

DB2 Universal Database Version 8.1.4



Release Notes (Version 8.1.4)

Version 8.1.4

DB2 Universal Database Version 8.1.4



Release Notes (Version 8.1.4)

Version 8.1.4

Contents

About the Release Notes	vii	License Center back-level versions not supported	18
4 What's New for Version 8.1.4	1	Microsoft Visual Studio, Visual C++	18
4 Application development enhancements	1	Microsoft XP fix needed on 64-bit operating systems	18
4 DB2 family enhancements	4	MVS operating system not supported	18
4 Manageability enhancements	4	SNA support limitations in Version 8.	18
4 Performance enhancements	5	Supported LDAP client and server configurations	19
4 Replication enhancements	7	Tivoli Storage Manager for LinuxAMD64 Windows XP operating systems	20
Known problems and workarounds (Version 8.1.4)	9	Application development	20
Product and product-level support	9	Asynchronous execution of CLI	20
3 Alternate FixPaks on Linux and UNIX-based systems	9	CLI and ODBC on Windows 64-bit operating systems	20
Back-level DB2 Universal Database server support	10	Configuration Assistant	20
Classic Connect unavailable	12	Unsupported bind options	20
Data Warehouse Center down-level server support restrictions	12	Configuration parameters	21
DB2 license policy for DB2 Universal Database Workgroup Server Edition	13	NUM_LOG_SPAN configuration parameter on a multipartition database	21
DB2 Universal Database Version 7 server access	13	dasdrop limitation in multiple FixPak environments	21
Data Warehouse Center not available in Simplified Chinese	13	Data Warehouse Center	22
DB2 Web Tools	13	ERwin 4.x metadata bridge	22
2 Development Center APARs required for SQLJ and SQL Assist support on DB2 UDB for OS/390, Version 6 and DB2 UDB for z/OS, Version 7.	14	Japanese names of remote objects	23
Development Center limitations for 64-bit operating systems	14	Restrictions for the Clean Data transformer Using the warehouse agent for replication and accessing Client Connect warehouse sources.	23
1 Development Center now supports Actual Cost information for SQL statements run on OS/390 or z/OS servers	14	Scheduling a warehouse process to run at intervals	24
1 Development Center on the Linux operating system	15	Replication Center restrictions on iSeries systems	24
2 Federated systems restrictions	15	Restriction for import and export	25
High availability function is included in DB2 Universal Database Workgroup Server Edition.	17	The Visual Warehouse 5.2 DB2 program "VW 5.2 Load flat file into DB2 UDB EEE (AIX only)" is not supported	25
2 IBM DB2 Development Add-In for Microsoft Visual Studio .NET	17	Limited support for CURSOR load	26
2 Installation on AIX.	17	DB2 Cube Views	26
2 Installation on Linux	18	Cube Views sample application differs from sample source	26
		DB2 Data Links Manager	26
		Data Links server backup fails using Tivoli Storage Manager archive server (AIX, Solaris Operating Environment)	26

DB2 Universal Database backup and restore	27	1	Migrating DB2 Universal Database when	
Backup and restore on Linux 390 operating		1	using DataJoiner or replication	34
systems	27		Query Patroller Center	34
DB2 Universal Database replication	28	4	Enhanced filtering abilities	34
Java Administrative API documentation for		1	Secure Windows environments.	37
DB2 data replication	28		SQL Assist	38
Column mapping restrictions and the			SQL Assist button disabled in the	
Replication Center	28		Command Center	38
Development Center	28		Two versions of SQL Assist are launched	
4			from DB2	38
4			System Monitor	38
4		3	Event record size limitation	38
Debugging Stored Procedures with double			Snapshot UDF restrictions	38
quotation marks	28	4	Amount of Log to be Redone for Recovery	
SQLFLAG(STD) precompiler option error	28	4	monitor element undefined for V8.1.4	38
Documentation	28	2	Throttled utility restrictions	39
DB2 Replication Guide and Reference			XML Extender	39
documentation	28	2	XML Extender sample programs renamed	39
DB2 Universal Database, Version 8 HTML		4	Decomposing documents in XML Extender	
documentation installation restriction		4	that contain non-unique attribute and	
(Windows)	28	4	element names	40
Documentation search may fail on AIX			Additional Information	43
unless all documentation categories are			Change in Unicode server behavior	43
installed	29		Full message text is not returned when	
Documentation search problem with Java 2			using SQLException.getMessage().	43
JRE1.4.0	29		IBM DB2 Universal JDBC driver	43
Installation of the DB2 Information Center			Java functions and routines on Linux,	
for languages that are not options during			UNIX, and Windows operating systems.	44
installation	29		English Microsoft Data Access	
Official naming convention for DB2			Components (MDAC) files are used for all	
Universal Database for Linux when used			national language versions of DB2	
on host systems.	30		Universal Database, V8.1, unless translated	
GUI tools	30		MDAC files are installed first.	44
Control Center plugin support	30		The Simplified Chinese locale on AIX	
Displaying Indic characters in the DB2			operating systems	44
GUI tools	30		The Simplified Chinese locale on Red Hat,	
GUI tools not supported for zSeries servers			V8 operating systems	45
running Linux operating systems	31			
Load and Import Columns page does not		3		
support DBCS characters in IXF files.	32	3		
Incorrect indicators given when a load				
operation fails	32			
Minimum display settings for GUI tools	32			
SQL1224N error when using the GUI tools				
on AIX	32			
Health Monitor	32			
Health Monitor off by default	32			
Health indicator restrictions.	32			
Information Catalog Center	33			
2				
4				
4				
Do not partition information catalog tables	33			
Log file not generated when importing tag				
language files	33			

Corrections to the documentation 47

4	Administration Guide: Planning	47
4	Application development: Call Level Interface	
4	(CLI)	50
4	Data Links Manager	51
4	Data Warehouse Center	52
4	Information Catalog Center	54
4	National languages.	54
4	Registry and environment variables	56
4	XML Extender	56

Online help corrections and updates 59

Configuring the C environment for SQL stored procedures in the Development Center	59	Runstats dialog – updated getting there information	61
2 Enabling view docking when accessing Development Center with Hummingbird		Specifying build options for a Java stored procedure in the Development Center	61
2 Exceed	59		
2 Microsoft Visual Studio .NET add-in information update in Development Center			
2 help.	60		
2 Migrating DB2 XML Extender to Version 8.1.2 Path settings to enable Java routines to compile in the Development Center	60 61	Appendix. Notices	63
		Trademarks	66

About the Release Notes

Content:

The release notes contain the latest information for the following DB2®, Version 8 products:

DB2 Universal Database™ Personal Edition
DB2 Universal Database Workgroup Server Edition
DB2 Universal Database Workgroup Server Unlimited Edition
DB2 Universal Database Enterprise Server Edition
DB2 Personal Developer's Edition
DB2 Universal Developer's Edition
DB2 Warehouse Manager
DB2 Warehouse Manager Sourcing Agent for z/OS™
DB2 Data Links Manager
DB2 Net Search Extender
DB2 Spatial Extender
DB2 Intelligent Miner™ Scoring
DB2 Intelligent Miner Modeling
DB2 Intelligent Miner Visualization
DB2 Connect™ Application Server Edition
DB2 Connect Enterprise Edition
DB2 Connect Personal Edition
DB2 Connect Unlimited Edition
DB2 Query Patroller

Structure:

- 3 The release notes are divided into four parts:
- 3 • The first part highlights what is new in this release.
 - 3 • The second part contains the details of the problems, limitations, and
3 workarounds known at the time of publication that affect the products
3 previously listed. Read these notes to become familiar with any known
3 outstanding issues with this release of the DB2 family of products.
 - 3 • The third part contains corrections to previously released HTML, PDF, and
3 printed documentation.
 - 3 • The fourth part contains corrections and updates to the information
3 available in the product GUI tools help.

1 The most up-to-date documentation is available in the latest version of the
1 DB2 Information Center, which is accessed through a browser. The URL for
1 downloading the latest documentation is provided in the Additional resources
1 section below.

1 Revision marks in the DB2 Information Center documentation indicate text
1 that has been added or changed since the PDF information for version 8.1 was
1 originally made available. A vertical bar (|) indicates information that was
1 added at the time that version 8.1 was first released. A numeric indicator, such
1 as a 1 or a 2, indicates that the information was added for the FixPak or level
1 ending in the same number. For example, a 1 indicates that the information
1 was added or changed in FixPak 1, a 2 indicates that the information was
1 changed for Version 8.1.2.

1 The *Data Links Manager Administration Guide and Reference* was updated in PDF
1 form (book number SC27-1221-01) at the time of FixPak 1, and is available for
1 download at the DB2 support site:

1 <http://www.ibm.com/software/data/db2/udb/winos2unix/support>

Additional resources:

Documentation for the DB2 Life Sciences Data Connect product is available
for download from the IBM software site:

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

If you want to view the DB2 documentation in HTML format, you can access
the DB2 HTML Information Center online from
<http://publib.boulder.ibm.com/infocenter/db2help/>. Alternatively, if you
want to install the DB2 HTML Information Center on your system, a *DB2
HTML Documentation* CD-ROM image is available for download from the
same site. Updates are made to the DB2 HTML documentation with every
release. For the latest documentation, access the DB2 HTML Information
Center online or download the *DB2 HTML Documentation* CD-ROM image for
installation on your system. PDF documentation is updated less frequently.

More information on the DB2 Development Center and DB2 for z/OS is
available at <http://www.ibm.com/software/data/db2/os390/spb/>.

For the latest information about the DB2 family of products, obtain a free
subscription to *DB2 Magazine*. The online edition of the magazine is available
at <http://www.db2mag.com>; instructions for requesting a subscription are
also posted on this site.

4 What's New for Version 8.1.4

4 This section provides a brief summary of the enhancements available with
4 DB2 Universal Database Version 8.1.4. Detailed information on these new
4 features can be found in the DB2 Information Center under **Release**
4 **information for Version 8.1.4**. If you do not have Version 8.1.4 installed, you
4 can access this information at
4 <http://publib.boulder.ibm.com/infocenter/db2help/>.

4 Application development enhancements

4 **KEEPDYNAMIC support for DB2 CLI/ODBC applications**

4 DB2 CLI/ODBC applications on the client can now take advantage of
4 the KEEPDYNAMIC feature on servers that are DB2 UDB for z/OS
4 and OS/390 Version 7 and later.

4 **DB2 CLI dynamic scrollable cursor support**

4 Dynamic scrollable cursors are now supported in DB2 CLI when
4 accessing servers which are DB2 UDB for z/OS Version 8.1 and later.

4 **CLI LOAD fullselect support**

4 The CLI LOAD utility has been enhanced to support a fullselect on
4 the INSERT statement. This feature allows you to load data from a
4 table and execute the INSERT statement once instead of multiple
4 times, as may be required when the VALUES clause is used.

4 **CLI/ODBC configuration keyword RETURNOLEDBCHARASWCHAR**

4 The IBM DB2 OLE DB Provider describes CHAR, VARCHAR, LONG
4 VARCHAR, and CLOB data as DBTYPE_WSTR by default as of DB2
4 UDB Version 8.1.2. The CLI/ODBC configuration keyword
4 OLEDBRETURNCHARASWCHAR allows you to change this default
4 to have the previously stated character data types reported as
4 DBTYPE_STR.

4 **Result set retrieval from SQL data-change operations**

4 In DB2 Version 8.1.4, you can use SELECT and SELECT INTO
4 statements to retrieve result sets from SQL data-change operations
4 (INSERT, UPDATE, and DELETE) embedded in the FROM clause.

4 **Searched INSERT, UPDATE, DELETE, and MERGE statements against fullselects**

4 As of DB2 Version 8.1.4 you can issue searched INSERT, UPDATE,
4 DELETE, and MERGE statements on the results of fullselects.

4 **Support of SQL statements that modify SQL data within SQL Table**
4 **Functions**

4 DB2 now supports the use of SQL statements that modify data within
4 the body of SQL table functions.

4 **EXECUTE privilege on routine is required to invoke a routine**

4 As of DB2 Version 8.1 routine level EXECUTE authorization exists to
4 ensure safe and manageable use of routines. The following
4 authorizations are required so that a user can successfully invoke a
4 routine (method, stored procedure, or user defined function). The
4 EXECUTE authority on all routines migrated to DB2 version 8.1 from
4 earlier releases has been GRANTED TO PUBLIC. Database
4 administrators can manage the EXECUTE privileges on these routines
4 with the GRANT and REVOKE commands.

4 Any user that will be an invoker of the routine must have at least one
4 of the following:

- 4 • EXECUTE privilege on the routine
- 4 • SYSADM or DBADM authority

4 The definer of a routine (the user who issues the CREATE statement
4 to register the routine) must have at least one of the following:

- 4 • EXECUTE privilege on the routine
- 4 • SYSADM or DBADM authority

4 Any user with EXECUTE WITH GRANT OPTION privilege on the
4 routine (this includes the routine definer unless the privilege has been
4 explicitly revoked), SYSADM or DBADM authority must explicitly
4 GRANT EXECUTE on a routine to PUBLIC or to the authorization IDs
4 of users that will be invokers of the routine. Caution should be taken
4 when granting EXECUTE privilege on a routine to PUBLIC as this
4 allows any database user to invoke the routine.

4 The owner of the package or packages associated with the routine, or
4 any user with SYSADM or DBADM authority must GRANT
4 EXECUTE on the packages to the routine definer. The package owner
4 is the user who executes the BIND command to bind a routine
4 package to the database (or the user specified in the OWNER
4 precompile/bind option if it was overridden). The package owner,
4 upon completion of a BIND, has EXECUTE WITH GRANT authority
4 on the bound package. When a routine is invoked the authorization to
4 run the routine is checked against the definers', and not the invokers',
4 authorization to EXECUTE the packages associated with the routine.
4 Thus, the definers' role is to encapsulate under one authorization ID
4 the privileges of running the packages associated with a routine and
4 to be a user capable of granting EXECUTE privileges to users that

4 need to invoke the routine. To determine which users will be invokers
4 of a routine, it is necessary to consider how a routine can be invoked.

4 Routines can be invoked from a command window or can be
4 embedded within an application program. In the case of methods and
4 UDFs the routine reference will be embedded in another SQL
4 statement whereas a procedure is invoked by using the CALL
4 statement. For dynamic SQL in an application, the invoker is the
4 runtime authorization ID of the immediately higher-level routine or
4 application containing the routine invocation (however, this ID can
4 also depend on the DYNAMICRULES option, with which the higher
4 level routine/application was bound). For static SQL, the invoker is
4 the value of the OWNER precompile/bind option of the package that
4 contains the reference to the routine. These users will require
4 EXECUTE authority on the routine.

4 If a user attempts to invoke a routine that the user is not authorized
4 to EXECUTE, an error (SQLSTATE 42501) is returned. This error is
4 also returned when the definer of the routine is missing EXECUTE
4 privilege on a package associated with the routine.

4 **Note:** For SQL routines the routine definer is also implicitly the
4 package owner. Therefore the definer will have EXECUTE
4 WITH GRANT OPTION on the routine and on the routine
4 package.

4 **External routines migrated to DB2 Version 8.1**

4 Prior to DB2 Version 8.1, the privilege to invoke a routine by a user
4 was managed by granting and revoking EXECUTE privileges on the
4 packages associated with the routine. Upon database migration the
4 following actions are taken:

- 4 • GRANT EXECUTE WITH GRANT OPTION to the routine
4 DEFINER of all existing routines
- 4 • GRANT EXECUTE to PUBLIC for all existing functions, methods,
4 and external stored procedures
- 4 • For each SQL procedure, GRANT EXECUTE to all users who
4 previously had EXECUTE privilege on the routine package

4 These actions ensure that all users that were previously able to invoke
4 a routine continue to be able to do so. Database administrators can
4 REVOKE the EXECUTE privilege of routines from specific database
4 users if they now wish to restrict the use of the routine.

4 The db2undgp command, can be used by database administrators to
4 both identify external stored procedures that access SQL and revoke
4 from all database users the EXECUTE privilege on these procedures.
4 This is useful for reverting all external stored procedures that access
4 SQL data to a state where no one can execute them. The DBADM or
4 SYSADM can then choose to GRANT EXECUTE on each routine to
4 the particular users that are intended to invoke the procedure.

4 DB2 family enhancements

4 Data Links Manager online help and error message functionality

4 This release enhances existing and adds new online help functions for
4 the Data Links File Manager (DLFM) component of Data Links
4 Manager.

4 XML Extender environment variable DB2DXX_MIN_TMPFILE_SIZE

4 DB2 XML Extender may place large documents in temporary files to
4 avoid using too much memory during processing. The environment
4 variable DB2DXX_MIN_TMPFILE_SIZE instructs XML Extender to use
4 memory buffers, rather than temporary files, for processing
4 documents smaller than the specified value.

4 Control Center for DB2 Universal Database for z/OS and OS/390 4 enhancements

4 The Control Center now supports the following features of DB2
4 Universal Database for z/OS and OS/390 Version 8.1:

- 4 • support of long names in the DB2 catalog
- 4 • an increase of the maximum number of partitions for table spaces
4 from 254 to 4096
- 4 • 64-bit virtual addressing support

4 Manageability enhancements

4 Resource management (db2nodes.cfg)

4 A new column has been added to the db2nodes.cfg file. The
4 resourcename column is for multiple logical node (MLN)
4 configurations and can give each node its own set of resources.

4 Referential integrity

4 The data definition language restrictions on referential constraints
4 have been lifted (some restrictions still apply), such that it is now
4 possible to have multiple referential constraints with different
4 referential integrity actions on the same table in a delete-connected
4 relationship.

4 Row migration UNION ALL views

4 DB2 UDB now offers complete transparency when you update rows

4 through a UNION ALL view. In addition to supporting update, delete,
4 or insert operations, DB2 can now move rows (during update
4 operations) between the branches of a UNION ALL view.

4 **Online import**

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

4 **sqluimpr API replaced by db2Import API**

4 The sqluimpr API has been replaced by db2Import API. The sqluimpr
4 API is still supported, but new applications should call the db2Import
4 API.

4 **Compressed backup**

4 Database backups can now be compressed to conserve disk space.
4 Backup compression has been added as an option to the BACKUP
4 DATABASE command and the db2Backup API. Support for the
4 restoring of compressed backups has been added to the RESTORE
4 DATABASE command and the db2Restore API.

4 **Setting buffer pool memory allocation size**

4 You can set the size for buffer pool memory allocations using the
4 DB2_ALLOCATION_SIZE registry variable. Setting this variable to a
4 higher value means that it will require fewer allocations to reach the
4 desired amount of memory that is allocated to a buffer pool.

4 **Performance enhancements**

4 **Range-clustered tables**

4 Range-clustered tables provide fast, direct access to data.

4 **Asymmetric index splitting**

4 The CREATE INDEX statement now has new clauses to allow greater
4 control of the space used when index pages are split after becoming
4 full.

4 **Temporary tables in SMS**

4 Storage for temporary tables in SMS table spaces will not be deleted
4 by default once they are no longer needed. In cases where temporary
4 tables are used repeatedly, this new feature avoids some of the
4 overhead of deleting and reserving storage for temporary tables.

4 **Page cleaning enhancements**

4 Starting in Version 8.1.4, there is an alternate method of configuring
4 page cleaning in your system. This alternate method differs from the
4 default behavior in that page cleaners behave more proactively in
4 choosing which dirty pages get written out at any given point in time.

4 This new method of page cleaning differs from the default page
4 cleaning method in two major ways:

- 4 1. Page cleaners do not respect the `chnpggs_thresh` configuration
4 parameter.

4 In this alternative method of page cleaning, page cleaners will no
4 longer react in response to value of the `chnpggs_thresh`
4 configuration parameter. Instead of attempting to keep some
4 percentage of the buffer pool clean, the alternate method of page
4 cleaning provides a mechanism whereby the agents are informed
4 of the location of good victim pages that have just been written
4 out, so that agents do not have to search the buffer pool to look
4 for a victim. When the number of good victim pages drops below
4 an acceptable value, the page cleaners are triggered, and proceed
4 to search the entire buffer pool, writing out potential victim pages,
4 and informing the agents of the location of these pages.

- 4 2. Page cleaners no longer respond to LSN gap triggers issued by the
4 logger.

4 When the amount of log space encompassing the log record which
4 has updated the oldest page in the buffer pool and the current log
4 position exceeds that allowed by the `softmax` parameter, it is said
4 that the database is in an 'LSN gap' situation. Under the default
4 method of page cleaning, when the logger detects than an LSN
4 gap has occurred, it will trigger the page cleaners to write out all
4 the pages which are contributing to the LSN gap situation. That is,
4 it will write out those pages which are older than what is allowed
4 by the `softmax` parameter. Page cleaners will be idle for some
4 period of time while no LSN gap is occurring. Then, once an LSN
4 gap occurs, the page cleaners are activated to write a large number
4 of pages before going back to sleep. This can result in the
4 saturation of the I/O subsystem, which then affects other agents
4 which are reading or writing pages.

4 The alternate method of page cleaning will modulate this behavior
4 by spreading out the same number of writes over a greater period
4 of time. The cleaners do this by proactively determining not only
4 which pages are currently in an LSN gap, but also which pages
4 will be in the LSN gap some period in the future based on the
4 current level of activity.

4 To use the new method of page cleaning, set the
4 `DB2_USE_ALTERNATE_PAGE_CLEANING` registry variable to "ON".

4 **Monitoring network time**

4 New monitor elements for monitoring network time have been
4 introduced. These elements are used to get a better idea of the
4 database activity and network traffic at the database or application
4 level.

db2dml_upd_hostname command

In a Data Links Manager environment, the DB2 host name information is stored on the Data Links Manager server, and the Data Links Manager host name information is stored on the DB2 server in a configuration file. In order to update the DB2 or DLM host name, you can invoke the db2dml_upd_hostname - Data Links Update Host Name command.

Lock deferral

To improve concurrency, DB2 now permits the deferral of row locks for CS or RS isolation scans in some situations until a record is known to satisfy the predicates of a query.

Direct I/O support on AIX

Currently Direct I/O is already supported on Windows via registry variable DB2NTNOCACHE. Now, limited Direct I/O support has been added to AIX. This new support is for all SMS containers except for long fields, LOB, and temporary table spaces.

Replication enhancements

Object name length

Replication now supports schema and table names of up to 128 bytes on z/OS when DB2 UDB for z/OS Version 8 is run in new-function mode.

Password file

The asnpwd command allows you to list the aliases and user IDs contained in the password file. You can also use the encrypt parameter of the asnpwd command to encrypt either all of the entries in a file or just the password entry in a file.

Data on Oracle sources

The Apply program no longer needs to issue lock table statements for CCD tables on Oracle sources. To take advantage of this improvement, you must migrate any existing registrations and subscriptions for Oracle sources.

Known problems and workarounds (Version 8.1.4)

The following are the currently known limitations, problems, and workarounds for DB2® Universal Database Version 8.1.4. The information in this section applies only to the Version 8.1.4 release of DB2 Universal Database™, and its supporting products. Any limitations and restrictions might or might not apply to other releases of the product.

Product and product-level support

3 **Alternate FixPaks on Linux and UNIX-based systems**

3 Prior to DB2 Universal Database Version 8, FixPaks only functioned as
3 updates to installed DB2 Universal Database packages or filesets in one fixed
3 location. Essentially, this meant that the installation of FixPaks would replace
3 existing files with the updated ones provided within the FixPaks, and
3 multiple DB2 FixPak levels on a single system was not a possibility. DB2
3 Universal Database Version 8.1 Enterprise Server Edition (ESE) with multiple
3 FixPak levels can now coexist in the same system. This feature, supported in
3 production operating environments since Version 8.1.2, is achieved using the
3 following two FixPak types:

3 **regular FixPaks**

- 3 • Available not only for ESE, but for all supported DB2 V8.1 products
3 for the related platforms
- 3 • Can be installed directly on top of the existing installation either in
3 /usr/opt/db2_08_01 on AIX® or /opt/IBM/db2/V8.1 on other
3 platforms

3 **alternate FixPaks**

- 3 • Can be installed as a completely new copy of DB2 Universal
3 Database ESE
- 3 • Installed in a predefined location other than the location used for a
3 regular DB2 Universal Database installation

3 **Notes:**

- 3 1. You are *not* required to perform a multiple FixPak installation if you do
3 not feel it is necessary for your environment.
- 3 2. Starting with IBM DB2 Universal Database Enterprise Server Edition (ESE)
3 for Linux and UNIX®, Version 8.1.2, FixPaks are supported in production
3 operating environments when installed as Multiple FixPaks.

3 To update a multiple FixPak instance to a different FixPak level, perform one
3 of the following options:

- 3 • Install the appropriate regular FixPak on the General Availability (GA)
3 installation, and update the instance by running db2iupdt from the existing
3 GA path.
- 3 • Install the appropriate alternate FixPak to its own unique path, and update
3 the instance by running db2iupdt from this path.

3 For further information regarding downloading alternate FixPaks, visit the
3 IBM® support site at
3 <http://www.ibm.com/software/data/db2/udb/winos2unix/support>.

Back-level DB2 Universal Database server support

If you migrate your DB2 Universal Database client systems to version 8 before you have migrated all of your DB2 Universal Database servers to version 8, several restrictions and limitations apply.

For version 8 clients to work with version 7 servers, you need to configure and enable the use of DRDA® application server capability on the version 7 server. For information on how to do this, refer to the version 7 *Installation and Configuration Supplement*. You cannot access a DB2 Connect™ Version 7 server from a DB2 Universal Database Version 8 client.

When accessing version 7 servers from version 8 clients, there is no support available for:

- The following data types:
 - Large object (LOB) data types
 - User-defined distinct types
 - DATALINK data types
The DATALINK data type allows you to manage external data found in nonrelational storage. The DATALINK data type references files that physically reside on file systems external to DB2 Universal Database.
- The following security capabilities:
 - Authentication type SERVER_ENCRYPT
SERVER_ENCRYPT is a method to encrypt a password. The encrypted password is used with the user ID to authenticate the user.
 - Changing passwords
You are not able to change passwords on the version 7 server from a version 8 client.
- The following connections and communication protocols:
 - Instance requests that require an ATTACH instead of a connection
ATTACH is not supported from a version 8 client to a version 7 server.

- Network protocols other than TCP/IP.
(SNA, NetBIOS, IPX/SPX, and others)
- The following application features and tasks:
 - The DESCRIBE INPUT statement for all applications except ODBC/JDBC
To support version 8 clients running ODBC/JDBC applications that access version 7 servers, apply a fix for DESCRIBE INPUT support to all version 7 servers where this type of access is required. This fix is associated with APAR IY30655. Use the “Contacting IBM” information in any DB2 Universal Database documentation set (PDF or HTML) to find out how to get the fix associated with APAR IY30655.
The DESCRIBE INPUT statement is a performance and usability enhancement that allows an application requestor to obtain a description of input parameter markers in a prepared statement. For a CALL statement, this includes the parameter markers associated with the IN and INOUT parameters for the stored procedure.
 - Two-phase commit
A version 7 server cannot be used as a transaction manager database when using coordinated transactions that involve version 8 clients, nor can a version 7 server participate in a coordinated transaction where a version 8 server may be the transaction manager database.
 - XA-compliant transaction managers
An application using a version 8 client cannot use a version 7 server as an XA resource. This includes WebSphere[®], Microsoft[®] COM+/MTS, BEA WebLogic, and others that are part of a transaction management arrangement.
 - Monitoring
 - Utilities that can be initiated by a client to a server
 - SQL statements greater than 32 KB in size

There are similar limitations and restrictions for version 8 tools working with version 7 servers.

The following version 8 GUI tools, products, and Centers support only version 8 servers:

- Control Center
- Development Center
- Health Center (including the Web version of this center)
- Indoubt Transaction Manager
- Information Catalog Center (including the Web version of this center)
- Journal
- License Center

1

- Satellite Administration Center
- Spatial Extender
- Task Center
- Tools Settings

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

- Command Center (including the Web version of this center)
 - Saving, importing, and scheduling of scripts are not supported by the Command Center.
- Data Warehouse Center
- Replication Center
- The import/export configuration file function of the Configuration Assistant
- SQL Assist
- Visual Explain

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

Classic Connect unavailable

The Classic Connect product is *not* available. Although you may find references to the Classic Connect product in the Data Warehouse documentation and elsewhere, these should be ignored, as they are no longer applicable.

Data Warehouse Center down-level server support restrictions

The following limitations exist for down-level server support for DB2 Universal Database (DB2 UDB) Enterprise Server Edition Version 8 Data Warehouse Center:

Large Object (LOB) support

- If you are using a warehouse control database on a server that is earlier than DB2 UDB Enterprise Server Edition Version 8 you will not be able to work with LOBs. You must upgrade the warehouse control database to the correct level, or move the control database to the system on which the DB2 UDB Enterprise Server Edition Version 8 warehouse server is installed and use it locally from that system.
- To move LOBs between the Data Warehouse Center and DB2, you must upgrade to DB2 UDB Enterprise Server Edition Version 8.

Systems Network Architecture (SNA) support

If you use SNA to connect to your warehouse sources and targets, you must change the configuration to be TCP/IP over SNA, or use the Windows NT[®] warehouse agent.

Support for EXPORT and LOAD utilities

When you upgrade your warehouse agent, you must also upgrade your source target databases, or replace the EXPORT and LOAD utilities in your warehouse processes with SQL Select and Insert steps. SQL Select and Insert steps use a DELETE* command followed by SELECT and INSERT commands. SQL Select and Insert steps require the database to log all transactions. As a result, the performance for SQL Select and Insert steps is not as efficient as it is for EXPORT and LOAD utilities.

DB2 license policy for DB2 Universal Database Workgroup Server Edition

Although the *DB2 Quick Beginnings for Servers* book and the online tools help for the License Center indicate otherwise, the Internet license policy is *not* valid for DB2 Universal Database Workgroup Server Edition. If you require a license for Internet users, you need to purchase DB2 Universal Database Workgroup Server Unlimited Edition.

DB2 Universal Database Version 7 server access

To access a DB2 Universal Database Version 7 server on a Linux, UNIX, or Windows[®] operating system from a version 8 client, you must have version 7 FixPak 8 or later installed on your server and have run the **db2updv7** command. For instructions on installing the version 7 FixPaks, refer to the version 7 FixPak Readme and Release Notes.

You cannot access a DB2 Connect Version 7 server from a DB2 Universal Database Version 8 client.

Data Warehouse Center not available in Simplified Chinese

The Data Warehouse Center is not available in Simplified Chinese. Therefore, the following dependent DB2 components are also unavailable in this environment:

- The DB2 Warehouse Manager.
- The DB2 Warehouse Manager Connector for the Web and the DB2 Warehouse Manager Connector for SAP, which are dependent on the installation of the DB2 Warehouse Manager.
- The Information Catalog Center, which is dependent on the Manage Information Catalog wizard component of DB2 Warehouse Manager.

DB2 Web Tools

The application servers supported by DB2 Web Tools for the following languages are required to be compliant with the Servlet 2.3 specification:

- Japanese
- Korean
- Simplified Chinese
- Traditional Chinese
- Russian
- Polish

Development Center APARs required for SQLJ and SQL Assist support on DB2 UDB for OS/390, Version 6 and DB2 UDB for z/OS, Version 7

When using the Development Center on an Application Development client for DB2 Universal Database Version 8 on Windows or UNIX operating systems, the following APARs need to be installed on the server to enable SQLJ and SQL Assist support:

DB2 UDB for z/OS, Version 7

- PQ65125 - Provides SQLJ support for building Java SQLJ stored procedures
- PQ62695 - Provides SQL Assist support

DB2 UDB for OS/390®, Version 6

- PQ62695 - Provides SQL Assist support

Development Center limitations for 64-bit operating systems

Debugging of Java stored procedures against a 64-bit server is not supported by the Development Center. Debugging SQL stored procedures is supported on 64-bit Windows operating systems. OLE DB and XML are not supported on 64-bit servers.

Development Center now supports Actual Cost information for SQL statements run on OS/390 or z/OS servers

The DB2 Development Center now provides Actual Cost information for SQL statements run on DB2 Universal Database for OS/390 and z/OS, Version 6 and Version 7.. The following Actual Cost information is provided:

- CPU time
- CPU time in external format
- CPU time as an integer in hundredths of a second
- Latch/lock contention wait time in external format
- Number of getpages in integer format
- Number of read i/o in integer format
- Number of write i/o in integer format

This functionality also allows you to view multiple sets of Actual Cost results for a single SQL statement with different host variable values.

1 Actual Cost information is available in the Development Center from the
1 Create SQL Stored Procedure and Create Java Stored Procedure wizards,
1 within the SQL Statement window for OS/390 and z/OS connections. To use
1 the Actual Cost functionality, click the **Actual Cost** button in the OS/390 and
1 z/OS connection window in either of the Stored Procedure wizards. To use
1 Actual Cost functionality, you must have the Stored Procedure Monitor
1 Program (DSNWSPM) installed on your DB2 OS/390 server .

2 **Development Center on the Linux operating system**

2 You cannot use the Development Center to debug Java™ stored procedures
2 running on any of the Linux distributions (32-bit, 64-bit, Intel, zSeries, or
2 iSeries).

Federated systems restrictions

For users of DB2 Universal Database for Linux, UNIX, and Windows Version 7.2 federated databases:

To create nicknames for tables and views in DB2 Universal Database (DB2 UDB) for UNIX and Windows, Version 8, you need to apply DB2 UDB for UNIX and Windows Version 7.2 FixPak 8 on your DB2 UDB for UNIX and Windows Version 7.2 federated database. If you do not apply FixPak 8 on your DB2 UDB for UNIX and Windows Version 7.2 federated database, an error will occur when you access the nicknames.

LONG VARCHAR and LONG VARGRAPHIC support:

The federated documentation indicates that the LONG VARCHAR and LONG VARGRAPHIC data types used by the DB2 family of products are unsupported. This is not entirely accurate. You can create nicknames for DB2 Universal Database for Linux, UNIX, and Windows data source objects that contain LONG VARCHAR and LONG VARGRAPHIC data type columns. These remote columns will be mapped to DB2 Universal Database for Linux, UNIX, and Windows LOB data types. For the other DB2 family of products, you can create a view that omits or recasts these data types and then create a nickname for the view.

WITH HOLD cursors:

1 You can use the WITH HOLD semantics on a cursor that is defined on a
1 nickname or in a PASSTHRU session. However, you will receive an error if
1 you try to use the semantics (with a COMMIT) and the data source does not
1 support the WITH HOLD semantics.

Data sources:

3 Previously unsupported data sources are supported as of Version 8.1.2 using
3 DB2 Information Integrator relational and nonrelational wrappers. For a full
3 list of supported data sources, see the DB2 Information Center for DB2
3 Information Integrator: Product Overview -> Federated systems - overview ->
3 Data sources -> Supported data sources

Support for DB2 Universal Database Server for VM and VSE:

Federated support for DB2 Universal Database Server for VM and VSE has been added in Version 8.1.2.

Product support:

3 Previously unsupported products are now supported through Information
3 Integrator:

- 3 • DB2 Relational Connect is supported using DB2 Information Integrator
3 relational wrappers.
- 3 • DB2 Life Sciences Data Connect is supported using DB2 Information
3 Integrator non relational wrappers.

Unsupported operating systems:

Federated systems are not supported on the Windows ME operating system.

Setting up the federated server to access data sources:

3 The COMPACT installation option does not install the necessary access to the
3 DB2 family or Informix™ data sources. You must use the TYPICAL or
3 CUSTOM installation to access the DB2 family data sources. The CUSTOM
3 installation option is the only option you can use to install access to both the
3 DB2 family and the Informix data sources.

Update federated database in order to create wrapper:

3 If you are using a federated database with DB2 Universal Database (DB2
3 UDB) Version 8.1.2 or later that was created using DB2 UDB Version 8.1 or
3 DB2 UDB Version 8.1 FixPak 1, you must update your federated database
3 using the db2updv8 command.

Syntax:

3 **db2updv8 -d database-name** -u userid -p password

3 If you do not update your database to Version 8.1.2, you will get one of the
3 following error messages when you try to create a wrapper from the
3 Federated Database Objects folder in the Control Center:

- 3 • java.lang.NullPointerException
- 3 • [IBM][CLI Driver][DB2/NT] SQL0444N Routine "GET_WRAP_CFG_C"
3 (specific name "SQL030325095829810") is implemented with code in
3 library or path "\GET_WRAP_CFG_C", function "GET_WRAP_CFG_C"
3 which cannot be accessed. Reason code: "4". SQLSTATE=42724

3 **Cataloging DB2 family data sources in the federated system database 3 directory:**

When the name of the remote database is more than 8 characters, it is
necessary to create a database connection services (DCS) directory entry.

An example of cataloging an entry in the DCS directory for the database
using the CATALOG DCS DATABASE command is:

```
CATALOG DCS DATABASE SALES400 AS SALES_DB2DB400
```

where:

SALES400

Is the name of the remote database you entered in the CATALOG
DATABASE command.

AS SALES_DB2DB400

Is the name of the target host database that you want to catalog.

High availability function is included in DB2 Universal Database Workgroup Server Edition

Though it is not mentioned explicitly in the topic DB2 Workgroup Server
Edition, the high availability function of the DB2 Universal Database
Enterprise Server Edition that is referred to in the topic DB2 Enterprise Server
Edition is included in the DB2 Universal Database Workgroup Server Edition.

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

2 The IBM DB2 Development Add-In for Microsoft Visual Studio .NET does not
2 support the following:

- 2 • DB2 Universal Database for z/OS and OS/390, Version 8
- 2 • DB2 Universal Database for iSeries, Version 8

3 **Installation on AIX**

3 If the db2setup program is run from a directory whose path includes a blank,
3 the setup will fail with the following error:

```
3 <file>: not found
```

3 Place the installable image in a directory whose path does not include spaces.

Installation on Linux

When you install version 8.1 of DB2 Universal Database on Linux, the RPM-based installation attempts to install the IBM Java RPM (IBMJava2-SDK-1.3.1.-2.0.i386.rpm). If a later level of the RPM (such as IBMJava2-SDK-1.4.0.-2.0.i386.rpm) already exists, the back-level RPM is not installed.

However, in this case, the installation leaves the JDK_PATH database configuration parameter pointing to the Java 1.3 path, /opt/IBMJava2-14/. As a result, none of the Java-dependant functionality, including the installation of the DB2 Tools Catalog, will work.

To solve this problem, run the following command as the instance owner:

```
db2 update dbm cfg using JDK_PATH /opt/IBMJava2-14
```

This will point DB2 Universal Database to the correct IBM Developer Kit.

License Center back-level versions not supported

If a version 7 License Center attempts a connection to a version 8 server, the License Center receives an "SQL1650 - Function not supported" error message indicating that the connection is not supported.

Microsoft Visual Studio, Visual C++

Despite being mentioned in the DB2 Development Center online help as a possible solution for the Build not successful: -1 error, Microsoft Visual Studio Visual C++ Version 5.0 is not supported for the development of SQL stored procedures. However, Microsoft Visual Studio Visual C++ Version 6.0 is supported. Additional configuration information is available in the *IBM DB2 Application Development Guide: Building and Running Applications*.

Microsoft XP fix needed on 64-bit operating systems

If you are using the Microsoft XP operating system (2600) configured to use the NETBIOS protocol with the DB2 family of products, you need to obtain a hotfix from Microsoft. Contact Microsoft with the Knowledge Base article number Q317437.

MVS operating system not supported

Despite being mentioned in the documentation, the MVS™ operating system is no longer supported by DB2 Universal Database. MVS has been replaced with z/OS.

SNA support limitations in Version 8

The following support has been withdrawn from DB2 Universal Database Enterprise Server Edition (ESE) for Windows and UNIX-based operating systems, Version 8, and from DB2 Connect Enterprise Edition (CEE) for Windows and UNIX-based operating systems, Version 8:

- Multisite update (two-phase commit) capability using SNA cannot be used. Applications that require multisite update (two-phase commit) *must* use TCP/IP connectivity. Multisite update (two-phase commit) using TCP/IP to a host or iSeries™ database server has been available for several releases. Host or iSeries applications that require multisite update (two-phase commit) support can use the new capability of TCP/IP multisite update (two-phase commit) support within DB2 Universal Database ESE, Version 8.
- DB2 Universal Database ESE or DB2 CEE servers no longer accept client connections using SNA. As of version 8, FixPak 1, DB2 Universal Database allows the 32-bit version of AIX, Solaris™ Operating Environment, HP-UX, and Windows-based applications to access host- or iSeries-based database servers using SNA. This support will allow applications access to host or iSeries database servers using SNA, but using only one-phase commit.
- Sysplex support with DB2 Universal Database for z/OS™ is only available using TCP/IP. Sysplex support is not provided using SNA connectivity.
- Change password support is no longer available using SNA connectivity to host database servers.
- All SNA support will be withdrawn in the next version of DB2 Universal Database and DB2 Connect.

Supported LDAP client and server configurations

The following table summarizes the supported LDAP client and server configurations:

Table 1. Supported LDAP client and server configurations

	IBM SecureWay Directory	Microsoft Active Directory	Netscape LDAP server
IBM LDAP client	Supported	Supported	Supported
Microsoft LDAP/ADSI client	Supported	Supported	Supported

IBM SecureWay Directory Version 3.1 is an LDAP Version 3 server available for Windows NT, Windows 2000, and Windows 2003, AIX, and Solaris. SecureWay Directory is shipped as part of the base operating system on AIX and iSeries (AS/400), and with OS/390 Security Server.

DB2 supports IBM LDAP client on AIX, Solaris, Windows 98, Windows XP, Windows NT, Windows 2000, and Windows 2003.

DB2 supports IBM LDAP Version 3.2.2 on Linux IA32 and Linux/390.

Microsoft Active Directory is an LDAP Version 3 server and is available as part of the Windows 2000 Server operating system.

4 The Microsoft LDAP client is included with the Windows operating system.

4 When running on Windows operating systems, DB2 supports using either the
4 IBM LDAP client or the Microsoft LDAP client to access the IBM SecureWay
4 Directory Server. To explicitly select the IBM LDAP client, use the **db2set**
4 command to set the DB2LDAP_CLIENT_PROVIDER registry variable to the
4 value IBM.

4 **Tivoli Storage Manager for LinuxAMD64**

4 Tivoli Storage Manager support is now available for Linux on AMD Opteron -
4 64-bit systems. The minimum required level of TSM client API is TSM 5.2.0.

4 **Windows XP operating systems**

2 The Windows XP Professional operating system is only supported by Personal
2 Edition and Workgroup Server Edition products. The Windows XP Home
2 Edition operating system is only supported by Personal Edition products.

Application development

Asynchronous execution of CLI

Asynchronous execution with CLI is not available.

2 **CLI and ODBC on Windows 64-bit operating systems**

2 You cannot use an application with mixed ODBC and DB2 CLI on a Windows
2 64-bit operating system.

Configuration Assistant

Unsupported bind options

The Configuration Assistant does not support the following bind options:

- CALL RESOLUTION
- CLIPKG
- CNULREQD
- DBPROTOCOL
- ENCODING
- MESSAGES
- OPTHINT
- OS400NAMING
- GENERIC
- IMMEDIATE
- KEEP DYNAMIC
- PATH

- SORTSEQ
- TRANSFORM_GROUP
- VALIDATE
- VARS

Configuration parameters

2 NUM_LOG_SPAN configuration parameter on a multipartition database

2 The NUM_LOG_SPAN specifies the maximum number of log files that a
 2 transaction can span. If at any time a transaction violates the
 2 NUM_LOG_SPAN setting, it is rolled back, and forces the application that
 2 caused the transaction off the database.

2 However, in a multipartition system, the db2loggr process can only force an
 2 application if the coordinating node for the application and the node of the
 2 db2loggr process that detects the error are both the same. For example, you
 2 have a system with three nodes (0,1 and 2) and have set the
 2 NUM_LOG_SPAN parameter to 2 on all nodes. An application connects to
 2 node 2 of the database, and begins a long-running transaction, which spans
 2 more than two log files. If the db2loggr process on node 1 is the first to detect
 2 this error, then nothing will happen. However, if the violation also occurs on
 2 node 2, then the db2loggr process will notice the error, the transaction will be
 2 rolled back, and the application will be forced off.

dasdrop limitation in multiple FixPak environments

Alternate FixPaks install their own version of the **dasdrop** command. On AIX, it is installed in the `/usr/opt/db2_08_FPn/` path. On other UNIX systems, it is installed in the `/opt/IBM/db2/V8.FPn/` path. In both cases, *n* is the number of the FixPak.

In a multiple FixPak Environment, you can only have one DAS set up at any one time. You can create the DAS against version 8.1 of the product, or against any of the alternate FixPaks. To drop a DAS that was created against version 8.1 of the product, you can use any version of **dasdrop** to drop it. However, to drop a DAS that was created against an alternate FixPak, you must use an alternate FixPak version of **dasdrop**.

For example, consider the following scenario on the AIX operating system:

- You install DB2 Universal Database, Version 8.1.
- You install alternate FixPak 1.
- You create a DAS using the version 8.1 code, with the following command:
`/usr/opt/db2_08_01/instance/dascrt dasusr1`
- You want to drop the DAS.

You can drop this DAS by using either of the following commands:

```
/usr/opt/db2_08_01/instance/dasdrop
```

```
/usr/opt/db2_08_FP1/instance/dasdrop
```

Both will work properly.

However, in this example:

- You install DB2 Universal Database, Version 8.1.
- You install the alternate FixPak 1.
- You create a DAS using the alternate FixPak 1 code, with the following command:

```
/usr/opt/db2_08_FP1/instance/dascrt dasusr1
```

- You want to drop this DAS.

You must use the alternate FixPak 1 **dasdrop** command:

```
/usr/opt/db2_08_FP1/instance/dasdrop
```

Trying to use the version 8.1 **dasdrop** command will cause an error.

This limitation only applies to version 8.1 of the product, and not to any of the regular FixPaks. For example:

- You install DB2 Universal Database, Version 8.1.
- You apply regular FixPak 1, which corrects the problem with the version 8.1 **dasdrop**.
- You install the alternate FixPak 1.
- You create a DAS using the alternate FixPak 1 code, with the following command:

```
/usr/opt/db2_08_FP1/instance/dascrt dasusr1
```

- You want to drop this DAS.

You can drop this DAS by using either of the following commands:

```
/usr/opt/db2_08_01/instance/dasdrop
```

```
/usr/opt/db2_08_FP1/instance/dasdrop
```

Both will work properly because the version of dasdrop in the `/usr/opt/db2_08_01/` path was corrected when you applied the regular FixPak.

Data Warehouse Center

ERwin 4.x metadata bridge

ERwin 4.0 metadata cannot be imported on the Linux operating system.

The ERwin 4.x bridge is supported on Windows 98 and WinME with the following restrictions:

- The **db2erwinimport** command can only be run from the db2 command line processor.
- You must fully qualify the XML and trace file names for the -x and -t parameters.

Japanese names of remote objects

Schema, table, and column names of remote sources in Japanese cannot contain certain characters. Unicode mapping differences may cause the names to be null. See <http://www.ingrid.org/java/i18n/encoding/ja-conv.html> for more information.

Restrictions for the Clean Data transformer

Linking restrictions:

You cannot link any OS/390 data resources such as tables or views with a new Clean Data step. You can still link OS/390 data resources with the deprecated programs Clean Data step.

Parameter restrictions:

For the Find and Replace parameter: If your new Clean Data transformer rules table contains different data types for the Find and Replace columns, you must change the Target column data type in both the Target Table properties page and the Column Mapping page before you promote the transformer to test mode.

For the Discretize parameter: If your new Clean Data rules table contains different data types for the Bounds and Replace columns, you must change the Target column data type in both the Target Table properties page and the Column Mapping page before you promote the transformer to test mode.

iSeries platform restrictions:

On the iSeries platform, the new Clean Data transformer does not perform error processing. You can generate the All Matches match type only on the iSeries platform.

Using the warehouse agent for replication and accessing Client Connect warehouse sources

Using the warehouse agent for replication

If the source, target, capture control, or apply control servers (databases) are remote to the client system, you must catalog the database with the same

name, user ID, and password on both the client and the warehouse agent systems. After you catalog the source on the client and warehouse agent systems, verify that you can connect to the source, target, capture, and apply databases.

If you cannot connect to the warehouse source, warehouse target, replication capture, or replication apply databases, verify that the environment variable DB2COMM on the remote system is set to TCP/IP, and that the port number matches the port number of the node that is cataloged on the client system.

To verify the port number on the remote system, type the following command at a DB2 Universal Database command prompt:

```
get dbm cfg | grep SVCENAME
```

You specify the port number of the client system when you catalog the node.

Accessing Client Connect warehouse sources using the warehouse agent

When you access a warehouse source that was defined using Client Connect with a warehouse agent, the source must be cataloged with the same name, user ID, and password on both the client and warehouse agent systems. If you are using the ODBC version of the warehouse agent, you must also catalog the source as an ODBC source on both the warehouse agent and client sites. Otherwise, actions that require the warehouse agent to access the warehouse source will fail.

Scheduling a warehouse process to run at intervals

When scheduling a warehouse process to run at intervals, you must determine the longest time it takes to run all the production steps in the process and schedule the intervals accordingly. If a process exceeds the scheduled time interval, all subsequently scheduled occurrences of that process will not run and will not be rescheduled.

Replication Center restrictions on iSeries systems

Administrative tasks in IASPs:

When using the Replication Center, you cannot perform administrative tasks in IASPs on iSeries systems.

Restrictions for replication steps that use iSeries control, source and target servers:

The iSeries control, source, and target servers are supported only on the DB2 Universal Database Enterprise Server Edition.

2 For both default and remote agents, the iSeries servers must be cataloged on
2 the local machine. For a remote agent, the iSeries servers must also be
2 cataloged on the machine where the agent resides. If the source or the target
2 servers are on an iSeries operating system, then you must specify the system
2 name on the Database page of the Source or Target notebook.

3 **Restriction for import and export**

3 If a process with unlinked shortcuts is exported and then imported as a .tag
3 file into another control database, the unlinked shortcut data will cause error
3 DWC3142:

3 <dirID> was not found in the Data Warehouse Center control database.

3 This error displays when the unlinked shortcut dirIDs are untranslated and
3 they refer back to the original control database.

4 **The Visual Warehouse 5.2 DB2 program "VW 5.2 Load flat file into DB2 4 UDB EEE (AIX only)" is not supported**

4 The Visual Warehouse 5.2 DB2 EEE load step is not supported in DB2 Version
4 8. To load a delimited file into a partitioned table in version 8, perform the
4 following steps:

- 4 1. Migrate the target database or table to DB2 Version 8 if it is not already
4 at DB2 Version 8. One way to migrate the database is from the command
4 line with the **db2move** command.
- 4 2. Open the Properties window of the Visual Warehouse 5.2 EEE load step,
4 select the **Parameters** tab, and note the values for **Column delimiter**,
4 **String delimiter**, and **Decimal delimiter** parameters.
- 4 3. Create a new process for the new load step or use the original process. If
4 you plan to create a new process, add the source file and target table to
4 the new process. Also add the new process to the appropriate warehouse
4 security group.
- 4 4. Create a DB2 load step in the process that you are using.
- 4 5. Connect the source and target to the step.
- 4 6. Open the Properties window of the load step and select **PARTITIONED**
4 in the **Load mode** field. If necessary, update the **Column**, **Character**
4 **strings**, and **Decimal point** fields with the values used in the **Column**
4 **delimiter**, **String delimiter**, and **Decimal delimiter** of the old step.
- 4 7. Click **Advanced** to start the Load wizard. From the Operation page,
4 select **Split and load data**.
- 4 8. From the Type page, select **Replace table data**.
- 4 9. Accept the Load wizard's remaining default values.
- 4 10. The Summary page shows the final load command. Review the final load
4 command and click **Finish**.
- 4 11. Close the Properties window.

Limited support for CURSOR load

The DB2 UDB Load step now allows a view or a table to be used as the source to the step, resulting in a LOAD FROM CURSOR.

In order to map columns in the wizard for CURSOR load, the **Map columns based on column positions found in the input file** radio button must be selected.

DB2 Cube Views

Cube Views sample application differs from sample source

DB2 Cube Views V8.1 provides a sample application called db2mdapiclient.exe that demonstrates some important functions of Cube Views. The source code for this application, db2mdapiclient.cpp, is included. In the initial release of DB2 Cube Views V8.1, the application was built using the source code. In version 8.1.4 the source code is unchanged, but the application has been modified. The source code is still valid but it does not exactly match the application.

DB2 Data Links Manager

Data Links server backup fails using Tivoli Storage Manager archive server (AIX, Solaris Operating Environment)

Problem: When installing or migrating to DB2 Data Links Manager, Version 8.1, a Data Links File Manager (DLFM)-initiated backup of Data Links server data to a Tivoli® Storage Manager archive server fails. One of the following sets of error messages display, either on-screen or on the installation status report:

DLFM129I: Automatic backup of DLFM_DB database has been triggered.
Please wait for the backup to complete.

DLFM901E: A system error occurred. Return code = "-2062".
The current command cannot be processed.
Refer to the db2diag.log file for additional information.

— or —

DLFM811E: The current DLFM database could not be backed up.
SQL code = "-2062", Return code = "-2062"

DLFM901E: A system error occurred. Return code = "-2062".
The current command cannot be processed.
Refer to the db2diag.log file for additional information.

Cause: The DB2 Data Links Manager Installer program failed to set the variables required to use Tivoli Storage Manager as an archive (backup) server for a Data Links server.

Tip: If you want to use Tivoli Storage Manager as an archive server, and you have not yet installed or migrated to DB2 Data Links Manager, Version 8.1, you can prevent this problem from occurring. First, do not use the "Tivoli Storage Manager" backup option from the Installer program. Then, manually configure the Data Links Manager Administrator's profile to include the appropriate Tivoli Storage Manager variables, as described in step 2 that follows. After you have completed both tasks, you can proceed with the installation or migration.

Workaround: Perform the following tasks in the order listed.

1. Back up the DLFM database using this command: `db2 backup <dl_fm_db><path>` where:
 - `<dl_fm_db>` is the name of the DLFM database. By default, the database is called `DLFM_DB`.
 - `<path>` is the directory path to the backup storage location of your choice.
2. Manually configure the Data Links Manager Administrator's profile to include the appropriate Tivoli Storage Manager variables. The manual configuration procedure and the required variables are described in the following documentation topics:
 - Using Tivoli Storage Manager as an archive server (AIX)
 - Using Tivoli Storage Manager as an archive server (Solaris Operating Environment)

You can find these topics either online in the DB2 Information Center, or in the "System Management Options" chapter of the *DB2 Data Links Manager Administration Guide and Reference*.

- If you are completing a new installation of DB2 Data Links Manager, Version 8.1, you are finished.
- If you are migrating to DB2 Data Links Manager Version 8.1, rerun the Migration Utility program, **db2dlmmg**.

DB2 Universal Database backup and restore

Backup and restore on Linux 390 operating systems

Backup and restore operations to and from multiple tape devices may not work if you are using the Linux 390 operating system.

DB2 Universal Database replication

Java Administrative API documentation for DB2 data replication

If you are developing applications using the administrative functions available in DB2 DataPropagator, you can obtain documentation for the relevant administrative Java APIs by contacting IBM support.

Column mapping restrictions and the Replication Center

You cannot map an expression in a source table to a key column in a target table if the TARGET_KEY_CHG column of the IBMSNAP_SUBS_MEMBER table is "Y" for that target table. This means that when using the Replication Center to create a subscription-set member, you should not select the option **Let the Apply program use before-image values to update target-key columns** if a key column in the target table is mapped to an expression in the source table.

Development Center

4 Debugging Stored Procedures with double quotation marks

4 The IBM DB2 Universal Database Development Center 8.1.4 and all previous
4 releases does not support debugging for any stored procedure with double
4 quotation marks (") in the stored procedure name, schema, or specific name.

4 SQLFLAG(STD) precompiler option error

4 Remove the SQLFLAG (STD) precompile option when using Development
4 Center to create SQL stored procedure to run on DB2 for z/OS, Version 8. If
4 the SQLFLAG(STD) precompile option is enabled it will cause the following
4 error: Abend C6 occurred while running Precompile program DSNHPC

Documentation

DB2 Replication Guide and Reference documentation

The solutions information at <http://www.ibm.com/software/data/dbtools/datarepl.htm> is no longer available. This information is referred to in the preface to the *Replication Guide and Reference*.

DB2 Universal Database, Version 8 HTML documentation installation restriction (Windows)

On Windows, do not install the DB2 Universal Database, Version 8 HTML documentation on a workstation or server where a DB2 Universal Database, Version 7 (or earlier) product is already installed. The installer detects the earlier version and removes the earlier product.

A workaround does exist. If you need to install the DB2 Universal Database, Version 8 HTML documentation on a workstation where an earlier version of DB2 Universal Database is installed, you can manually copy the files and directories from the DB2 Universal Database, Version 8 HTML Documentation CD, rather than using the installer. The DB2 Information Center and full text search will work, but you will not be able to apply any HTML documentation FixPaks.

Documentation search may fail on AIX unless all documentation categories are installed

If you do not to install all of the categories of documentation on the DB2 HTML documentation CD, an attempt to search against “All topics” may fail with an `InvalidParameterException` reported in the Java console of your browser and no search results.

To work around the documentation search problem do one of the following:

- Narrow the scope of your search by selecting from the **Search scope** list box in the Search window.
- Install all the documentation categories from the DB2 HTML documentation CD.

Documentation search problem with Java 2 JRE1.4.0

If your browser uses Java 2 JRE V1.4.0 and your documentation is installed in a path that contains spaces (e.g., `C:\Program Files\SQLLIB\doc\`), the documentation search applet may fail with an `InvalidParameterException` reported in the Java console of your browser and no search results. This problems is fixed with JRE V1.4.1.

To work around the documentation search problem, do one of the following:

- Upgrade your browser’s JRE version to 1.4.1, available at <http://java.sun.com/j2se/1.4.1/download.html>
- Downgrade your browser’s JRE version to 1.3.x, available at <http://www.ibm.com/software/data/db2/udb/ad/v8/java/>

Installation of the DB2 Information Center for languages that are not options during installation

The DB2 Setup wizard can only install the DB2 HTML documentation for languages that it also installs with the DB2 product. As a result, the DB2 HTML documentation cannot be installed using the DB2 Setup wizard in the following languages:

- Portuguese (restriction for UNIX only)
- Danish, Finnish, Norwegian, Swedish (restriction for Linux only)
- Dutch, Turkish (restriction for HP-UX, Solaris, Linux only)
- Arabic (restriction for UNIX only)

To install the DB2 Information Center for one of the languages listed previously:

1. Insert the *DB2 HTML Documentation CD* in your CD-ROM drive.
2. Copy the following directory to your computer:
 - `/cdrom/program files/IBM/SQLLIB/doc/htmlcd/language`

where *cdrom* is where you have mounted the CD and *language* is the code for the language you want to use.

It does not matter where you place the folder. You can also view the DB2 HTML documentation directly from the CD. For instructions on how to do so, see the "Viewing technical information online directly from the DB2 HTML Documentation CD" topic in the appendix of any DB2 Version 8 manual.

Notes:

1. To view the documentation, you must use Microsoft Internet Explorer 5.0 or later, or Netscape 6.1 browsers or later.
2. As well, if you launch the documentation from your product, it will go to the documentation installed as part of your product install, and not to the documentation that you copied over manually.

Official naming convention for DB2 Universal Database for Linux when used on host systems

The official naming convention for DB2 Universal Database for Linux on host systems is *DB2 on Linux for S/390® and zSeries™*. *S/390* refers to 32-bit and *zSeries* to 64-bit. In addition, it is important to note that the following terms are also obsolete:

- 64-bit Linux/390
- Linux/SGI

GUI tools

Control Center plugin support

The Control Center now supports custom folders. Custom folders can contain user-selected system or database objects. Creating Control Center plugins specifically for a custom folder is not supported, but plugins can be created for the object contained in custom folders. For more information about Control Center plugins, see the topic *Introducing the plug-in architecture for the Control Center*.

Displaying Indic characters in the DB2 GUI tools

If you have problems displaying Indic characters when using the DB2 GUI tools, you might not have the required fonts installed on your system.

DB2 Universal Database has packaged the following IBM TrueType and OpenType proportional Indic language fonts for your use. You can find these fonts in the font directory on either of the following CDs:

- IBM Developer Kit, Java Technology Edition, Version 1.3.1 for AIX operating systems on 64-bit systems
- Java application development and Web administration tools supplement for DB2, Version 8.1

These fonts are to be used only in conjunction with DB2. You cannot engage in the general or unrestricted sale or distribution of these fonts:

Table 2. Indic fonts packaged with DB2 Universal Database

Typeface	Weight	Font File Name
Devanagari MT for IBM	Medium	devamt.ttf
Devanagari MT for IBM	Bold	devamtb.ttf
Tamil	Medium	TamilMT.ttf
Tamil	Bold	TamilMTB.ttf
Telugu	Medium	TeluguMT.ttf
Telugu	Bold	TeleguMTB.ttf

Detailed instructions on how to install the fonts and modify the font.properties file can be found in the Internationalization section of the IBM development kit for Java documentation.

In addition, the following Microsoft products also come with Indic fonts that can be used with our GUI tools:

- Microsoft Windows 2000 operating system
- Microsoft Windows XP operating system
- Microsoft Publisher
- Microsoft Office

GUI tools not supported for zSeries servers running Linux operating systems

With the exception of the DB2 Setup wizard, GUI tools will not work on zSeries servers running the Linux operating system. This limitation includes any items normally launched from the Install launchpad, such as the Quick Tour.

If you want to use the GUI tools with one of these systems, install the administrative tools on a client system with a different system configuration, and use this client to connect to your zSeries server.

Load and Import Columns page does not support DBCS characters in IXF files

If you use the Load wizard or Import notebook to set up a load or import from an IXF input file containing DBCS characters, the Columns page will not correctly display the column names contained in the file.

Incorrect indicators given when a load operation fails

If a load fails, but only warnings were returned (not errors), then the task icon will still appear with a green checkmark in the Task Center. Be sure to double-check the success of any loads you carry out.

Minimum display settings for GUI tools

For the GUI tools such as the Control Center to work properly, you must have a screen resolution of at least 800 by 600 dpi, and be using a display palette of at least 32 colors.

SQL1224N error when using the GUI tools on AIX

If you are using the GUI tools on an AIX operating system, you may receive an SQL1224N error. This error is caused by a memory handling problem in DB2. The following workaround will help eliminate the error:

Procedure:

To stop the SQL1224N error on AIX operating systems:

1. As the instance owner, run the following commands:

```
export EXTSHM=ON
db2set DB2ENVLIST=EXTSHM
```

2. Restart the instance with the following commands:

```
db2stop
db2start
```

Once the instance restarts with the new environment variable settings, the SQL1224N errors should stop.

Health Monitor

Health Monitor off by default

The default value for the database manager switch for the health monitor (HEALTH_MON) is OFF.

Health indicator restrictions

- 2 The health monitor is unable to execute actions for the db2.db2_op_status
2 health indicator if the indicator enters the down state. This state can be
2 caused, for example, when an instance that the indicator is monitoring
2 becomes inactive because of an explicit stop request or an abnormal

2 termination. If you want to have the instance restart automatically after any
2 abnormal termination, you must configure the fault monitor to keep the
2 instance highly available.

Information Catalog Center

2 Do not partition information catalog tables

2 Tables that the Information Catalog Manager uses must be contained within a
2 single database partition. Numerous methods are available to put the tables
2 within a single partition. The following procedure is one approach that works.

- 2 1. Open a DB2 Command Line Processor and issue these commands:
 - 2 a. CREATE DATABASE PARTITION GROUP *pname* ON DBPARTITIONNUM *pnumber*
 - 2 b. CREATE REGULAR TABLESPACE *tsname* IN DATABASE PARTITION GROUP *pname*
2 MANAGED BY SYSTEM USING ('*cname*')

2 Click Start -> Programs -> IBM DB2 -> Set-up Tools -> Manage
2 Information Catalog Wizard.

2 On the Options page, specify the table space name in the **Table space**
2 field.

4 Log file not generated when importing tag language files

4 If an Information Catalog Center log file is not generated when importing tag
4 language files to the Information Catalog Center, perform the following
4 troubleshooting steps:

4 When running db2icmimport from a command line:

- 4 • If output files were not generated (.xml, .out, .err, .log), then there is
4 probably an error on the command line. Verify that the first five
4 arguments, which are UserId, Password, Database, Catalog, and
4 Tagfile, are correct. View the syntax by entering db2icmimport. If
4 this does not solve the problem, modify db2icmimport to capture
4 the output of db2javit by using the -g option to save the output to a
4 file (for example, db2javit
4 -j:com.ibm.db2.common.icm.tag.IcmImport -w: -i: -o:"-Xmx128m
4 -Xms32m" -g:"d:\temp\myimport.trc" . . .).
- 4 • If a log file was not generated, it is usually a parsing error. Look at
4 the .xml file and the .out file. If you can, insert a
4 ":COMMIT.CHKPID(DEBUG)" command at the beginning of the tag
4 language file. This command will generate debug report messages
4 and check the .xml and .out files for parsing errors.
- 4 • After parsing, errors should appear in the .log file. When the debug
4 report is being generated, look at the .log file and the .out file for
4 information.
- 4 • Always check the .err file to see if there is a run-time error.

4 **When importing tag language files using the Information Catalog Center**
4 **GUI:**

- 4 • When you import tag language files using the GUI interface, no
- 4 .out or .err files are generated.
- 4 • If a .log or .xml file is generated, try to debug by using those files.
- 4 • If they are not generated or do not help, run the import process
- 4 from a command line to get more information.

1 **Migrating DB2 Universal Database when using DataJoiner or replication**

1 If you want to migrate an instance of DataJoiner[®] or DB2 Universal Database
1 for Linux, UNIX, and Windows on which you are running the Capture or
1 Apply programs for DB2 Universal Database replication, you must prepare to
1 migrate your replication environment before you migrate the DB2 Universal
1 Database or DataJoiner instance. Detailed instructions for carrying out the
1 required preparation are included in the migration documentation for DB2
1 DataPropagator[™], Version 8. You can find migration documentation for DB2
1 DataPropagator, Version 8 at
1 <http://www.ibm.com/software/data/dpropr/replmigration.htm>.

Query Patroller Center

4 **Enhanced filtering abilities**

4 When you are using the Query Patroller Center to view the Managed Queries
4 report and the Query Activity Over Time report, a significant number of rows
4 might be displayed. A filter window can be used to display only the queries
4 that meet the conditions based on the filtering criteria supplied for each
4 report.

4 For the Managed Queries report, you can filter on:

- 4 • ID
- 4 • Status
- 4 • Submitter ID
- 4 • Time Created
- 4 • Time Completed
- 4 • Query Class of the query

4 For the Query Activity Over Time report, you can filter on:

- 4 • Query ID
- 4 • Submitter ID
- 4 • Explain Run

4 **Notes:**

- 4 1. For all Historical Analysis reports (such as Query Activity Over Time), you
4 can also filter on completion time by using the time range fields that
4 appear at the top of the Query Patroller Center when a historical object is
4 selected.

4 To open the Filter notebook from the Managed Queries folder:

- 4 1. Open the Query Patroller Center.
4 2. Expand the object tree under the Monitoring folder until you find the
4 Managed Queries folder.
4 3. Right-click the Managed Queries folder and select Filter in the pop-up
4 menu. The Filter notebook opens.

4 To open the Filter notebook from the Queries folder:

- 4 1. Open the Query Patroller Center.
4 2. Expand the object tree under the Historical Analysis folder until you find
4 the Queries folder.
4 3. Right-click the Queries folder and select Filter in the pop-up menu. The
4 Filter notebook opens.

4 The Filter notebook will also open automatically when a Managed Queries or
4 Queries folder is selected, if both of the following conditions are met:

- 4 • The number of objects in the folder exceeds the object count specified in the
4 Filter notebook
4 • The Filter notebook's **Automatically display filter when object count is
4 exceeded** option is enabled

4 **Procedure:**

4 The procedure for filtering tables for historical analysis is described in
4 *DB2 Query Patroller Guide: Installation, Administration, and Usage*
4 manual.

4 The following table shows the values expected for each column listed
4 in the Filter dialog for Managed Queries:

4 *Table 3. Valid values for Filter dialog (Managed Queries)*

Column	Value
ID	Numeric value representing the query ID

Table 3. Valid values for Filter dialog (Managed Queries) (continued)

Column	Value
Status	A single character representing the status. Possible values are: <ul style="list-style-type: none"> • H (for queries whose status is Held) • Q (for queries whose status is Queued) • R (for queries whose status is Running) • A (for queries whose status is Aborted) • C (for queries whose status is Cancelled) • D (for queries whose status is Done) • U (for queries whose status is Unknown)
Submitter ID	A character string representing the ID of the submitter
Created	A timestamp representing the time the query was created. For Example: 2003-07-29-00.00.00
Completed	A timestamp representing the time the query was completed. For Example: 2003-07-29-00.00.00
Query Class	Numeric value representing the ID of the query class used for this query

The following table shows the values expected for each column listed in the Filter dialog for Historical Queries:

Table 4. Valid values for Filter dialog (Historical Queries)

Column	Value
ID	Numeric value representing the query ID
Submitter ID	A character string representing the ID of the submitter
Explain Run	A single character that indicates whether the Historical Analysis Data Generator has been run on this query. Possible values are: <ul style="list-style-type: none"> • N (Historical Analysis Data Generator has not yet run) • S (Historical Analysis Data Generator ran successfully) • F (Historical Analysis Data Generator ran unsuccessfully)

1 Secure Windows environments

1 You may experience file permission problems if you are using DB2 Universal
1 Database on Windows and are not an administrator on the Windows system.
1 If you receive an SQL1035N, SQL1652N, or SQL5005C error message, possible
1 causes and workarounds are as follows:

1 User does not have sufficient authority on the sqllib directory:

1 Problem

1 Received an SQL1035N or SQL1652N error when trying to open DB2
1 CLP or command window. The DB2 Universal Database code (core
1 files) are installed into a directory structure where write privileges are
1 limited, but some DB2 Universal Database tools need to write and
1 create files in the DB2INSTPROF directory.

1 Workaround

1 Create a new directory where you can grant users, at minimum, the
1 MODIFY permission and use either **db2set -g db2tempdir** to point to
1 the new directory, or set the db2tempdir variable in the Windows
1 system environment.

1 User does not have sufficient authority to write to the sqllib*<instance_dir>* 1 directory even though user belongs to SYSADM_GROUP:

1 Problem

1 Received an SQL5005C system error when trying to update the
1 database manager configuration file (update dbm cfg). The user does
1 not have the required NTFS permissions to write to the
1 sqllib*instance_dir* directory even though you have added this user to
1 the SYSADM_GROUP.

1 First workaround

1 Grant the users, at minimum, the MODIFY permission on the
1 *instance_dir* directory at the file system level.

1 Second workaround

1 Create a new directory where you can grant the user, at minimum, the
1 MODIFY permission. Use the **db2set db2instprof** to point to the new
1 directory. You will either need to recreate the instance so that the
1 information is stored under the new instance directory specified by
1 db2instprof, or you will need to move the old instance directory to the
1 new directory.

SQL Assist

SQL Assist button disabled in the Command Center

In the Command Center, the SQL Assist button only becomes enabled once a connection has been established.

Two versions of SQL Assist are launched from DB2

You can invoke both version 7 and version 8 of SQL Assist from within DB2 Universal Database, Version 8.1. You can launch version 7 from the DB2 Data Warehouse Center. All other centers launch the latest version 8. The product online help has additional information for SQL Assist, Version 7.

System Monitor

Event record size limitation

For deadlock event monitors and global detailed deadlock event monitors, the event record is limited by the size of a nonconfigurable internal buffer. If db2diag.log cannot write this logging entry due to the record size, the log will record a message such as "event record is larger than BUFFERSIZE."

Snapshot UDF restrictions

Snapshot user-defined functions (UDF) are intended to be used on databases whose **Directory entry type** value displays as Indirect or Home when the LIST DB DIRECTORY command is issued. If a UDF is used against a remote database, the UDF will fail with the following error:

```
SQL1427N An instance attachment does not exist.
```

The snapshot UDFs, introduced in V8.1 cannot be used in conjunction with the monitor switches commands and APIs or monitor reset commands and APIs. This restriction includes:

- GET MONITOR SWITCHES
- UPDATE MONITOR SWITCHES
- RESET MONITOR

This limitation is due to the fact that such commands use an INSTANCE ATTACH, while the snapshot UDFs make use of DATABASE CONNECTs.

Amount of Log to be Redone for Recovery monitor element undefined for V8.1.4

In the Release information section of the Information Center for V8.1.4, log_to_redo_for_recovery is introduced as a new monitor element. The value of this monitor element is undefined for V8.1.4. This functionality is planned in a future release.

2 Throttled utility restrictions

2 The simultaneous execution of multiple throttled utilities is not supported. For
2 example:

- 2 • If you are running three online backups, only one can be throttled. The
2 other two must have a priority of 0.
- 2 • You can invoke a rebalance and backup at the same time, but either the
2 rebalance or backup must have a priority of 0.

2 If you invoke multiple throttled utilities at the same time, it can cause the
2 utilities to run for excessively long periods of time. It can also cause the
2 system impact to be more severe than the limit set by the impact policy
2 (UTIL_IMPACT_LIM).

XML Extender

2 XML Extender sample programs renamed

2 Conflicts between the system operations and XML Extender can cause some
2 XML Extender sample programs to seriously damage your files. The following
2 list shows the conflicting XML Extender sample programs as well as new
2 replacement programs that are less likely to cause conflicts. Make sure to use
2 the new sample programs in place of the old ones.

2 *Table 5. Replacement sample programs for XML Extender (Windows)*

2 Old Program (Do not use)	2 New Program (Use)
2 insertx.exe	2 dxxisrt.exe
2 retrieve.exe	2 dxxretr.exe
2 retrieve2.exe	2 dxxretr2.exe
2 retrievec.exe	2 dxxretrc.exe
2 shred.exe	2 dxxshrd.exe
2 tests2x.exe	2 dxxgenx.exe
2 tests2xb.exe	2 dxxgenxb.exe
2 tests2xc.exe	2 dxxgenxc.exe

2 *Table 6. Replacement sample programs for XML Extender (UNIX)*

2 Old Program (Do not use)	2 New Program (Use)
2 insertx	2 dxxisrt
2 retrieve	2 dxxretr
2 retrieve2	2 dxxretr2
2 retrievec	2 dxxretrc

Table 6. Replacement sample programs for XML Extender (UNIX) (continued)

Old Program (Do not use)	New Program (Use)
shred	dxxshrd
tests2x	dxxgenx
tests2xb	dxxgenxb
tests2xc	dxxgenxc

Using the new sample programs with the sample sqx files

The source code (.sqx files) for the executables listed above are located in the samples\db2xml\c directory of your installation. The source files are still labeled with their old names. If you make changes to the source code, copy your newly compiled executables (with the old names) to the sqllib\bin directory. On Windows platforms, you must make an additional copy, rename it with its new name above, and copy it to the bin directory. Both copies replace the existing files in the bin directory. For example, after compiling your new version of shred.exe, you need to make two copies and replace the files in the bin directory: one labeled shred.exe and the other renamed dxxshrd.exe. On UNIX platforms, you only need to replace the file with the old name with your newly compiled version.. If you create new executable files from these samples, you must copy the new files from the \SQLLIB\samples\db2xml\c\ directory into the \SQLLIB\bin\ directory, and then make an additional copy, renaming them according to the table above.

Decomposing documents in XML Extender that contain non-unique attribute and element names

You can now decompose documents that contain non-unique attributes and/or non-unique element names that map to different columns (of the same or different tables) without receiving the DXXQ045E error. The following is an example of an XML document with non-unique attributes and non-unique element names:

```
<Order ID="0001-6789">
  <!-- Note: attribute name ID is non-unique -->
  <Customer ID = "1111">
    <Name>John Smith</Name>
  </Customer>
  <!-- Note: element name Name is non_unique -->
  <Salesperson ID = "1234">
    <Name>Jane Doe</Name>
  </Salesperson>
  <OrderDetail>
    <ItemNo>xxxx-xxxx</ItemNo>
    <Quantity>2</Quantity>
    <UnitPrice>12.50</UnitPrice>
  </OrderDetail>
</OrderDetail>
```

```

4         <ItemNo>yyyy-yyy</ItemNo>
4         <Quantity>4</Quantity>
4         <UnitPrice>24.99</UnitPrice>
4     </OrderDetail>
4 </Order>

```

The accompanying DAD, which maps the duplicate elements/attributes to different columns, looks like this:

```

4 <element_node name="Order">
4   <RDB_node>
4     <table name="order_tab" key="order_id"/>
4     <table name="detail_tab"/>
4     <condition>
4       order_tab.order_id = detail_tab.order_id
4     </condition>
4   </RDB_node>
4
4   <!-- attribute ID duplicated below, but mapped to a different column -->
4   <attribute_node name="ID">
4     <RDB_node>
4       <table name="order_tab" />
4       <column name="order_id" type="char(9)"/>
4     </RDB_node>
4   </attribute_node>
4
4   <element_node name="Customer">
4     <!-- attribute ID duplicated above, but mapped to a different column -->
4     <attribute_node name="ID">
4       <RDB_node>
4         <table name="order_tab" />
4         <column name="cust_id" type="integer"/>
4       </RDB_node>
4     </attribute_node>
4
4     <!-- element name duplicated below, but mapped to a different column -->
4     <element_node name="Name">
4       <text_node>
4         <RDB_node>
4           <table name="order_tab" />
4           <column name="cust_name" type="char(20)" />
4         </RDB_node>
4       </text_node>
4     </element_node>
4   </element_node>
4
4   <element_node name="Salesperson">
4     <!-- attribute ID duplicated above, but mapped to a different column -->
4     <attribute_node name="ID">
4       <RDB_node>
4         <table name="order_tab" />
4         <column name="salesp_id" type="integer"/>
4       </RDB_node>
4     </attribute_node>

```

```

4
4      <!-- element name duplicated above, but mapped to a different column -->
4      <element_node name="Name">
4          <text_node>
4              <RDB_node>
4                  <table name="order_tab" />
4                  <column name="salesp_name" type="char(20)" />
4              </RDB_node>
4          </text_node>
4      </element_node>
4  </element_node>
4
4      <element_node name="OrderDetail" multi_occurrence="YES">
4          <element_node name="ItemNo">
4              <text_node>
4                  <RDB_node>
4                      <table name="detail_tab" />
4                      <column name="itemno" type="char(9)"/>
4                  </RDB_node>
4              </text_node>
4          </element_node>
4          <element_node name="Quantity">
4              <text_node>
4                  <RDB_node>
4                      <table name="detail_tab" />
4                      <column name="quantity" type="integer"/>
4                  </RDB_node>
4              </text_node>
4          </element_node>
4          <element_node name="UnitPrice">
4              <text_node>
4                  <RDB_node>detail_tab" />
4                  <table name="detail_tab" />
4                  <column name="unit_price" type="decimal(7,2)"/>
4                  </RDB_node>
4              </text_node>
4          </element_node>
4      </element_node>
4  </element_node>

```

The contents of the tables would look like the following after the document above is decomposed:

ORDER_TAB:

ORDER_ID	CUST_ID	CUST_NAME	SALESP_ID	SALESP_NAME
0001-6789	1111	John Smith	1234	Jane Doe

DETAIL_TAB:

ORDER_ID	ITEMNO	QUANTITY	UNIT_PRICE
0001-6789	xxxx-xxxx	2	12.50
0001-6789	yyyy-yyyy	4	24.99

4 **Note:** Multiple element/attribute mappings to the same column of the same
4 table are not allowed regardless of whether the element/attribute
4 names are different or the same.

4 **Additional Information**

Change in Unicode server behavior

In version 7, Unicode servers ignored any graphic code pages by applications at connect time and assumed that UCS2 Unicode (code page 1200) was being used. Version 8 Unicode servers now respect the code page sent by the client.

Full message text is not returned when using `SQLException.getMessage()`

By default, the `DB2BaseDataSource.retrieveMessagesFromServerOnGetMessage` property is disabled. If you enable this property, all calls to the standard JDBC `SQLException.getMessage()` invoke a server-side stored procedure, which retrieves the readable message text for the error. By default, the full message text is not returned to the client when a server-side error occurs.

You can use the proprietary method `DB2Sqlca.getMessage()` to retrieve the fully formatted message text. A call to the `SQLException.getMessage()` method will start a unit of work only if `retrieveMessagesFromServerOnGetMessage` is enabled. A call to the `DB2Sqlca.getMessage()` method results in a stored procedure call, which starts a unit of work. Prior to FixPak 1, the `DB2Sqlca.getMessage()` method may throw an exception.

IBM DB2 Universal JDBC driver

The IBM DB2 Universal JDBC Driver cannot connect to databases that were created using the HP default character set, `roman8`. All SQLJ and JDBC applications that use the universal JDBC driver, must connect to a database created with a different character set. If your `LANG` is set to `"C"` or to a `"roman8"` locale, you must change it to the corresponding ISO locale. For instance, if your `LANG` is set to `de_DE.roman8`, it must be changed to `de_DE.iso88591`:

```
export LANG=de_DE.iso88591
```

To run the DB2 SQLJ and JDBC sample programs with the universal JDBC driver, you can create the sample database with the following commands (in this case, using the ISO locale for US English):

```
export LANG=en_US.iso88591
db2 terminate
db2samp1
```

Note that if the sample database already exists, you must drop it before executing these commands.

Java functions and routines on Linux, UNIX, and Windows operating systems

Due to limitations in the JVM, a Java routine defined as NOT FENCED will be invoked as if it had been defined as FENCED THREADSAFE. In version 8.1, all Java routines specified as NOT FENCED are treated as FENCED routines. Java UDFs defined with parameter style DB2GENERAL that contain a LOB locator in the argument definition of the UDF will not function. You must modify these functions to use a BLOB or CLOB argument instead of a locator. For Java UDFs, LOB locators are only supported as input arguments when the DB2JAVA parameter style is specified.

English Microsoft Data Access Components (MDAC) files are used for all national language versions of DB2 Universal Database, V8.1, unless translated MDAC files are installed first.

If you do not install the national language version of MDAC 2.7 prior to installing the national language version of DB2, then DB2 Universal Database installs English MDAC files by default. This causes the Windows ODBC Data Source Administrator panels to appear untranslated if your operating system is non-English. To fix this problem, you can install the “MDAC 2.7 RTM - Refresh” bundle from the Microsoft Web site at http://www.microsoft.com/data/download_270RTM.htm. Choose the language that you want to install, download the required executable, and run it. This will install the translated ODBC Data Source Administrator files.

The Simplified Chinese locale on AIX operating systems

AIX has changed the code set bound to the Simplified Chinese locale Zh_CN on:

- AIX Version 5.1.0000.0011 or later
- AIX Version 5.1.0 with maintenance level 2 or later

The code set has been changed from GBK (code page 1386) to GB18030 (code page 5488 or 1392). Since DB2 Universal Database for AIX supports the GBK code set natively and the GB18030 code set via Unicode, DB2 Universal Database will default the Zh_CN locale’s code set to ISO 8859-1 (code page 819), and in some operations will also default the locale’s territory to the United States (US).

To work around this limitation, you have two options:

- You can override the locale’s code set from GB18030 to GBK and the territory from US to China (whose territory ID is CN and territory code is 86).
- You can use a different Simplified Chinese locale.

If you choose to use the first option, issue the following commands:

```
db2set DB2CODEPAGE=1386
db2set DB2TERRITORY=86
db2 terminate
db2stop
db2start
```

If you choose to use the second option, change your locale from Zh_CN to either ZH_CN or zh_CN. The ZH_CN locale's code set is Unicode (UTF-8), while the zh_CN locale's code set is eucCN (code page 1383).

The Simplified Chinese locale on Red Hat, V8 operating systems

Red Hat, Version 8 has changed the default code set for Simplified Chinese from GBK (code page 1386) to GB18030 (code page 5488 or 1392).

Since DB2 Universal Database for Linux supports the GBK code set natively and the GB18030 code set via Unicode, DB2 Universal Database will default its code set to ISO 8859-1 (code page 819), and in some operations will also default its territory to the United States (US).

To work around this limitation, you have two options:

- You can override the Red Hat default code set from GB18030 to GBK and the territory from US to China (whose territory ID is CN and territory code is 86).
- You can use a different Simplified Chinese locale.

If you choose to use the first option, issue the following statements:

```
db2set DB2CODEPAGE=1386
db2set DB2TERRITORY=86
db2 terminate
db2stop
db2start
```

If you choose to use the second option, issue any one of the following commands:

```
export LANG=zh_CN.gbk
export LANG=zh_CN
export LANG=zh_CN.utf8
```

where the code set associated with zh_CN is eucCN or code page 1383, and with zh_CN.utf8 is code page 1208.

Corrections to the documentation

This section describes corrections to previously released HTML, PDF, and printed documentation. Updated versions of the affected topics will be available in a future version of the DB2® documentation.

4 Administration Guide: Planning

4 Topic Title: Space requirements for temporary tables

4 Location in the DB2 Information Center accessed from a browser

4 Concepts -> Administration -> Database design -> Physical
4 -> Space requirements for database objects -> Temporary
4 tables

4 Correction

4 Replace the first paragraph with the following:

4 Some SQL statements require temporary tables for processing
4 (such as a work file for sorting operations that cannot be done
4 in memory). These temporary tables require disk space; the
4 amount of space required is dependent upon the size, number,
4 and nature of the queries, and the size of the returned tables.
4 Your work environment is unique which makes the
4 determination of your space requirements for temporary
4 tables difficult to estimate. For example, more space may
4 appear to be allocated for system temporary table spaces than
4 is actually in use due to the longer life of various system
4 temporary tables on disk. This could occur when the registry
4 variable DB2_SMS_TRUNC_TEMP_THRESH is used.

4 For more information on this registry variable, refer to the
4 "Registry and Environment Variables" section in this release
4 notes.

4 Topic Title: Temporary table space design

4 Location in the DB2 Information Center accessed from a browser

4 Concepts -> Administration -> Database design -> Physical
4 -> Table space design -> Temporary table space design

4 Correction

4 Add the following information:

4 If you are reliant on system temporary tables in SMS system
4 temporary table spaces because of your work environment,

4 you may want to consider using the registry variable
4 DB2_SMS_TRUNC_TEMP_THRESH. In the past when system
4 temporary tables were no longer needed, they were truncated
4 to a file size of zero. This was the default. The default has
4 changed so that the file size is no longer truncated to zero but
4 rather the file size is truncated to a non-zero size once the
4 default threshold value is exceeded. The need for space for a
4 new system temporary table would have a performance cost
4 associated with reserving the disk space for it. Using this
4 registry variable allows for leaving disk space allocated for
4 system temporary tables even after they have been dropped,
4 to avoid the performance cost of repeated space reservation
4 and truncation of storage on disk for system temporary tables.
4 This is particularly beneficial in environments where large
4 numbers of system temporary tables are created.

4 **Topic Title: Enabling bidirectional support**

4 **Location in the DB2 Information Center accessed from a browser**

4 Tasks -> Administering database systems -> Creating a
4 database -> Enabling language support -> Enabling
4 bidirectional support

4 **Correction**

4 Change the Restrictions section to include the following:

4 When converting from one Arabic CCSID to another Arabic
4 CCSID, DB2 UDB employs the following logic to deshape (or
4 expand) the lam-alef ligature. Deshaping will occur when the
4 Text Shaping attribute of the source Arabic CCSID is shaped
4 but the Text Shaping attribute of the target Arabic CCSID is
4 unshaped.

4 The logic to deshape the lam-alef ligature is:

- 4 1. If the last character of the data stream is a blank character,
4 then every character after the lam-alef ligature will be
4 shifted to the end of the data stream, therefore making
4 available an empty position for the current lam-alef
4 ligature to be deshaped (expanded) into its two constituent
4 characters: lam and alef.
- 4 2. Otherwise, if the first character of the data stream is a
4 blank character, then every character before the lam-alef
4 ligature will be shifted to the beginning of the data stream,
4 therefore making available an empty position for the
4 current lam-alef ligature to be deshaped (expanded) into
4 its two constituent characters: lam and alef.
- 4 3. Otherwise, there is no blank character at the beginning
4 and end of the data stream, and the lam-alef ligature

4 cannot be deshaped. If the target CCSID does have the
4 lam-alef ligature, then the lam-alef ligature remains as is;
4 otherwise, the lam-alef ligature is replaced by the target
4 CCSID's substitution character.

4 Conversely when converting from an Arabic CCSID whose
4 Text Shaping attribute is unshaped to an Arabic CCSID whose
4 Text Shaping attribute is shaped, the source lam and alef
4 characters will be contracted to one ligature character, and a
4 blank character is inserted at the end of the target area data
4 stream.

4 Change the final paragraph in the Procedure to the following:

4 For DRDA environments, if the HOST EBCDIC platform also
4 supports these bidirectional CCSIDs, you only need to set the
4 DB2CODEPAGE value. Note that you must not further specify
4 the same CCSID on the BIDI parameter in the PARMS field of
4 the DCS database directory entry for the server database,
4 otherwise an extra bidi layout conversion would occur, and
4 render any Arabic data to be incorrectly reversed. However, if
4 the HOST platform does not support these CCSIDs, you must
4 also specify a CCSID override for the HOST database server
4 to which you are connecting. This is accomplished through
4 the use of the BIDI parameter in the PARMS field of the DCS
4 database directory entry for the server database. The override
4 is necessary because, in a DRDA environment, code page
4 conversions and layout transformations are performed by the
4 receiver of data. However, if the HOST server does not
4 support these bidirectional CCSIDs, it does not perform layout
4 transformation on the data that it receives from DB2 UDB. If
4 you use a CCSID override, the DB2 UDB client performs
4 layout transformation on the outbound data as well. Please
4 consult the "Handling BiDi data" section of the IBM DB2
4 Connect User's Guide for more information.

4 **Topic Title: Version 8 incompatibilities**

4 **Location in the DB2 Information Center accessed from a browser**
4 Release information -> Incompatibilities between releases

4 **Correction**

4 Within the Messages subsection, the following item "DB2
4 Connect messages returned instead of DB2 messages" needs
4 to add the following information:

4 The messages affected by this change are related to bind,
4 connection, or security errors. SQL errors for queries and
4 other SQL requests are not affected by this change.

4 **Application development: Call Level Interface (CLI)**

4 **Topic Title: Cursor Considerations in CLI Applications**

4 **Location in the DB2 Information Center accessed from a browser**

4 Concepts -> Application development -> Programming
4 features -> Cursors

4 **Correction**

4 The discussion of which cursor type to use should include
4 dynamic scrollable cursors as a possible cursor type. Dynamic
4 scrollable cursors are only supported when accessing servers
4 which are DB2 UDB for z/OS Version 8.1 and later.
4 Keyset-driven cursors are recommended over dynamic
4 scrollable cursors when the application does not need to
4 detect rows that are inserted by other cursors, and when it is
4 acceptable for deleted rows to exist as "holes" in the result set.
4 Dynamic cursors do not recognize deleted rows as "holes", but
4 instead replace the deleted row with the next row in the result
4 set, if one exists.

4 **Topic Title: Importing Data with the CLI LOAD Utility in CLI Applications**

4 **Location in the DB2 Information Center accessed from a browser**

4 Tasks -> Developing applications -> Programming with
4 application programming interfaces (APIs) -> CLI ->
4 Processing transactions -> Inserting data

4 **Correction**

4 The restriction requiring parameter markers for each column
4 of the target table in the prepared INSERT statement does not
4 apply when a fullselect is used instead of the VALUES clause
4 in the INSERT statement. The CLI LOAD utility now supports
4 a fullselect on the INSERT statement, allowing you to load
4 data from one table into another while executing the INSERT
4 statement only once.

4 **Topic Title: Statement Attributes (CLI) List**

4 **Location in the DB2 Information Center accessed from a browser**

4 Reference -> APIs (application programming interfaces) ->
4 CLI attributes

4 **Correction**

4 For the SQL_ATTR_CURSOR_TYPE statement attribute, the

SQL_CURSOR_DYNAMIC value is supported by DB2 CLI if the server is DB2 UDB for z/OS Version 8.1 and later.

Data Links Manager

Topic Title: Changing the diagnostic level for error message log files

Location in the DB2 Information Center accessed from a browser

Tasks -> Data Links Manager-> Working with the Data Links File Manager-> Logging

Correction

1. Replace the first sentence of the first paragraph, which currently states: "Data Links maintains error message log files from the Data Links server and from the DB2 databases that manage DATALINK data files." with the following sentence: "DB2 Data Links Manager maintains error message log files from the Data Links server and from the DB2 host databases that manage linked files."
2. Remove all references to the DLFM_LOG_LEVEL registry variable from this entire topic. The DLFM_LOG_LEVEL registry variable is no longer supported. DB2 Data Links Manager information is now included in the DB2 trace functionality in this release.
 - Replace the first paragraph of the **Procedure** section with the following paragraph:

"You control the level of detailed information that is written to the db2diag.log file on DB2 host and Data Links server computers by using the *DIAGLEVEL* configuration parameter."
 - Remove the entire third paragraph of the **Procedure** section.
 - Add Related reference links to the "DB2 trace (db2trc)" and "db2trc - Trace Command" topics. Add a Related concept link to the "DB2 traces" topic.

Topic Title: Using Tivoli Storage Manager as an archive server (AIX)

Location in the DB2 Information Center accessed from a browser

Tasks -> Data Links Manager-> Configuring system management options

Correction

Replace the text in step 2 with the following text:

"Register the Data Links server machine on which you installed the Tivoli Storage Manager client application with

4 the Tivoli Storage Manager server. For more information, see
4 your Tivoli Storage Manager product documentation.”

4 **Topic Title: Using Tivoli Storage Manager as an archive server (Solaris**
4 **Operating Environment)**

4 **Location in the DB2 Information Center accessed from a browser**

4 Tasks -> Data Links Manager-> Configuring system
4 management options

4 **Correction**

4 Replace the text in step 2 with the following text:

4 “Register the Data Links server machine on which you
4 installed the Tivoli Storage Manager client application with
4 the Tivoli Storage Manager server. For more information, see
4 your Tivoli Storage Manager product documentation.”

4 **Topic Title: Using Tivoli Storage Manager as an archive server (Windows)**

4 **Location in the DB2 Information Center accessed from a browser**

4 Tasks -> Data Links Manager-> Configuring system
4 management options

4 **Correction**

4 1. Replace the text in step 2 with the following text:

4 “Register the Data Links server machine on which you
4 installed the Tivoli Storage Manager client application with
4 the Tivoli Storage Manager server. For more information,
4 see your Tivoli Storage Manager product documentation.”

4 2. Remove step 5 entirely.

4 3. In step 7, change the reference to the `dsm.sys` file to be to
4 the `dsm.opt` file. The revised step reads as follows:

4 “Set the `PASSWORDACCESS` option to generate in the
4 `c:\tsm\baclient\dsm.opt` Tivoli Storage Manager system
4 options file.”

4 **Data Warehouse Center**

4 **Topic Title: Syntax for starting the external trigger client**

4 **Location in the DB2 Information Center accessed from a browser**

4 Reference -> Data warehouses -> Administration

4 **Correction**

4 When running a step from outside the Data Warehouse Center
4 using the external trigger program, a seventh value for the

4 Information Catalog Center

4 **Topic Title: Configuring the Information Catalog Center for the Web with**
4 **WebSphere Application Server 5**

4 **Location in the DB2 Information Center accessed from a browser**
4 Tasks -> Data warehouses -> Installing Warehouse Manager
4 components -> Information Catalog Manager tools

4 **Correction**

4 In the fifth step "Update the application classpath," the
4 following example is an incorrect classpath for Windows:

4 C:\Program Files\ibm\sqllib\java\db2cmn.jar

4 The correct classpath example is:

4 C:\Program Files\ibm\sqllib\tools\db2cmn.jar

4 In addition, you can connect to only one Information Catalog
4 per Information Catalog Center for the Web application,
4 because all of the Web users log on with the same user ID and
4 password to access the catalog.

4 National languages

4 **Topic Title: String comparisons in a Unicode database**

4 **Location in the DB2 Information Center accessed from a browser**
4 Concepts -> Administration -> Unicode character encoding

4 **Correction**

4 Pattern matching is one area where the behavior of existing
4 MBCS databases is slightly different from the behavior of a
4 UCS-2 database.

4 For MBCS databases in DB2 UDB, the current behavior is as
4 follows: If the match-expression contains MBCS data, the
4 pattern can include both SBCS and non-SBCS characters. The
4 special characters in the pattern are interpreted as follows:

- 4 • An SBCS halfwidth underscore refers to one SBCS character.
- 4 • A non-SBCS fullwidth underscore refers to one non-SBCS
4 character.
- 4 • A percent (either SBCS halfwidth or non-SBCS fullwidth)
4 refers to zero or more SBCS or non-SBCS characters.

4 In a Unicode database, there is really no distinction between
4 "single-byte" and "non-single-byte" characters. Although the
4 UTF-8 format is a "mixed-byte" encoding of Unicode
4 characters, there is no real distinction between SBCS and

4 non-SBCS characters in UTF-8. Every character is a Unicode
4 character, regardless of the number of bytes in UTF-8 format.
4 In a Unicode graphic column, every non-supplementary
4 character, including the halfwidth underscore (U+005F) and
4 halfwidth percent (U+0025), is two bytes in width. For
4 Unicode databases, the special characters in the pattern are
4 interpreted as follows:

- 4 • For character strings, a halfwidth underscore (X'5F') or a
4 fullwidth underscore (X'EFBCBF') refers to one Unicode
4 character. A halfwidth percent (X'25') or a fullwidth percent
4 (X'EFBC85') refers to zero or more Unicode characters.
- 4 • For graphic strings, a halfwidth underscore (U+005F) or a
4 fullwidth underscore (U+FF3F) refers to one Unicode
4 character. A halfwidth percent (U+0025) or a fullwidth
4 percent (U+FF05) refers to zero or more Unicode characters.

4 **Note:** You need two underscores to match a Unicode
4 supplementary graphic character because such a
4 character is represented by two UCS-2 characters in a
4 GRAPHIC column. Only one underscore is needed to
4 match a Unicode supplementary character in a CHAR
4 column.

4 For the optional "escape expression", which specifies a
4 character to be used to modify the special meaning of the
4 underscore and percent sign characters, the expression can be
4 specified by any one of:

- 4 • A constant
- 4 • A special register
- 4 • A host variable
- 4 • A scalar function whose operands are any of the above
- 4 • An expression concatenating any of the above

4 with the restrictions that:

- 4 • No element in the expression can be of type LONG
4 VARCHAR, CLOB, LONG VARGRAPHIC, or DBCLOB. In
4 addition, it cannot be a BLOB file reference variable.
- 4 • For CHAR columns, the result of the expression must be
4 one character or a binary string containing exactly one (1)
4 byte (SQLSTATE 22019).
- 4 • For GRAPHIC columns, the result of the expression must
4 be one character (SQLSTATE 22019).

4 Registry and environment variables

4 Topic Title: Performance variables

4 Location in the DB2 Information Center accessed from a browser

4 Reference -> Registry and environment variables ->
4 Performance variables

4 Correction

4 Add the following performance variable:

4 *Table 7. Performance variable DB2_SMS_TRUNC_TEMP_THRESH*

Variable Name	Operating System	Values
Description		
DB2_SMS_TRUNC_TEMP_THRESH	All	Default=1 Values: -1, 0-n, where n=the number of extents per container that are to be maintained
Specifies a minimum file size threshold at which the file representing a temporary table will be maintained in SMS table spaces. Setting this variable to a value greater than 0 can reduce some of the system overhead involved in dropping and recreating the file each time a temporary table is used. By default, when a temporary table is no longer needed, the file for that table is truncated to 1 extent per container. If the file is already one extent or smaller in size, it is left as is. When the value of this variable is higher than 1, a larger file is maintained.		
If this variable is set to -1, then the file is not be truncated at all and the file will be allowed to grow indefinitely, restricted only by system resources.		
If this variable is set to 0, then no special threshold handling is done. Instead, once a temporary table is no longer needed, that file is truncated to 0.		

4 XML Extender

4 Topic Title: Planning for XML columns

4 Location in the DB2 Information Center accessed from a browser

4 Tasks -> XML Extender -> Managing data in XML columns
4 -> Planning for XML columns

4 Correction

4 The subsection of this topic titled 'The DAD file' incorrectly
4 states that the DAD file for XML columns specifies the
4 directory paths on the client for the DAD file and DTD.

4
4
4

The DAD file for XML columns and XML collections does not specify directory paths on the client or the server for either type of files.

Online help corrections and updates

Configuring the C environment for SQL stored procedures in the Development Center

If you are working with DB2® for Windows® on the server and you are using the Visual C++ compiler, you must configure your SQL build settings. You will not be able to build SQL stored procedures until you configure your SQL Build options.

Use the Database Connection Properties notebook in the Development Center to configure your SQL build settings.

To configure the C compiler environment for SQL stored procedures:

1. On the SQL Build Settings page of the notebook, specify a compiler environment that you want to use for building SQL objects.
 - Click **Refresh**.
 - In the **Compiler environment** field, type the location of the VC98\BIN\VCVARS32.BAT file on your Windows server.
2. Click **OK** to close the notebook and save your changes. If you click **Apply**, the changes are saved and you can continue changing the properties.

2 Enabling view docking when accessing Development Center with Hummingbird

2 Exceed

2 When accessing the Development Center on UNIX® with Hummingbird®
2 Exceed, the XTEST extension version 2.2 must be enabled before you can
2 move and dock views by dragging their title bars within the Development
2 Center.

2 To enable the XTEST extension:

- 2 1. From the Start menu, select **Programs** -> **Hummingbird Connectivity 7.0**
2 ->**Exceed**->**XConfig**. The XConfig window will open.
- 2 2. Optional: If your configuration requires a password, enter the XConfig
2 password.
- 2 3. Double click the **Protocol** icon. The Protocol window will open.
- 2 4. Check the **X Conformance Test Compatibility** checkbox.
- 2 5. In the **Protocol** window, click the **Extensions...** button. The Protocol
2 Extensions window will open.
- 2 6. In the Enable Extensions list, select the **XTEST(X11R6)** checkbox.

2

7. Click **OK**.

2 **Microsoft Visual Studio .NET add-in information update in Development Center** 2 **help**

2
2
2
2
2

The "About the Development Center" help topic does not include information about the new Microsoft® Visual Studio .NET add-in in the list of development environment add-ins provided. The following information describes the .NET add-in that supports Development Center functionality in the Microsoft Visual Studio .NET development environment:

2
2

DB2 Development Add-In for the Microsoft Visual Studio .NET development environment:

2
2
2
2
2
2

A new component of the DB2 Application Development Client is the IBM® DB2 Development Add-In for Microsoft Visual Studio .NET for the .NET framework version 1.0. This add-in extends the Visual Studio .NET IDE to provide tightly integrated DB2 application development support using the DB2 .NET Managed Provider as well as DB2 server-side development support. Using this add-in available in Microsoft Visual Studio .NET you can:

2
2
2
2
2
2
2
2
2
2
2
2

- Develop DB2 specific database projects from the new IBM Projects folder with advanced script generation wizards.
- Explore the DB2 catalogue information using DB2 data connections in the new IBM Explorer.
- Make use of the extended intelligence features for DB2 table/view columns and procedure/function parameters.
- Generate ADO.NET code for your windows forms using drag and drop.
- Configure your DB2 Managed Provider objects using properties custom editors and wizards.
- Launch various DB2 development and administration centers.
- View add-in help from the existing dynamic help window.

2
2
2

The DB2 Development Add-In for Microsoft Visual Studio .NET database connections are managed using the DB2 .NET managed provider and ADO.NET.

2 **Migrating DB2 XML Extender to Version 8.1.2**

2
2
2
2

If you are migrating from a Version 7 FixPak, refer to each of the release notes for the Version 7 FixPak for more information on what changes are included when upgrading to Version 8.1.2. Each new FixPak contains all of the previous FixPaks' updates.

2
2

To migrate DB2 XML Extender from previous versions to Version 8.1.2, complete the following steps.

- 2 1. From the DB2 Command Line, enter:
2 `db2 connect to database_name`
2 `db2 bind dxxinstall\@dxxMigv.lst`

- 2 where *dxxinstall* is the directory path in which you installed DB2 Universal
2 Database .
- 2 2. From the DB2 Command Line, enter:
2 `dxxMigv database_name`

Path settings to enable Java routines to compile in the Development Center

The Development Center cannot compile Java™ routines unless it knows where your developer kit versions are installed. Default locations for these will be written to your \$HOME/IBM/DB2DC/DB2DC.settings file when the Development Center starts for the first time. You may copy these into your \$USER.settings file and modify them with a Unicode editor, or you may create symbolic links to your developer kit directories in the default locations.

Runstats dialog – updated getting there information

To open the Runstats notebook:

1. From the Control Center, expand the object tree until you find the Tables folder.
2. Click the Tables folder. Any existing tables are displayed in the contents pane.
3. Right-click all the tables that you want to run statistics on, and select Run Statistics from the pop-up menu. The Runstats notebook opens.

Specifying build options for a Java stored procedure in the Development Center

Use the Stored Procedure Properties notebook to specify the compile options that will be used when building a Java stored procedure.

These steps are part of the larger task of changing stored procedure properties.

To specify the build options of a stored procedure:

1. On the Build page of the Stored Procedure Properties notebook, specify the compile options for building the stored procedure. See your compiler documentation for information about available options.
 - a. In the Precompile options field, type the DB2 Universal Database™ precompiler options that you want to use when building stored procedures. The package name must not exceed 7 characters.

- b. In the Compile options field, type the compiler options that you want to use when building stored procedures.
2. Click **OK** to close the notebook and save your changes. If you click **Apply**, the changes are saved and you can continue changing the properties.

Appendix. Notices

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

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

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

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

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

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

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

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

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

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

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

IBM Canada Limited
Office of the Lab Director
8200 Warden Avenue
Markham, Ontario
L6G 1C7
CANADA

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

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

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

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

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

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

COPYRIGHT LICENSE:

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

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

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

Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both, and have been used in at least one of the documents in the DB2 UDB documentation library.

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

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

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

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

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

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

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



Printed in USA