *` statement is used by an application.

Related concepts:

- “Deprecated functionality summary” on page 139

Related reference:

- “Deprecated SQL administrative routines and their replacement routines or views” in *Administrative SQL Routines and Views*
- “Supported administrative SQL routines and views” in *Administrative SQL Routines and Views*

ADD PARTITIONING KEY clause of the ALTER TABLE statement is deprecated

As part of the new table partitioning functionality, the clause for adding a distribution key (known as a partitioning key in previous releases) using the `ALTER TABLE` statement has changed from `ADD PARTITIONING KEY` to `ADD DISTRIBUTE BY HASH`. The `ADD PARTITIONING KEY` clause is deprecated; it is supported for backwards compatibility. There is no restriction on using this clause with the new `PARTITION BY RANGE` clause.

The `DISTRIBUTE BY REPLICATION` clause, as per previous releases, is only supported with materialized query tables (MQTs). An error is returned if you specify it for any other type of table.

Related concepts:

- “Distribution keys” in *Administration Guide: Planning*
- “Table partitioning keys” in *Administration Guide: Planning*

- “Deprecated functionality summary” on page 139

Related tasks:

- “Changing distribution keys” in *Administration Guide: Implementation*
- “Altering a table” in *Administration Guide: Implementation*

Related reference:

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

DROP PARTITIONING KEY clause of the ALTER TABLE statement is deprecated

As part of the new table partitioning functionality, the clause for dropping a distribution key (known as a partitioning key in previous releases) using the ALTER TABLE statement has changed from DROP PARTITIONING KEY to DROP DISTRIBUTION. The DROP PARTITIONING KEY clause is deprecated; it is supported for backwards compatibility. There is no restriction on using this old clause with the new PARTITION BY RANGE clause.

The DISTRIBUTE BY REPLICATION clause, as per previous releases, is only supported with materialized query tables (MQTs). An error is returned if you specify it for any other type of table.

Related concepts:

- “Keys” in *SQL Reference, Volume 1*
- “Distribution keys” in *Administration Guide: Planning*
- “Table partitioning keys” in *Administration Guide: Planning*
- “Deprecated functionality summary” on page 139

Related tasks:

- “Altering a table” in *Administration Guide: Implementation*
- “Changing distribution keys” in *Administration Guide: Implementation*

Related reference:

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

Type 1 indexes are deprecated

In DB2 Version 9.1, Type 1 indexes are deprecated. Type 1 indexes are still supported in DB2 Version 9.1, but are no longer recommended. Type 2 indexes were introduced in DB2 UDB Version 8.1, and all new indexes created since then are Type 2. You can migrate manually from Type 1 to Type 2 indexes during an index reorganization.

Related concepts:

- “Index structure” in *Performance Guide*
- “Deprecated functionality summary” on page 139

Related reference:

- “REORG INDEXES/TABLE command” in *Command Reference*

DB2 JDBC Type 2 driver is deprecated

The DB2 JDBC Type 2 driver lets Java applications make calls to DB2 through JDBC. The DB2 JDBC Type 2 driver was deprecated in version 8.2, and remains deprecated in version 9.1. Support for the driver will be removed in a future release.

Use the IBM DB2 Driver for JDBC and SQLJ instead. For information on how to make the IBM DB2 Driver for JDBC and SQLJ the default driver for routines and applications, follow the related links.

Related concepts:

- “Specification of a driver for Java routines” in *SQL Guide*

Related tasks:

- “Installing the IBM DB2 Driver for JDBC and SQLJ” in *Developing Java Applications*
- “Migrating Java applications that use DB2 JDBC Type 2 or 3 driver” in *Migration Guide*
- “Migrating Java routines” in *Migration Guide*

Related reference:

- “Miscellaneous variables” in *Performance Guide*

The -w option of the db2icrt, db2ilist, and db2iupdt commands is ignored

The **-w** option, which specifies the instance width in bits, is deprecated and is ignored in DB2 Version 9.1 because the instance bit size is now determined by the operating system where the product installation is done. This change affects the following commands:

- **db2icrt**
- **db2iupdt**
- **db2ilist**

If you specify this option, it has no effect. This option is only valid on AIX 5L, HP-UX, Linux, and Solaris operating systems.

Related concepts:

- “Changes in existing functionality summary” on page 105
- “db2ilist command options are deprecated (Linux and UNIX)” on page 140

Related reference:

- “db2icrt - Create instance command” in *Command Reference*
- “db2ilist - List instances command” in *Command Reference*
- “db2iupdt - Update instances command” in *Command Reference*

Chapter 18. Discontinued functionality

Discontinued functionality summary

For information about discontinued functionality in DB2 Version 9.1, follow the links in the “Related concepts” section.

Related concepts:

- “Alternate FixPak images are discontinued (UNIX)” on page 147
- “Audio, Image, and Video (AIV) Extenders are no longer supported” on page 148
- “Autoloader utility (db2atld) is no longer supported” on page 148
- “db2profc and db2profp utilities are discontinued” on page 150
- “db2reg2large utility for converting DMS table space size is discontinued” on page 150
- “Data Links Manager no longer supported” on page 148
- “DB2 Administration Tools are no longer supported on some platforms” on page 152
- “DB2_SCATTERED_IO registry variable is discontinued (Linux)” on page 151
- “Desktop icon and folder making utilities are no longer supported (Linux)” on page 151
- “Extended Storage option for buffer pools is discontinued” on page 151
- “NetBIOS and SNA communication protocols are no longer supported” on page 152
- “Text Extender is no longer supported” on page 153
- “The Data Warehouse Center and the Information Catalog Center are no longer included” on page 150
- “Type 3 JDBC support is discontinued” on page 153
- “Vendor load API (sqluvtld) is discontinued” on page 153
- “VSE and VM objects are no longer supported in the DB2 Control Center” on page 154
- “Fast communications manager (FCM) no longer uses virtual interface (VI) architecture” on page 152
- “Version 9 incompatibilities with previous releases and changed behaviors” in *Administration Guide: Planning*

Alternate FixPak images are discontinued (UNIX)

The Alternate FixPack images, which allowed multiple levels of DB2 Enterprise Server Edition in Version 8 to coexist are no longer provided. Instead, this capability is replaced with the ability to install DB2 Version 9 images to multiple locations that can be serviced independently of each other.

Related concepts:

- “Coexistence of multiple DB2 versions and fix packs enhancements (Linux and UNIX)” on page 69

Related tasks:

- “Installing DB2 servers (Linux and UNIX)” in *Quick Beginnings for DB2 Servers*

Related reference:

- “Multiple DB2 copies roadmap” in *Administration Guide: Implementation*

Autoloader utility (db2atld) is no longer supported

The Autoloader utility (db2atld) is no longer supported.

You should use the load utility for distributing and loading data in partitioned database environments.

Related concepts:

- “Load overview” in *Data Movement Utilities Guide and Reference*
- “Optimizing load performance” in *Data Movement Utilities Guide and Reference*
- “Load in a partitioned database environment - overview” in *Data Movement Utilities Guide and Reference*
- “Discontinued functionality summary” on page 147

Related reference:

- “Load configuration options for partitioned database environments” in *Data Movement Utilities Guide and Reference*

Audio, Image, and Video (AIV) Extenders are no longer supported

AIV Extenders are no longer supported. You might consider implementing your own extensions similar to the AIV Extenders by using DB2 user-defined functions and third-party software. For an example of a user-defined function to manage and process images stored in a DB2 database system, see “A DB2 UDB still image extender” at www.ibm.com/developerworks/db2/library/techarticle/dm-0504stolze/.

Related concepts:

- “User-defined functions (UDFs) or methods” in *Administration Guide: Implementation*
- “Discontinued functionality summary” on page 147

Data Links Manager no longer supported

DB2 Data Links Manager is not supported in this release. Also, the following components of a Data Links server are not supported:

- Data Links File Manager (DLFM)
- Data Links Filesystem Filter (DLFF) controlling a Data Links File System (DLFS)
- DB2 Logging Manager

There is no support for Data Links Manager between DB2 Version 9.1 and DB2 Universal Database Version 8. There is no support for Data Links Manager between a DB2 Version 9 client and a DB2 Universal Database Version 8 server, or a DB2 Universal Database Version 8 server and a DB2 Version 9 client.

For the DB2 Version 9 client, the command `SQLGetTypeInfo` will not report `DATALINKS` as a supported data type for the DB2 Version 9.1 for Linux, UNIX, and Windows server. Existing applications that access other versions of servers that might still support Data Links Manager, such as iSeries and z/OS, will still work.

In addition, you cannot create any new database objects that reference `DATALINK` columns. Any attempts to use the `DATALINK` data type will result in the following error: `SQL0104N: An unexpected token "<token>" was found following "<text>".` This restriction includes any objects that are created or modified by the following SQL statements:

- `CREATE/ALTER TABLE`
- `ALTER TABLE`
- `ALTER TYPE`
- `CREATE DISTINCT TYPE`
- `CREATE FUNCTION MAPPING`
- `CREATE FUNCTION`
- `DELETE`
- `DROP`
- `INSERT`
- `SET INTEGRITY`
- `UPDATE`
- `RECONCILE`

Note: The `db2_recon_aid` utility that you use to run the `RECONCILE` statement on multiple tables is no longer available.

Any database that has a Data Links server will be affected in the following ways:

- If you try to migrate an instance that contains Data Links Manager, the migration will fail with error `DBI1139E`. Instance migration must succeed before you can migrate a database. Database migration modifies the data on the disk.
- If you issue the `RESTORE` command and a `datalinks.cfg` file exists in the image, which indicates that Data Links Manager is installed, then the `RESTORE` command will fail with `SQL20158N` before implicit database migration.
- During database migration, the database manager configuration file is updated so that `DATALINKS` is set to `NO`.
- Database migration will fail with `SQL20158N` if the database contains `DATALINK` columns, including columns that are based on the `DATALINK` data type, which exists in functions, methods, distinct types, or structured types. You must drop all `DATALINK` data types before migrating the database.
- If you use DB2 Net Search Extender and you have Data Links Manager installed, you must drop all `DATALINK` features before attempting to migrate.

The following database manager configuration parameters will no longer be displayed:

- Data Links access token expiry interval - (`DL_EXPINT`) = 60
- Data Links write token initial expiry interval - (`DL_WT_IEXPINT`) = 60
- Data Links number of copies - (`DL_NUM_COPIES`) = 1
- Data Links time after drop - (`DL_TIME_DROP`) = 1
- Data Links token in uppercase - (`DL_UPPER`) = `NO`
- Data Links token algorithm - (`DL_TOKEN`) = `MAC0`

Related concepts:

- “Discontinued functionality summary” on page 147

Related tasks:

- “Migrating DB2 Data Links Manager environments” in *Migration Guide*

The Data Warehouse Center and the Information Catalog Center are no longer included

The Data Warehouse Center and the Information Catalog Center are not included with Version 9.1. Also, DB2 Warehouse Manager Standard Edition is not available in Version 9.1.

Related concepts:

- “Discontinued functionality summary” on page 147

db2reg2large utility for converting DMS table space size is discontinued

The db2reg2large utility, which you could previously use for converting DMS table spaces defined as REGULAR to DMS table spaces defined as LARGE, is discontinued in DB2 Version 9.1. It is replaced with the new CONVERT TO LARGE option of the ALTER TABLESPACE SQL statement.

Related concepts:

- “Discontinued functionality summary” on page 147
- “DMS table spaces” in *Administration Guide: Planning*

Related reference:

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

db2profc and db2profp utilities are discontinued

The DB2 JDBC Type 2 Driver originally used the name db2profc for the SQLJ profile customizer command, and the name db2profp for the SQLJ profile printer command. For the IBM DB2 Driver for JDBC and SQLJ, the SQLJ profile customizer command is named db2sqljcustomize, and the SQLJ profile printer command is named db2sqljprint.

In previous releases, db2profc was accepted as an alternative name for db2sqljcustomize, and db2profp was accepted as an alternative name for db2sqljprint. In version 9, these alternative names are no longer accepted.

Related reference:

- “db2sqljcustomize - SQLJ profile customizer” in *Developing Java Applications*
- “db2sqljprint - SQLJ profile printer” in *Developing Java Applications*

DB2_SCATTERED_IO registry variable is discontinued (Linux)

In DB2 Version 9.1, the DB2_SCATTERED_IO registry variable, which enabled a DB2 system to use `readv()` to read from disk on Linux, is no longer needed and is discontinued. In DB2 Version 9.1, the default behavior is to always use `readv()`, when appropriate, to take advantage of associated performance improvements.

Related concepts:

- “Discontinued functionality summary” on page 147

Related reference:

- “Performance variables” in *Performance Guide*

Desktop icon and folder making utilities are no longer supported (Linux)

Unlike previous versions, DB2 Version 9.1 does not include a set of utilities for creating DB2 desktop folders and icons for launching commonly used DB2 tools on the Gnome and KDE desktops on supported Intel-based Linux operating systems.

Related concepts:

- “Discontinued functionality summary” on page 147

Extended Storage option for buffer pools is discontinued

On 32-bit operating systems, the amount of virtual addressable memory is usually limited to between 2 GB and 4 GB. In previous versions, if DB2 was running on a computer that used a 32-bit operating system and that had more real addressable memory than this amount, you could use the Extended Storage option for buffer pools to take advantage of the extra memory for better performance. In Version 9.1, the Extended Storage option for buffer pools is discontinued, and two configuration parameters for extended storage (ESTORE_SEG_SZ and NUM_ESTORE_SEGS) are deprecated.

If you use a Windows 32-bit operating system and need to use more memory, you should consider instead using a Windows 64-bit operating system or using the Address Windowing Extensions (AWE) feature, which allows a 32-bit system to use more memory, via the DB2_AWE performance variable.

Related concepts:

- “Secondary buffer pools in extended memory on 32-bit platforms” in *Performance Guide*
- “Discontinued functionality summary” on page 147

Related reference:

- “Performance variables” in *Performance Guide*

Fast communications manager (FCM) no longer uses virtual interface (VI) architecture

Fast communications manager (FCM) has been re-architected and no longer uses virtual interface (VI) architecture.

The following three registry variables used to control the use of virtual interface (VI) architecture within the product are discontinued:

- DB2_VI_ENABLE
- DB2_VI_DEVICE
- DB2_VI_VIPL

Related concepts:

- “Fast communications manager (FCM) communications” in *Administration Guide: Implementation*
- “Fast communications manager (FCM) enhancements” on page 82

NetBIOS and SNA communication protocols are no longer supported

NetBIOS is no longer supported. SNA, including its APIs APPC, APPN, and CPI-C, is also no longer supported. If you use these protocols, you must recatalog your nodes and databases using a supported protocol such as TCP/IP.

Related concepts:

- “Migration overview for DB2 clients” in *Migration Guide*
- “Migration overview for DB2 servers” in *Migration Guide*
- “Discontinued functionality summary” on page 147

Related tasks:

- “Recataloging nodes and databases that use NetBIOS and SNA protocols” in *Migration Guide*

DB2 Administration Tools are no longer supported on some platforms

In previous releases, the DB2 Administration Tools—for example, the Control Center— were supported on all platforms. In Version 9.1, the DB2 Administration Tools are supported only on the Windows x86, Windows x64 (for AMD64 and Intel EM64T), 32-bit Linux x86, and Linux for AMD64 and Intel EM64T operating systems.

Related concepts:

- “Control Center overview” in *Administration Guide: Implementation*
- “Discontinued functionality summary” on page 147

Related tasks:

- “Finding service level information about the DB2 administration tools environment” in *Administration Guide: Implementation*
- “Setting startup and default options for the DB2 administration tools” in *Administration Guide: Implementation*

- “Shutting down server DB2 administration tools” in *Administration Guide: Implementation*

Text Extender is no longer supported

Text Extender is no longer supported. A replacement with all of the equivalent functionality is not available. However, there are other capabilities for efficiently searching text documents stored in columns (also called doing a full-text search), such as:

- DB2 Net Search Extender, a DB2-integrated search engine similar to Text Extender. For more information, see the DB2 Net Search Extender home page at www.ibm.com/software/data/db2/extenders/netsearch.
- WebSphere Information Integrator OmniFind™ Edition, an enterprise search solution for finding the most relevant information not only in relational databases but also across corporate or governmental public Web sites and a wide range of content repositories. For more information, see the WebSphere Information Integrator OmniFind Edition home page at www.ibm.com/software/data/integration/db2ii/editions_womnifind.html.

Related concepts:

- “Discontinued functionality summary” on page 147

Type 3 JDBC support is discontinued

The DB2 JDBC Type 3 driver is no longer supported.

Use the IBM DB2 Driver for JDBC and SQLJ instead. For information on how to make the IBM DB2 Driver for JDBC and SQLJ the default driver, follow the related links.

Related concepts:

- “Discontinued functionality summary” on page 147
- “JDBC and SQLJ enhancements” on page 52

Related tasks:

- “Migrating Java applications that use DB2 JDBC Type 2 or 3 driver” in *Migration Guide*
- “Installing the IBM DB2 Driver for JDBC and SQLJ” in *Developing Java Applications*

Vendor load API (sqluvtd) is discontinued

The vendor load API (sqluvtd) is no longer available. The load utility is the only supported bulk loader. You can run the load utility by using the db2Load API.

Related concepts:

- “Load overview” in *Data Movement Utilities Guide and Reference*
- “Loading data in a partitioned database environment - hints and tips” in *Data Movement Utilities Guide and Reference*
- “Discontinued functionality summary” on page 147

Related tasks:

- “Loading data” in *Data Movement Utilities Guide and Reference*
- “Loading data in a partitioned database environment” in *Data Movement Utilities Guide and Reference*

Related reference:

- “db2Load API - Load data into a table” in *Administrative API Reference*

VSE and VM objects are no longer supported in the DB2 Control Center

As of DB2 Version 9.1, you can no longer connect to or disconnect from VSE and VM databases from the DB2 Control Center. Also, when adding an instance, you can no longer select the VSE and VM operating systems. However, you can still display cataloged VSE and VM databases.

Related concepts:

- “Deprecated functionality summary” on page 139
- “Control Center overview” in *Administration Guide: Implementation*
- “DB2 Connect” in *DB2 Connect User’s Guide*

Appendix A. DB2 Database technical information

Overview of the DB2 technical information

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

- DB2 Information Center
 - Topics
 - Help for DB2 tools
 - Sample programs
 - Tutorials
- DB2 books
 - PDF files (downloadable)
 - PDF files (from the DB2 PDF CD)
 - printed books
- Command line help
 - Command help
 - Message help
- Sample programs

IBM periodically makes documentation updates available. If you access the online version on the DB2 Information Center at ibm.com[®], you do not need to install documentation updates because this version is kept up-to-date by IBM. If you have installed the DB2 Information Center, it is recommended that you install the documentation updates. Documentation updates allow you to update the information that you installed from the *DB2 Information Center CD* or downloaded from Passport Advantage as new information becomes available.

Note: The DB2 Information Center topics are updated more frequently than either the PDF or the hard-copy books. To get the most current information, install the documentation updates as they become available, or refer to the DB2 Information Center at ibm.com.

You can access additional DB2 technical information such as technotes, white papers, and Redbooks™ online at ibm.com. Access the DB2 Information Management software library site at <http://www.ibm.com/software/data/sw-library/>.

Documentation feedback

We value your feedback on the DB2 documentation. If you have suggestions for how we can improve the DB2 documentation, send an e-mail to db2docs@ca.ibm.com. The DB2 documentation team reads all of your feedback, but cannot respond to you directly. Provide specific examples wherever possible so that we can better understand your concerns. If you are providing feedback on a specific topic or help file, include the topic title and URL.

Do not use this e-mail address to contact DB2 Customer Support. If you have a DB2 technical issue that the documentation does not resolve, contact your local IBM service center for assistance.

Related concepts:

- “Features of the DB2 Information Center” in *Online DB2 Information Center*
- “Sample files” in *Samples Topics*

Related tasks:

- “Invoking command help from the command line processor” in *Command Reference*
- “Invoking message help from the command line processor” in *Command Reference*
- “Updating the DB2 Information Center installed on your computer or intranet server” on page 161

Related reference:

- “DB2 technical library in PDF format” on page 156

DB2 technical library in PDF format

The following tables describe the DB2 library available from the IBM Publications Center at www.ibm.com/shop/publications/order.

Although the tables identify books available in print, the books might not be available in your country or region.

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 or other DB2 products.

Table 6. DB2 technical information

| Name | Form Number | Available in print |
|----------------------------------------------------------------|--------------------|---------------------------|
| <i>Administration Guide: Implementation</i> | SC10-4221 | Yes |
| <i>Administration Guide: Planning</i> | SC10-4223 | Yes |
| <i>Administrative API Reference</i> | SC10-4231 | Yes |
| <i>Administrative SQL Routines and Views</i> | SC10-4293 | No |
| <i>Call Level Interface Guide and Reference, Volume 1</i> | SC10-4224 | Yes |
| <i>Call Level Interface Guide and Reference, Volume 2</i> | SC10-4225 | Yes |
| <i>Command Reference</i> | SC10-4226 | No |
| <i>Data Movement Utilities Guide and Reference</i> | SC10-4227 | Yes |
| <i>Data Recovery and High Availability Guide and Reference</i> | SC10-4228 | Yes |
| <i>Developing ADO.NET and OLE DB Applications</i> | SC10-4230 | Yes |
| <i>Developing Embedded SQL Applications</i> | SC10-4232 | Yes |
| <i>Developing SQL and External Routines</i> | SC10-4373 | No |

Table 6. DB2 technical information (continued)

| Name | Form Number | Available in print |
|----------------------------------------------------------------------------------------------------------------------------------------------------|-------------|--------------------|
| <i>Developing Java Applications</i> | SC10-4233 | Yes |
| <i>Developing Perl and PHP Applications</i> | SC10-4234 | No |
| <i>Getting Started with Database Application Development</i> | SC10-4252 | Yes |
| <i>Getting started with DB2 installation and administration on Linux and Windows</i> | GC10-4247 | Yes |
| <i>Message Reference Volume 1</i> | SC10-4238 | No |
| <i>Message Reference Volume 2</i> | SC10-4239 | No |
| <i>Migration Guide</i> | GC10-4237 | Yes |
| <i>Net Search Extender Administration and User's Guide</i> Note: HTML for this document is not installed from the HTML documentation CD. | SH12-6842 | Yes |
| <i>Performance Guide</i> | SC10-4222 | Yes |
| <i>Query Patroller Administration and User's Guide</i> | GC10-4241 | Yes |
| <i>Quick Beginnings for DB2 Clients</i> | GC10-4242 | No |
| <i>Quick Beginnings for DB2 Servers</i> | GC10-4246 | Yes |
| <i>Spatial Extender and Geodetic Data Management Feature User's Guide and Reference</i> | SC18-9749 | Yes |
| <i>SQL Guide</i> | SC10-4248 | Yes |
| <i>SQL Reference, Volume 1</i> | SC10-4249 | Yes |
| <i>SQL Reference, Volume 2</i> | SC10-4250 | Yes |
| <i>System Monitor Guide and Reference</i> | SC10-4251 | Yes |
| <i>Troubleshooting Guide</i> | GC10-4240 | No |
| <i>Visual Explain Tutorial</i> | SC10-4319 | No |
| <i>What's New</i> | SC10-4253 | Yes |
| <i>XML Extender Administration and Programming</i> | SC18-9750 | Yes |
| <i>XML Guide</i> | SC10-4254 | Yes |
| <i>XQuery Reference</i> | SC18-9796 | Yes |

Table 7. DB2 Connect-specific technical information

| Name | Form Number | Available in print |
|----------------------------------------------------------|-------------|--------------------|
| <i>DB2 Connect User's Guide</i> | SC10-4229 | Yes |
| <i>Quick Beginnings for DB2 Connect Personal Edition</i> | GC10-4244 | Yes |

Table 7. DB2 Connect-specific technical information (continued)

| Name | Form Number | Available in print |
|-------------------------------------------------|-------------|--------------------|
| <i>Quick Beginnings for DB2 Connect Servers</i> | GC10-4243 | Yes |

Table 8. WebSphere Information Integration technical information

| Name | Form Number | Available in print |
|---------------------------------------------------------------------------------------------------------|-------------|--------------------|
| <i>WebSphere Information Integration: Administration Guide for Federated Systems</i> | SC19-1001 | Yes |
| <i>WebSphere Information Integration: ASNCLP Program Reference for Replication and Event Publishing</i> | SC19-1000 | Yes |
| <i>WebSphere Information Integration: Configuration Guide for Federated Data Sources</i> | SC19-1034 | No |
| <i>WebSphere Information Integration: SQL Replication Guide and Reference</i> | SC19-1002 | Yes |

Note: The DB2 Release Notes provide additional information specific to your product's release and fix pack level. For more information, see the related links.

Related concepts:

- "Overview of the DB2 technical information" on page 155
- "About the Release Notes" in *Release notes*

Related tasks:

- "Ordering printed DB2 books" on page 158

Ordering printed DB2 books

If you require printed DB2 books, you can buy them online in many but not all countries or regions. You can always order printed DB2 books from your local IBM representative. Keep in mind that some softcopy books on the *DB2 PDF Documentation CD* are unavailable in print. For example, neither volume of the *DB2 Message Reference* is available as a printed book.

Printed versions of many of the DB2 books available on the DB2 PDF Documentation CD can be ordered for a fee from IBM. Depending on where you are placing your order from, you may be able to order books online, from the IBM Publications Center. If online ordering is not available in your country or region, you can always order printed DB2 books from your local IBM representative. Note that not all books on the DB2 PDF Documentation CD are available in print.

Note: The most up-to-date and complete DB2 documentation is maintained in the DB2 Information Center at <http://publib.boulder.ibm.com/infocenter/db2help/>.

Procedure:

To order printed DB2 books:

- To find out whether you can order printed DB2 books online in your country or region, check the IBM Publications Center at <http://www.ibm.com/shop/publications/order>. You must select a country, region, or language to access publication ordering information and then follow the ordering instructions for your location.
- To order printed DB2 books from your local IBM representative:
 - Locate the contact information for your local representative from one of the following Web sites:
 - The IBM directory of world wide contacts at www.ibm.com/planetwide
 - The IBM Publications Web site at <http://www.ibm.com/shop/publications/order>. You will need to select your country, region, or language to the access appropriate publications home page for your location. From this page, follow the "About this site" link.
 - When you call, specify that you want to order a DB2 publication.
 - Provide your representative with the titles and form numbers of the books that you want to order.

Related concepts:

- "Overview of the DB2 technical information" on page 155

Related reference:

- "DB2 technical library in PDF format" on page 156

Displaying SQL state help from the command line processor

DB2 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 command help from the command line processor" in *Command Reference*
- "Invoking message help from the command line processor" in *Command Reference*

Accessing different versions of the DB2 Information Center

For DB2 Version 9 topics, the DB2 Information Center URL is <http://publib.boulder.ibm.com/infocenter/db2luw/v9/>.

For DB2 Version 8 topics, go to the Version 8 Information Center URL at: <http://publib.boulder.ibm.com/infocenter/db2luw/v8/>.

Related tasks:

- “Setting up access to DB2 contextual help and documentation” in *Administration Guide: Implementation*

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. Clear the browser cache and then refresh the page to display the DB2 Information Center in your preferred language.

To display topics in your preferred language in a Firefox or Mozilla browser:

1. Select the **Tools** —> **Options** —> **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. Clear the browser cache and then refresh the page to display the DB2 Information Center in your preferred language.

On some browser and operating system combinations, you might have to also change the regional settings of your operating system to the locale and language of your choice.

Related concepts:

- “Overview of the DB2 technical information” on page 155

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

If you have a locally-installed DB2 Information Center, updated topics can be available for download. The 'Last updated' value found at the bottom of most topics indicates the current level for that topic.

To determine if there is an update available for the entire DB2 Information Center, look for the 'Last updated' value on the Information Center home page. Compare the value in your locally installed home page to the latest value which is available on the IBM hosted Information Center home page. If they are the same, you have the latest documentation level and no update is required. If they are not the same, you should update your locally-installed Information Center.

Updating your locally-installed DB2 Information Center requires that you:

1. Stop the DB2 Information Center on your computer, and restart the Information Center in stand-alone mode. Running the Information Center in stand-alone mode prevents other users on your network from accessing the Information Center, and allows you to download and apply updates.
2. Use the Update feature to determine if update packages are available from IBM. If update packages are available, use the Update feature to download the packages. (The Update feature is only available in stand-alone mode.)
3. Stop the stand-alone Information Center, and restart the DB2 Information Center service on your computer.

Procedure:

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

1. Stop the DB2 Information Center service.
 - On Windows, click **Start** → **Control Panel** → **Administrative Tools** → **Services**. Then right-click on **DB2 Information Center** service and select **Stop**.
 - On Linux, enter the following command:

```
/etc/init.d/db2icdv9 stop
```
2. Start the Information Center in stand-alone mode.
 - On Windows:
 - a. Open a command window.
 - b. Navigate to the path where the Information Center is installed. By default, the DB2 Information Center is installed in the C:\Program Files\IBM\DB2 Information Center\Version 9 directory.
 - c. Run the help_start.bat file using the fully qualified path for the DB2 Information Center:

```
<DB2 Information Center dir>\doc\bin\help_start.bat
```
 - On Linux:
 - a. Navigate to the path where the Information Center is installed. By default, the DB2 Information Center is installed in the /opt/ibm/db2ic/V9 directory.
 - b. Run the help_start.sh file using the fully qualified path for the DB2 Information Center:

```
<DB2 Information Center dir>/doc/bin/help_start
```

The systems default Web browser launches to display the stand-alone Information Center.

3. Click the Update button (🔄). On the right hand panel of the Information Center, click **Find Updates**. A list of updates for existing documentation displays.
4. To initiate the download process, check the selections you want to download, then click **Install Updates**.
5. After the download and installation process has completed, click **Finish**.
6. Stop the stand-alone Information Center.
 - On Windows, run the help_end.bat file using the fully qualified path for the DB2 Information Center:

```
<DB2 Information Center dir>\doc\bin\help_end.bat
```
 - On Linux, run the help_end.sh file using the fully qualified path for the DB2 Information Center:

```
<DB2 Information Center dir>/doc/bin/help_end
```
7. Restart the DB2 Information Center service.
 - On Windows, click **Start** → **Control Panel** → **Administrative Tools** → **Services**. Then right-click on **DB2 Information Center** service and select **Start**.
 - On Linux, enter the following command:

```
/etc/init.d/db2icdv9 start
```

The updated DB2 Information Center displays the new and updated topics.

Related concepts:

- “DB2 Information Center installation options” in *Quick Beginnings for DB2 Servers*

Related tasks:

- “Installing the DB2 Information Center using the DB2 Setup wizard (Linux)” in *Quick Beginnings for DB2 Servers*
- “Installing the DB2 Information Center using the DB2 Setup wizard (Windows)” in *Quick Beginnings for DB2 Servers*

DB2 Visual Explain tutorial

The DB2 Visual Explain tutorial helps you learn about analyzing, optimizing, and tuning SQL statements for better performance. Lessons provide step-by-step instructions.

Before you begin:

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

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

DB2 Visual Explain tutorial:

To view the tutorial, click on the title.

Visual Explain Tutorial

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

Related concepts:

- “Visual Explain overview” in *Administration Guide: Implementation*

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 in the DB2 Troubleshooting Guide or the Support and Troubleshooting section of the DB2 Information Center. There you will find information on how to isolate and identify problems using DB2 diagnostic tools and utilities, solutions to some of the most common problems, and other advice on how to solve problems you might encounter with your DB2 products.

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 or bug fixes), fix packs, 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/support.html>

Related concepts:

- “Introduction to problem determination” in *Troubleshooting Guide*
- “Overview of the DB2 technical information” on page 155

Terms and Conditions

Permissions for the use of these publications is granted subject to the following terms and conditions.

Personal use: You may reproduce these Publications for your personal, non commercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these Publications, or any portion thereof, without the express consent of IBM.

Commercial use: You may reproduce, distribute and display these Publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these Publications, or reproduce, distribute or display these Publications or any portion thereof outside your enterprise, without the express consent of IBM.

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the Publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the Publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

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

Company, product, or service names identified in the documents of the DB2 Version 9 documentation library may be trademarks or service marks of International Business Machines Corporation or other companies. Information on the trademarks of IBM Corporation in the United States, other countries, or both is located at <http://www.ibm.com/legal/copytrade.shtml>.

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 documentation library:

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

Intel, Itanium[®], Pentium[®], and Xeon[®] 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.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

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

Index

Special characters

- .NET applications
 - new data type support 43
- .NET Framework 2.0
 - DB2 .NET Data Provider enhancements 45
- A**
 - ADD DISTRIBUTE BY HASH clause
 - replaces ADD PARTITIONING KEY clause 143
 - ADD PARTITIONING KEY clause
 - deprecated 143
 - add-ins
 - IBM Database Add-Ins for Microsoft Visual Studio 2005
 - enhancements 49
 - administration
 - DB2 Client replaced DB2 Administration Client 64
 - native XML data store tools support 39
 - summary of changed functionality 123
 - administration commands
 - SQL access enhancements 81
 - administration tools
 - changed platform support 152
 - agents
 - new
 - db2stmm and db2taskd 128
 - ALTER TABLE statement
 - ADD PARTITIONING KEY clause deprecated 143
 - DROP PARTITIONING KEY clause deprecated 144
 - enhancements 78
 - Alternate FixPak images (UNIX)
 - replaced with multiple DB2 copies support 147
 - APIs
 - APPC (Advanced Program-to-Program Communication)
 - no longer supported 152
 - APPN (Advanced Peer-to-Peer Networking)
 - no longer supported 152
 - autoloader (db2atld)
 - discontinued 148
 - db2Load
 - new `sqlu_remotefetch_entry` 90
 - SNA (Systems Network Architecture)
 - no longer supported 152
 - vendor load (`sqluvtld`)
 - discontinued 153
 - `app_ctl_heap_sz` configuration parameter
 - default value changes 117
 - APPC (Advanced Program-to-Program Communication)
 - no longer supported 152
 - application development
 - .NET Framework 2.0 support 45
 - changed functionality 132
 - DB2 .NET Data Provider enhancements 45
 - DB2 Client replaced DB2 Application Development Client 64
 - development software support 53
 - enhancements 43
 - IBM Database Add-Ins for Microsoft Visual Studio 2005
 - enhancements 49
 - IBM Software Development Kit (SDK) for Java 5.x 135

- application development (*continued*)
 - JDBC enhancements 52
 - native XML data store enhancements 32
 - new sample applications 55
 - SQLJ enhancements 52
 - tools enhancements 9
- application ID format
 - changed 123
- application ID monitor element 123
- application servers
 - DB2 Embedded Application Server enhancements 134
- applications
 - ADO.NET
 - migration changes 121
 - CLI
 - migration changes 121
 - embedded SQL
 - migration changes 121
 - feature support
 - changed 132
 - Java
 - migration changes 121
 - migration changed 106
 - new connection timeout support 63
- APPN (Advanced Peer-to-Peer Networking)
 - no longer supported 152
- Audio, Image, and Video (AIV) Extenders
 - discontinued 148
- autoloader utility
 - discontinued 148
- automated evaluation
 - of object maintenance policies 78
- automatic configuration
 - changed Configuration Advisor default 116
 - page cleaners 113
 - prefetchers 113
- automatic features
 - RUNSTATS 114
 - statistics collection 114
- automatic maintenance
 - diagnostic messages
 - location changes 127
- automatic self tuning memory
 - changed default setting 113
- automatic storage
 - creating using Control Centre 79
 - enabled by default 115
 - multi-partition databases 79
- autonomic computing
 - features enabled by default 114
 - new features 10
- `autonomic_switches` configuration parameter
 - default value changes 117
- `avg_appls` configuration parameter
 - default value changes 117

B

- backup
 - enhancements 59

- backup image
 - changed naming convention
 - Windows 124
- BINARY collating sequence value
 - replaced with IDENTITY 125
- BINARY SQL data type
 - support added in .NET and CLI applications 43
- blocking bind option
 - ignored in SQL procedures 134
- blocking precompile option
 - ignored in SQL procedures 134
- buffer pool identifiers
 - database memory output
 - enhancement 81
- buffer pools
 - extended storage option
 - no longer supported 151

C

- catalog tables
 - new table partitioning support 124
- character data
 - multiple byte
 - processing strings along character boundaries
 - support 87
- character-based string functions
 - new Unicode support 87
- check pending state
 - replaced with set integrity pending state 139
- CLI
 - new trusted context for DB2 for z/OS databases
 - support 45
- CLI applications
 - new data type support 43
 - new support to run without DB2 client 65
- CLI/ODBC keywords
 - CLISchema
 - deprecated 140
- clients
 - changed application support 121
 - DB2 Connect 22
 - DB2 Runtime Client enhancements 64
 - enhancements 63
 - hardware support 17
 - new DB2 Client 64
 - operating system support 17
- code units
 - string functions support 87
- coexistence
 - support for multiple DB2 versions
 - Linux and UNIX 69
 - Windows 70
- collating sequence value
 - BINARY
 - replaced with IDENTITY 125
- COLNAMES column
 - deprecated 140
- column type
 - changing with the ALTER TABLE statement 78
- command line processor (CLP)
 - native XML data store support 40
 - removed 64 KB limit for SQL statement 45
- commands
 - db2_deinstall
 - a option required 106
 - changed 73

- commands (*continued*)
 - db2_install
 - changed 73
 - db2batch
 - changed for native XML data store 41
 - changed parameters and output 133
 - db2extsec
 - replaced db2secv82 141
 - db2fs 83
 - db2icrt
 - w option ignored and deprecated 145
 - migration support changes (Windows) 109
 - db2ilist
 - w option ignored and deprecated 145
 - deprecated options 140
 - db2imigr
 - migration support changes (Windows) 109
 - db2iupdt
 - w option ignored and deprecated 145
 - db2licm
 - changed 108
 - db2look
 - changed collating sequence display 125
 - changed for multi-partition databases 115
 - changed for native XML data store 41
 - db2ls
 - added (Linux and UNIX) 69
 - db2mtrk
 - buffer pool identifiers 81
 - changed parameters 126
 - db2pd
 - new high water mark option 101
 - db2sampl
 - overview 56
 - db2secv82
 - replaced with db2extsec 141
 - db2trc
 - new trace mask support 102
 - db2unins
 - added 74
 - doce_deinstall
 - added 108
 - REORGCHK
 - changed output 129
 - ROLLFORWARD DATABASE
 - changed output 129
 - RUNSTATS
 - changed for native XML data store 38
- communication protocols
 - new IPv6 communication protocol support 65
- communications speed
 - using fast communications manager (FCM) 82
- components
 - renamed and replaced 26
- compression dictionaries
 - static dictionary-based lossless row compression 92
- Configuration Advisor
 - changed default setting 116
- configuration parameters
 - database_memory changed 119
 - default value changes 117
 - estore_seg_sz
 - deprecated 151
 - fcm_num_anchors
 - replaced with fcm_num_channels 82
 - fcm_num_buffers
 - changed 82

- configuration parameters (*continued*)
 - fcm_num_channels
 - added 82
 - fcm_num_connect
 - replaced with fcm_num_channels 82
 - fcm_num_rqb
 - replaced with fcm_num_channels 82
 - max_connections
 - changed 128
 - num_estore_segs
 - deprecated 151
 - sheapthres_shr
 - changed 130
- connection timeout support
 - added 63
- connections
 - DB2 Connect support 22
 - timeout support added 63
- connectivity
 - enhancements 63
- contacting IBM 171
- Control Center
 - new native XML data store support 39
 - VSE and VM objects no longer supported 154
- copying
 - schemas
 - between databases 81
- CREATE DATABASE statement
 - new RESTRICT option 98
- CREATE INDEX statement
 - new default 136
- cursor
 - file type modifiers 127
- cursor blocking
 - disabled in SQL procedures 134

D

- data access
 - label-based access control (LBAC) 97
 - options 15
- data collection
 - trap enhancements 101
- data compression
 - enhancement 92
- Data Links Manager
 - migrating 148
 - no longer supported 148
- data loading for partitioned tables
 - support added 84
- data movement
 - REMOTEFETCH media type 90
- data partitioning
 - see table partitioning 12
- data source
 - multivendor
 - two-phase commit 68
- data types
 - BINARY support added 43
 - DECFLOAT support added 43
 - VARBINARY support added 43
 - XML
 - parameter support in routines 33
- Data Warehouse Center
 - no longer available 150
- database administration
 - autonomic computing 10

- database applications
 - new connection timeout support 63
 - new samples 55
- database creation
 - new default setting for statistics collection 114
 - new RESTRICT option 98
- database logs
 - raw devices are deprecated 142
- database management
 - automatic self tuning memory 113
 - Configuration Advisor
 - new default 116
- database objects
 - new identifier length limit 84
 - permissions
 - new db2extsec command 141
- database partition servers
 - fast communications manager (FCM) enhancements 82
- database partitions
 - execution of routines supported on database partitions 49
 - new external table function support 49
- database servers
 - hardware support 17
 - instance support
 - changed 121, 136
 - operating system support 17
- database_memory configuration parameter
 - changed 119
 - default value changes 117
- databases
 - creating sample 56
 - migration
 - support and restrictions 110
 - rebuild database function 60
 - redirected restore 59
 - schema copy enhancement 81
- DB2 .NET Data Provider
 - enhancements 45
- DB2 Administration Tools
 - changed platform support 152
- DB2 Client
 - hardware support 17
 - operating system support 17
- DB2 Connect
 - application development changes
 - application ID format changed 123
 - db2batch command changes 133
 - SDK for Java 5.x support 135
 - application development enhancements
 - BINARY, VARBINARY, and DECFLOAT data type support 43
 - DB2 .NET Data Provider enhancements 45
 - Developer Workbench 47
 - development software support 53
 - IBM Database Add-Ins for Microsoft Visual Studio 2005 49
 - IBM DB2 Driver for JDBC and SQLJ 52
 - removed 64 KB SQL statement limit 45
 - XML support in Developer Workbench 34
 - XML support in SQL Assist 36
- component replacements 26
- connectivity enhancements
 - DB2 Runtime Client 64
 - new connection timeout support 63
 - new DB2 Client 64
 - new IPv6 communication protocol support 65

- DB2 Connect (*continued*)
 - database setup changes
 - w option ignored for db2icrt, db2ilist, and db2iupdt commands 145
 - configuration parameter default values changes 117
 - deprecated functionality
 - CLISchema CLI keyword 140
 - DB2 JDBC Type 2 driver 145
 - db2secv82 command 141
 - discontinued functionality
 - DB2 Administration Tools not supported on some platforms 152
 - db2profc utility replaced with db2sqljcustomize 150
 - db2profp utility replaced with db2sqljprint 150
 - desktop icon utility 151
 - folder making utility 151
 - JDBC Type 3 driver 153
 - NetBIOS communication protocol 152
 - SNA communication protocol 152
 - VSE and VM objects 154
 - federation enhancements
 - statement level isolation for nicknames 68
 - two-phase commit for multivendor data sources 68
 - user mapping retrieval from an external repository 67
 - installation changes
 - Administrator user ID not required (Windows) 72
 - db2_deinstall command parameter changes (Linux and UNIX) 106
 - Information Center restrictions 108
 - install images package format (Linux and UNIX) 106
 - licensing support 108
 - migration changes (Windows) 109
 - multiple CD installation (Linux and UNIX) 111
 - national language support (NLS) 108
 - installation enhancements
 - db2unins command (Windows) 74
 - documentation updates between product releases 107
 - manual installation (Linux and UNIX) 73
 - multiple DB2 versions and fix packs (Linux and UNIX) 69
 - multiple DB2 versions and fix packs (Windows) 70
 - response file 74
 - manageability enhancements
 - automated evaluation of object maintenance policies 78
 - EXEC SQL utility support for loading data into DB2 UDB for z/OS tables 82
 - First Steps interface 83
 - migration changes 106
 - support changes 110
 - migration enhancements
 - DB2 Client upgrade 21
 - name changes 26
 - product line changes 27
 - summary of enhancements 19
 - support changes 22
 - troubleshooting enhancements
 - data collection for database system failures 101
 - db2trc command 102
- DB2 Connect Personal Edition
 - upgrade enhancement 21
- DB2 database system
 - default DB2 copy
 - Windows 70
 - support for multiple copies
 - Windows 70
- DB2 database system (*continued*)
 - support for multiple versions
 - Windows 70
- DB2 Information Center
 - updating 161
 - versions 159
 - viewing in different languages 160
- DB2 Information Integrator
 - product name change 25
- DB2 JDBC Type 2 driver
 - deprecated 145
- DB2 Runtime Client
 - enhancements 64
- DB2 UDB for z/OS health monitor
 - automated evaluation of object maintenance policies 78
- DB2 Universal Database Version 8
 - component replacements 26
 - name changes 26
 - product name change 25
- DB2 Version 9.1
 - changes in existing functionality 105
 - enhancements 15
 - highlights 3
 - new features 27
 - product line changes 27
- DB2 versions
 - support for multiple copies
 - Windows 70
- DB2 Web Tools
 - updated 134
- DB2_ALLOCATION_SIZE
 - changed default value 121
- DB2_COLLECT_TS_REC_INFO
 - changed default value 121
- db2_deinstall command
 - a option required 106
 - changed 73
- DB2_FORCE_FCM_BP
 - changed default value 121
- db2_install command
 - changed 73
- DB2_MDC_ROLLOUT
 - changed default value 121
- DB2_OPT_MAX_TEMP_SIZE
 - added 91
- DB2_REDUCED_OPTIMIZATION
 - new NO_SORT_MGJOIN keyword 91
 - new NO_SORT_NLJOIN keyword 91
- DB2_RESOLVE_CALL_CONFLICT registry variable
 - new 135
- DB2_SCATTERED_IO
 - discontinued 151
- DB2_SERVER_CONTIMEOUT
 - new registry variable 127
- DB2_SMS_TRUNC_TMPTABLE_THRESH
 - changed default value 121
- db2atld utility
 - discontinued 148
- db2batch command
 - changed parameters and output 133
 - XML support changes 41
- DB2CHKSQLDA
 - changed default value 121
- db2icrt command
 - w option ignored and deprecated 145
- db2ilist command
 - w option ignored and deprecated 145

- db2ilist command (*continued*)
 - deprecated options
 - Linux and UNIX 140
- db2iupdt command
 - w option ignored and deprecated 145
- DB2LINUXAIO registry variable
 - deprecated 141
- db2look command
 - changed collating sequence display 125
 - changed for multi-partition databases 115
 - XML support changes 41
- db2mtrk command
 - buffer pool identifiers 81
 - changed parameters
 - Windows 126
- db2pd command
 - new FCM option 101
- db2profc utility
 - replaced with db2sqljcustomize 150
- db2profp utility
 - replaced with db2sqljprint 150
- db2reg2large utility
 - discontinued 150
- db2sampl command
 - overview 56
- db2secv82 command
 - deprecated 141
- db2trc command
 - new trace mask support 102
- DECFLOAT data type
 - support added in .NET and CLI applications 43
- decomposition
 - annotated XML schema 37
- deprecated functionality
 - ADD PARTITIONING KEY clause 143
 - CLISchema keyword 140
 - COLNAMES column 140
 - DB2 JDBC Type 2 driver 145
 - db2ilist command options
 - Linux and UNIX 140
 - DB2LINUXAIO registry variable 141
 - db2secv82 command 141
 - DROP PARTITIONING KEY clause 144
 - external routine default entry point support 141
 - iCheckPending parameter 139
 - overview 139
 - raw devices for database logging 142
 - SQL administrative routines 143
 - table functions and procedures 143
 - Type 1 indexes 144
- desktop icon utility
 - no longer supported on Linux 151
- Developer Workbench
 - new XML support 34
 - new XQuery builder 31
 - overview 47
- Development Center
 - replaced by Developer Workbench 47
- development software support
 - changed 53
- discontinued functionality
 - Audio, Image, and Video (AIV) Extenders 148
 - DB2 Administration Tools
 - not supported on some platforms 152
 - db2profc utility 150
 - db2profp utility 150
 - db2reg2large utility 150

- discontinued functionality (*continued*)
 - desktop icon utility
 - Linux 151
 - Extended Storage option for buffer pools 151
 - folder making utility
 - Linux 151
 - JDBC Type 3 driver 153
 - NetBIOS communication protocol 152
 - overview 147
 - SNA communication protocol 152
 - VSE and VM objects 154
- disk space
 - storage savings using row compression 92
- distribution key terminology
 - changed 27
- DMS table spaces
 - default behavior 126
 - larger record identifiers 96
 - size conversion 150
- documentation 155, 156
 - enhancement 107
 - Information Center installation
 - changed 108
 - new installation path defaults 108
 - migration resources 73
 - new command to remove Information Center (Linux) 108
 - new feature for updating Information Center 75
 - terms and conditions of use 163
 - updates 107
- drivers
 - CLI
 - new driver added 65
 - DB2 JDBC Type 2 driver deprecated 145
 - DB2 JDBC Type 3
 - discontinued 153
 - ODBC
 - new driver added 65
- drop distribution syntax
 - replaces drop partitioning key syntax 144
- drop partitioning key syntax
 - deprecated 144
- dropping a column
 - ALTER TABLE statement 78

E

- enhancements
 - federated systems 67
 - information as a service 6
 - manageability 77
 - performance 89
 - scalability 95
 - security 97
 - serviceability 101
- environment variables
 - changed default values 121
- estore_seg_sz configuration parameter
 - deprecated 151
- event monitors
 - write-to-table changes 131
- EXEC SQL utility
 - new support for loading data into DB2 UDB for z/OS tables 82
- Explain
 - new SQL/XML and XQuery statement support 40
- export utility
 - changes to syntax and behavior 125

- export utility (*continued*)
 - new XML support 42
- extended security
 - authority changes for starting database manager
 - Windows 113
- extender support
 - Audio, Image, and Video (AIV) Extenders
 - discontinued 148
 - Text Extender discontinued 153
- external functions
 - execution supported across database partitions 49
- external procedures
 - new XML data type support 33
- external repositories
 - new plug-in for retrieving user mappings 67
- external routines
 - deprecated default entry point support 141

F

- fast communications manager (FCM)
 - enhancements 82
 - no longer uses virtual interface (VI) architecture 152
- federated systems
 - enhancements 67
 - two-phase commit 68
- file type modifiers
 - cursor 127
 - LOAD command 127
- First Steps
 - enhancements 83
- fix packs
 - Alternate FixPak images discontinued (UNIX) 147
 - changed functionality 105
 - new functionality 69
- FixPaks
 - Alternate FixPak images replaced with multiple DB2 copies support (UNIX) 147
- folder making utility
 - no longer supported on Linux 151

G

- GET DATABASE CONFIGURATION command
 - changed collating sequence display 125

H

- hardware support
 - DB2 database servers and clients 17
- health monitor
 - new support for DB2 UDB for z/OS 78
- help
 - displaying 160
 - for SQL statements 159

I

- IBM Database Add-Ins
 - Microsoft Visual Studio 2005 49
- IBM Software Development Kit (SDK)
 - new Java 5.x support 135
- iCheckPending parameter
 - replaced with iSetIntegrityPending parameter 139

- identifier length limits
 - increased to 128 bytes for additional objects and names 84
- identity column values
 - generated by SET INTEGRITY 91
- import utility
 - changes to syntax and behavior 125
 - new XML support 42
- INCURSOR option
 - supported in EXEC SQL utility 82
- indexes
 - automatic reorganization enhancements 80
 - CREATE INDEX statement
 - ALLOW REVERSE SCANS 136
 - defined on more columns 95
 - enhancements 95
 - larger key parts 95
 - new performance guidelines for XML 38
 - new XML support 37
 - Type 1 indexes deprecated 144
- Information Catalog Center
 - no longer available 150
- Information Center
 - installation restrictions 108
 - uninstallation changes 108
 - updating 161
 - new feature 75
 - versions 159
 - viewing in different languages 160
- install images
 - changed format
 - Linux and UNIX 106
- installation
 - changed functionality 105, 112
 - Information Center
 - restrictions 108
 - installing National Language Package
 - Linux and UNIX 108
 - manual installation enhancements
 - Linux and UNIX 73
 - multiple CD installation
 - Linux and UNIX 111
 - new concurrent installation support
 - Linux and UNIX 69
 - Windows 70
 - new functionality 69
 - new uninstall features
 - Windows 74
 - non-administrator installation enhancement
 - Windows 72
 - package format changes
 - Linux and UNIX 106
 - response file keywords 74
- instances
 - 32-bit and 64-bit support 121
 - DB2 database server instance support changes 136
 - migration
 - support and restrictions 110
- integrity processing
 - online 91
- iSetIntegrityPending parameter
 - replaces iCheckPending parameter 139
- isolation clauses
 - federated SQL statement changes 68
- IXF files
 - import utility changes 125

J

- Java
 - applications
 - migration changes 121
 - IBM Software Development Kit (SDK) for Java 5.x 135
 - procedures
 - class loader enhancements 50
 - stored procedures
 - class loader enhancements 50
- java_heap_sz configuration parameter
 - default value changes 117
- JDBC
 - DB2 JDBC Type 2 Driver
 - deprecated 145
 - enhancements 52
 - new trusted context for DB2 for z/OS databases
 - support 45
 - Type 3 support discontinued 153
- JVM (Java virtual machine)
 - DB2 Connect support 22
 - DB2 database instance support changes 121

K

- key parts
 - indexes
 - increased maximum size 95

L

- label-based access control (LBAC)
 - improved data access security 97
- large table spaces
 - changed default value for new DMS table spaces 126
- LDAP servers
 - new plug-in for retrieving user mappings 67
- licenses
 - support changes 108
- LOAD command
 - incompatibilities when working with previous releases 127
- load utility
 - autoloader utility (db2atld) discontinued 148
 - changed default paths for LOB 125
 - enhancements 84
- loading data
 - autoloader utility (db2atld) discontinued 148
 - enhancements 90
 - EXEC SQL utility support for DB2 UDB for z/OS tables 82
 - load from cursor incompatibilities 127
 - load from cursor with remote fetch support 90
 - partitioned table enhancements 84
 - SOURCEUSEREXIT customizable user exit 90
 - vendor load API (sqluvtld) discontinued 153
- lock lists
 - additional space requirement 120
- locklist configuration parameter
 - default value changes 117
- logging
 - enhancements 59
- logs
 - TSM class no longer used
 - restoring and retrieving 131

M

- manageability
 - enhancements 77
- mapping
 - new plug-in for retrieving user mappings 67
- materialized query tables (MQTs)
 - enhancements 90
 - maintenance 90
 - performance 90
 - REFRESH DEFERRED
 - considered during query optimization 91
- max_connections configuration parameter
 - changed 128
- maxlocks configuration parameter
 - default value changes 117
- memory
 - automatic self tuning 113
 - increased requirements 120
 - new autonomic features 77
 - self-tuning memory allocation 77
- Memory Tracker command
 - buffer pool identifiers 81
 - changed parameters
 - Windows 126
- Microsoft Visual Studio 2005
 - IBM Database Add-Ins 49
- migrating
 - DB2 Client to DB2 Connect Personal Edition
 - enhancement 21
 - instances
 - 32-bit and 64-bit considerations 121
- migration
 - application migration changed 106
 - changed functionality 105
 - changes 110
 - new functionality 69
 - resources 73
 - routine migration changed 106
 - support changes for 32-bit and 64-bit (Windows) 109
- monitoring
 - between partitioned database environments 85
- multi-partition databases
 - automatic storage support 79
- multiple byte character data
 - processing strings along character boundaries support 87
- multiple DB2 copies
 - new support
 - Windows 70

N

- naming conventions
 - backup image naming changed (Windows) 124
 - export LOB files naming changed 125
- naming rules
 - new identifier length limits for more database objects and names 84
- national language support (NLS)
 - enhancements 87
 - installation changes 108
- native XML data store
 - administration tools support 39
 - application development support 32
 - application programming language support 32
 - CLP support 40
 - Control Center support 39

- native XML data store *(continued)*
 - db2batch command 41
 - db2look command 41
 - Explain support 40
 - export utility support 42
 - external entity management 30
 - indexing XML data
 - performance guidelines 38
 - support 37
 - load utility support 42
 - Net Search Extender support 32
 - new XML schema repository 30
 - overview 29
 - performance considerations 37
 - query optimizer support 39
 - SQL statement support 35
 - SQL/XML functions 35
 - statistics collection support 38
 - system command support 40
 - Visual Explain support 40
 - XML DTD management 30
 - XML schema management 30
- Net Search Extender
 - new platforms supported 32
 - support for native XML data store 32
- NetBIOS
 - no longer supported 152
- nicknames
 - statement level isolation 68
- NO_SORT_MGJOIN
 - new keyword for DB2_REDUCED_OPTIMIZATION 91
- NO_SORT_NLJOIN
 - new keyword for DB2_REDUCED_OPTIMIZATION 91
- notices 165
- nullability attribute of a column
 - changing with the ALTER TABLE statement 78
- num_iocleaners configuration parameter
 - default value changes 117
- num_ioservers configuration parameter
 - default value changes 117

O

- Object Maintenance Policy wizards
 - automated evaluation 78
- ODBC
 - new trusted context for DB2 for z/OS databases support 45
- ODBC applications
 - new support to run without DB2 client 65
- online integrity processing
 - new support 91
- operating system support
 - DB2 database servers and clients 17
- optimizer support
 - native XML data store 39
- ordering DB2 books 158

P

- packaging
 - changed 25
- page cleaners configuration
 - changed 113
- page sizes
 - new 64 KB support on AIX (POWER5+) 93

- parallel communication
 - fast communications manager (FCM) 82
- partitioned databases
 - monitoring information flow 85
 - new SQL Monitor Area (SQLMA) interface 85
- partitioned tables
 - catalog table changes 124
 - improved large database management 12
- partitioning data
 - table partitioning support 12
- pckcachesz configuration parameter
 - default value changes 117
- performance
 - enhancements 89
 - guidelines for XML indexing 38
 - query performance using statistical views 89
 - tuning
 - DB2_SCATTERED_IO variable is discontinued 151
- permissions
 - database objects
 - new db2extsec command 141
- platform support
 - changed for DB2 Administration Tools 152
- POWER5+ processors
 - new 64 KB page size support
 - AIX 93
- prefetchers configuration
 - changed 113
- printed books
 - ordering 158
- privileges
 - new security administrator authority (SECADM) 99
 - new SETSESSIONUSER privilege 99, 129
- problem determination
 - data collection improvements for traps 101
 - online information 163
 - tutorials 163
- procedures
 - Java
 - class loader enhancements 50
 - parameters
 - XML data type support 33
- product changes
 - overview 25, 27

Q

- query optimization
 - considering REFRESH DEFERRED materialized query tables 91
- query performance
 - support for statistical views 89

R

- range partitioning
 - see table partitioning 12
- raw devices for database logging
 - deprecated 142
- rebuilding
 - rebuild database restore options 60
- record identifiers (RIDs)
 - enhancements 96
- RECOVER DATABASE command
 - new RESTART option 59

- recovery
 - continuing an operation that ended 59
 - enhancements 59
- redirected restore
 - enhancements 59
- REFRESH DEFERRED materialized query tables
 - considered during query optimization 91
- registry variables
 - changed default values 121
 - DB2_COPY_NAME
 - added 70
 - DB2_OPT_MAX_TEMP_SIZE
 - added 91
 - DB2_REDUCED_OPTIMIZATION
 - changed 91
 - DB2_RESOLVE_CALL_CONFLICT 135
 - DB2_SERVER_CONTIMEOUT 127
 - DB2_VI_DEVICE
 - discontinued 152
 - DB2_VI_ENABLE
 - discontinued 152
 - DB2_VI_VIPL
 - discontinued 152
 - DB2LINUXAIO
 - deprecated 141
- REMOTEFETCH media type 90
 - data movement 90
- reorganization
 - indexes 80
 - REORGCHK command output changes 129
 - tables 80
- REORGCHK command
 - output changes 129
- response files
 - enhancements 74
- restore
 - new restore options
 - rebuild database function 60
 - redirected
 - automatically generated scripts 59
- reverse scans
 - CREATE INDEX statement 136
- RID (record identifiers)
 - log space
 - increased requirements 120
- ROLLFORWARD DATABASE command
 - output enhancement 129
- rollforward recovery
 - restrictions 131
- routines
 - external
 - deprecated default entry point support 141
 - feature support
 - changed 132
 - migration
 - 32-bit and 64-bit considerations 121
 - migration changed 106
 - SQL administrative routines
 - deprecated 143
- row compression
 - benefits 92
 - support added 92
- RUNSTATS command
 - automatic statistics collection 114
 - XML support 38

S

- sample database
 - enhancements 56
- samples
 - new samples 55
- scalability
 - enhancements 95
- schemas
 - copying enhancement 81
 - new XML schema repository 30
- SDK for Java 5.x
 - new support 135
- searching
 - new native XML text search support 32
- SECADM authority 99
- security
 - authority changes for extended security
 - Windows 113
 - data access control enhancement 97
 - enhancements 15, 97
 - RESTRICT option on CREATE DATABASE 98
 - SECADM authority 99
 - TRANSFER OWNERSHIP SQL statement 100
- security administrator authority (SECADM)
 - XXXX 99
- serviceability
 - enhancements 101
- session authorization ID
 - new SETSESSIONUSER privilege 129
- set integrity pending state
 - replaces check pending state 139
- SET SESSION AUTHORIZATION statement
 - changed 129
- SETSESSIONUSER privilege 99
- setup
 - changed functionality 112
 - tool
 - Windows 83
- sheapthres configuration parameter
 - default value changes 117
- sheapthres_shr configuration parameter
 - changed 130
 - default value changes 117
- SNA (Systems Network Architecture)
 - no longer supported 152
- sorheap configuration parameter
 - default value changes 117
- SOURCEUSEREXIT option
 - data movement 90
- space requirements
 - increased for databases 119
 - increased for lock lists 120
 - increased for log space 120
 - increased for table spaces 120
- SQL
 - DB2 administration command access 81
 - new XML data type support 33
- SQL administrative routines
 - deprecated routines 143
- SQL Assist
 - new XML data type support 36
- SQL Monitor Area (SQLMA)
 - new functionality for monitoring information flow 85
- SQL statement length
 - command line processor (CLP) 45
- SQL statements
 - displaying help 159

- SQL statements (*continued*)
 - Explain and Visual Explain support 40
 - new XML support 35
- SQL/XML
 - functions
 - enhancement 35
- SQLJ
 - enhancements 52
 - new trusted context for DB2 for z/OS databases
 - support 45
- statistical views
 - enhanced query performance 89
- statistics collection
 - automatic
 - changed default setting 114
 - for tables containing XML columns 38
- storage
 - automatic storage
 - changed default setting 115
 - multi-partition databases 79
 - new automatic storage support 79
 - Tivoli Storage Manager support Windows x64 61
- storage paths
 - adding using Control Centre 79
- STRIP scalar function
 - added 57
- supported development software 53
- syntax
 - drop partitioning key syntax 144
- system temporary tables
 - enhancements 95

T

- table functions
 - external
 - execution supported across database partitions 49
- table partitioning
 - catalog table changes 124
 - changed distribution key terminology 27
 - improved large database management 12
 - support added 12
- table space
 - increased requirements 120
 - size enhancements 96
- tables
 - automatic reorganization enhancements 80
 - DB2 UDB for z/OS
 - EXEC SQL utility support for loading data 82
 - system temporary table enhancements 95
 - user temporary table enhancements 95
- TCP/IP
 - IPv6 support 65
- temporary tables
 - enhancements 95
- terms and conditions
 - use of publications 163
- Text Extender
 - no longer supported 153
- text search
 - new native XML text search support 32
- timeout
 - connect request
 - setting 127
- Tivoli Storage Manager (TSM)
 - new Windows x64 support 61

- trace mask
 - db2trc command 102
- TRANSFER OWNERSHIP SQL statement 100
- trap files
 - serviceability enhancement 101
- TRIM scalar function
 - added 57
- troubleshooting
 - data collection improvements for traps 101
 - online information 163
 - tutorials 163
- trusted context
 - support added for DB2 for z/OS databases 45
- TSM management class
 - no longer used for filtering
 - restoring and retrieving 131
- tuning
 - self-tuning memory 77
- tutorials
 - troubleshooting and problem determination 163
 - Visual Explain 162
- two-phase commit
 - federated systems enhancement 68
- Type 1 indexes
 - deprecated 144
- Type 3 JDBC driver
 - no longer supported 153

U

- uninstalling
 - all components (Linux and UNIX) 106
 - Information Center 108
 - new features
 - Windows 74
- updates
 - DB2 Information Center 161
 - documentation 107
 - Information Center 161
- user exit
 - loading data 90
- user mappings
 - new plug-in for retrieving from external repository or LDAP server 67
- user temporary tables
 - enhancements 95
- userexit configuration parameter
 - default value changes 117
- users
 - new security administrator authority (SECADM) 99
- utilities
 - autoloader utility (db2atld) discontinued 148

V

- VARBINARY SQL data type
 - support added in .NET and CLI applications 43
- vendor load API (sqluvtld)
 - discontinued 153
- virtual interface (VI) architecture
 - no longer used by fast communications manager (FCM) 152
- Visual Explain
 - new SQL/XML and XQuery statement support 40
 - tutorial 162

- VM databases
 - no longer supported in DB2 Control Center 154
- VSE databases
 - no longer supported in DB2 Control Center 154

W

- WebSphere Federation Server
 - product name change 25
- Windows
 - installation requirements 72
 - migration changes 109
 - new DB2 uninstall features 74
- WITH isolation clause
 - required with nicknames 68
- write-to-table event monitors
 - changed 131

X

- XA
 - new trusted context for DB2 for z/OS databases
 - support 45
- XML
 - new programming language support 32
 - new XQuery language support 30
 - SQL/XML support 35
 - XQuery builder overview 31
- XML columns
 - retrieving data
 - new XQuery language support 30
- XML data
 - new native XML text search support 32
 - statistics collection 38
- XML data retrieval
 - XQuery builder overview 31
- XML data store
 - administration tools support 39
 - application development support 32
 - command line processor (CLP) support 40
 - Control Center support 39
 - export utility support 42
 - import utility support 42
 - optimizer support 39
- XML data type
 - parameter support in routines 33
 - SQL Assist support 36
- XML Metadata Registry
 - no longer supplied 134
- XML schema repository
 - native XML data store 30
- XML schemas
 - new XML schema decomposition 37
- XML support
 - native XML data store 29
 - new Developer Workbench 34
- XQuery
 - new XQuery language support 30
 - XQuery builder overview 31
- XQuery statements
 - Explain and Visual Explain support 40

Contacting IBM

To contact IBM in your country or region, check the IBM Directory of Worldwide Contacts at <http://www.ibm.com/planetwide>

To learn more about DB2 products, go to <http://www.ibm.com/software/data/db2/>.



Printed in USA

SC10-4253-00



DB2 9 BETA

Spine information:

DB2 9.5

IBM DB2
DB2 Version 9

What's New

