

IBM DB2[®] Universal Database[™]



Release Notes

Version 8.1 FixPak 6

IBM DB2[®] Universal Database[™]



Release Notes

Version 8.1 FixPak 6

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

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

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

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

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

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

© Copyright International Business Machines Corporation 2002 - 2004. All rights reserved.

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

Contents

About the Release Notes	vii	6	Alternative Unicode conversion tables for the coded character set identifier (CCSID) 954	9
About this release	1	6	Replacing the Unicode conversion tables for coded character set (CCSID) 954 with the Microsoft conversion tables	10
6 New in this release	1	6	MVS operating system not supported	10
6 Support for additional collating sequences when a new Unicode database is created	1	6	DB2 UDB for AIX 4.3.3 and 5.1 and above require AIX C++ Version 6 Runtime	10
Product fix history	1	5	Accessing Java shared libraries (Linux)	11
Backward compatibility	2	5	Backup and restore (Linux 390)	12
3 Alternate FixPaks (Linux and UNIX)	2	5	ERwin 4.x metadata bridge	12
Data Warehouse Center previous server support restrictions	3	5	Enabling view docking when accessing the Development Center with Hummingbird	2
4 The Visual Warehouse 5.2 DB2 program "VW 5.2 Load flat file into DB2 UDB EEE (AIX only)" is not supported	3	2	Exceed	12
4 DB2 Universal Database Version 7 server access	4	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	4	3	Installation, migration, upgrade, and configuration information	13
5 Command Center and Version 7 servers	5	2	Installation information	13
Two versions of SQL Assist are launched from DB2 UDB	5	5	Installing DB2 UDB (AIX)	13
Change in Unicode server behavior	5	6	Installing DB2 UDB (Linux)	13
Classic Connect unavailable	5	6	Installing additional Asian fonts (Linux)	13
Product compatibility	5	6	Restrictions to adding products using the db2setup command	15
2 IBM DB2 Development Add-In for Microsoft Visual Studio .NET	5	6	Downloading and uncompressing Spatial Extender FixPaks	16
6 WebSphere Application Server's execution of SQL statements in an SQLJ program	5	6	Binding Query Patroller packages after applying fixpaks	16
1 Microsoft Visual Studio, Visual C++	6	6	Installing Query Patroller at FixPak 3 level or later.	17
Microsoft XP fix needed on 64-bit operating systems	6	6	DB2 Web Tools	17
2 CLI and ODBC on Windows 64-bit operating systems	6	6	Installing the DB2 Information Center for languages that are not options during installation	17
4 Supported LDAP client and server configurations.	6	6	DB2 Universal Database, Version 8 HTML documentation installation restriction (Windows)	18
4 Windows XP operating systems	7	6	Installing MDAC files for national language versions of DB2 UDB.	18
6 Alternative Unicode conversion tables for the coded character set identifier (CCSID) 5039	7	1	Migration information	18
6 Replacing the Unicode conversion tables for coded character set (CCSID) 5039 with the Microsoft conversion tables	8	5	1 Migrating DB2 Universal Database when using DataJoiner or replication	18
		5	5 Migrating a DB2 Version 8 Windows 32-bit database to Windows 64-bit	19

4	Decomposing documents in XML Extender	5	SQL message additions	51
4	that contain non unique attribute and	5	DBI message updates	52
4	element names	5	Query Patroller	54
5	Differences between SNA and TCP/IP	5	Create Explain tables before running	
5	when using DB2 Connect	5	Query Patroller Historical Data Generator .	54
	Documentation updates	5	Checking Query Patroller log files for	
	Administration: Performance tuning	5	historical analysis	54
6	Default value for DB2_ENABLE_BUFDPD	6	Dynamic query class updates	54
6	registry variable.	6	Nested query behavior	55
	Application development: Call Level Interface	6	Abnormal shutdown of the Historical Data	
	(CLI)	6	Generator	55
6	MapDateDescribe CLI/ODBC	6	Limitations by SQL statement type	56
6	configuration keyword	6	System Monitor	56
6	MapTimeDescribe CLI/ODBC	6	SYSMON_GROUP database manager	
6	configuration keyword	6	configuration parameter replaces	
6	MapTimestampDescribe CLI/ODBC	6	DB2_SNAPSHOT_NOAUTH registry	
6	configuration keyword	6	variable	56
	Command Reference	5	SQL Reference	56
6	db2inidb - Initialize a mirrored database	5	SQL limits	56
6	command	5	XML Extender	56
	Data Warehouse Center	6	XMLVARCHAR UDT size limit removed	56
6	Updates to the Business Intelligence	6	New environment variable:	
6	Tutorial	6	DB2DXX_MIN_TMPFILE_SIZE.	57
6	Setting the purge limit for warehouse log			
6	files			
4	Data Warehouse Center support for			
4	CURSOR load			
	GUI tools			
	Control Center plug-in support			
	Message Reference			
5	Version 8.1.4 Information Center message			
5	topics			
5	ADM message updates			

Appendix A. DB2 UDB FixPak CD directory structure 59

Appendix B. Contacting IBM 63
Product information 63

Appendix C. Notices. 65
Trademarks 68

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

Version information:

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 that follows this section.

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.

6 Directory paths:

6 Windows-based systems use backslashes (\) to delimit directories in a
6 directory path. UNIX and Linux-based systems use forward slashes (/) as

6 delimiters. The Release Notes follow this convention when the information is
6 platform specific. However, when the information is platform independent,
6 you might need to enter the directory path differently than shown. For
6 example, if you have a Windows-based system, you will have to enter the
6 directory path using backslashes (\) if the Release Notes show forward
6 slashes (/). Conversely, if you have a UNIX or Linux-based system, you will
6 have to enter the directory path using forward slashes (/) if the Release Notes
6 show backslashes (\).

6 **DB2 Information Integrator:**

6 For information about known issues with DB2 Information Integrator and its
6 associated technologies, including federated systems, Web services, and
6 metadata management, see the DB2 Information Integrator Release notes at:
6 <http://www.ibm.com/software/data/integration/db2ii/support.html>

Additional resources:

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>

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, you can install the DB2 HTML Information Center on your
system; a *DB2 HTML Documentation* CD image is available for download from
the same Web 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 image for
installation on your system. PDF documentation is updated less frequently
than the HTML Information Center.

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.

About this release

6 New in this release

6 This section provides a brief summary of the enhancements available with this
6 release of DB2 Universal Database.

6 **Support for additional collating sequences when a new Unicode database 6 is created**

6 DB2 UDB Version 8 FixPak 6 supports several new collating sequences when
6 a new Unicode database is created:

6 **UCA400_NO**

6 This collator implements the Unicode Collation Algorithm (UCA),
6 which is based on the Unicode Standard version 4.00 with
6 normalization implicitly set to ON.

6 **UCA400_LSK**

6 This collator also implements the UCA version 4.00, and is tailored to
6 sort the Slovakian characters.

6 **UCA400_LTH**

6 This collator also implements the UCA version 4.00, but it sorts all
6 Thai characters according to the Royal Thai Dictionary order.

6 Details of the UCA can be found in the Unicode Technical Standard #10,
6 available at the Unicode Consortium Web site: <http://www.unicode.org>.

6 When using the CREATE DATABASE command to create a new Unicode
6 database, you can now specify UCA400_NO, UCA400_LSK, or UCA400_LTH
6 for the COLLATE USING parameter.

6 When using the sqlcrea C API to create a new Unicode database, you can
6 now specify SQL_CS_UCA400_NO, SQL_CS_UCA400_LSK, or
6 SQL_CS_UCA400_LTH for the SQLDBCSS field of the SQLEDBDESC data
6 structure.

Product fix history

For a list of Authorized Program Analysis Reports (APARs) addressed by this
fix pack, visit the following Web page:

<http://www.ibm.com/cgi-bin/db2www/data/db2/udb/winos2unix/support/versionfixpak.d2w/report>

Backward compatibility

Alternate FixPaks (Linux and UNIX)

Prior to DB2 Universal Database Version 8, FixPaks functioned only as updates to installed DB2 Universal Database packages or file sets in one fixed location. Essentially, this meant that the installation of FixPaks replaced existing files with the updated ones provided within the FixPaks, and multiple DB2 FixPak levels on a single system was not possible. Now, DB2 UDB (ESE) can exist at multiple fix pack levels in the same system. This feature, supported in production operating environments since Version 8.1.2, is achieved using the following two FixPak types:

regular FixPaks

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

alternate FixPaks

- Can be installed as completely new copies of DB2 Universal Database ESE
- Are installed in a predefined location other than the location used for a regular DB2 Universal Database installation

Notes:

1. You are *not* required to perform a multiple FixPak installation if it is unnecessary for your environment.
2. Starting with IBM DB2 Universal Database Enterprise Server Edition (ESE) for Linux and UNIX[®], Version 8.1.2, fix packs are supported in production operating environments when they are installed as Multiple fix packs.

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

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

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

Data Warehouse Center previous server support restrictions

The following limitations exist for previous 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 cannot work with LOBs. You must upgrade the warehouse control database to the correct level, or you can move the control database to the system on which the DB2 UDB Enterprise Server Edition Version 8 warehouse server is installed and use the warehouse control database locally from that system.
- To move LOBs between the Data Warehouse Center and DB2 UDB, 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 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* statement followed by SELECT and INSERT statements. 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 the EXPORT and LOAD utilities.

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 UDB
4 Version 8.

4 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 UDB Version 8 if it is not
4 already at version 8 level. One way to migrate the database is from the
4 command 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.

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.

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

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

DB2 UDB for z/OS, Version 7

- 2 • PQ65125 - Provides SQLJ support for building Java SQLJ stored
2 procedures
- 6 • PQ76858 - Provides SQL Assist support

DB2 UDB for OS/390®, Version 6

- 2 • PQ76858 - Provides SQL Assist support
6

5 **Command Center and Version 7 servers**

5 The Version 8 Command Center might generate warning messages and dump
5 files when the [...] (ellipsis) button associated with the Database connection
5 field is clicked. This behavior is associated with the limitations and restrictions
5 for Version 8 tools working with Version 7 servers. Clicking [...] opens the
5 Select Database window. As you expand the systems and instances that
5 appear in this window, DB2 UDB generates internal actions to retrieve system,
5 instance, and database information in order to populate the tree. If DB2 UDB
5 encounters a Version 7 server during these internal activities, it generates a
5 warning and dump files.

Two versions of SQL Assist are launched from DB2 UDB

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

Change in Unicode server behavior

In version 7, Unicode servers ignored any graphic code pages sent 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.

Classic Connect unavailable

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

Product compatibility

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 product versions:

- 2 • DB2 Universal Database for z/OS and OS/390, Version 8
- 2 • DB2 Universal Database for iSeries, Versions 5.1 and 5.2

6 **WebSphere Application Server's execution of SQL statements in an SQLJ 6 program**

6 If you are using a version of WebSphere Application Server which is earlier
6 than version 5.0.1, all SQL statements in an SQLJ program are executed
6 dynamically, regardless of you customizing the SQLJ program.

6 For WebSphere Application Server 5.0.1 and later, if you customize your SQLJ
6 program, SQL statements are executed statically.

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.

CLI and ODBC on Windows 64-bit operating systems

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

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. It is available for Windows NT, Windows 2000, Windows 2003, and AIX operating systems, and Solaris Operating Environment. SecureWay Directory is shipped as part of the base operating system on AIX and iSeries (AS/400), and with OS/390 Security Server.

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

DB2 UDB 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.

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

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

Windows XP operating systems

The Windows XP Home Edition operating system is supported only by Personal Edition products.

The Windows XP Professional operating system is supported by the following products:

- Personal Edition
- Workgroup Server Edition
- DB2 Connect Personal Edition
- DB2 Connect Enterprise Edition

Note: DB2 Connect Enterprise Edition is supported on Windows XP for development or test purposes only. Production environments require Windows 2000 or Windows Server 2003.

Alternative Unicode conversion tables for the coded character set identifier (CCSID) 5039

The Microsoft Japanese Windows Shift-JIS code page is registered as the IBM coded character set identifier (CCSID) 943. However, the Shift-JIS code page on HP-UX platform is registered as CCSID 5039. CCSID 5039 contains characters in the Japanese Industry Standard (JIS) only, and does not have any vendor defined character. You can use a DB2 UDB database of CCSID 5039 on HP-UX to store Shift-JIS characters, but there will be code page conversion between CCSID 5039 and CCSID 943. When using Microsoft ODBC applications, you might encounter a potential problem when converting data in CCSID 5039 to Unicode, due to differences between IBM's code page conversion table and Microsoft's code page conversion table.

The following list of characters, when converted from CCSID 5039 to Unicode, will result in different code points depending on which conversion table is used (IBM or Microsoft). For these characters, the IBM conversion table conforms to the Japanese Industry Standard JISX0208 and JISX0221.

Table 2. CCSID 5039 to Unicode code point conversion

Shift-JIS code point (character name)	IBM primary code point (Unicode name)	Microsoft primary code point (Unicode name)
X'815C' (EM dash)	U+2014 (EM dash)	U+2015 (Horizontal bar)
X'8160' (Wave dash)	U+301C (Wave dash)	U+FF5E (Fullwidth tilde)

Table 2. CCSID 5039 to Unicode code point conversion (continued)

Shift-JIS code point (character name)	IBM primary code point (Unicode name)	Microsoft primary code point (Unicode name)
X'8161' (Double vertical line)	U+2016 (Double vertical line)	U+2225 (Parallel to)
X'817C' (Minus sign)	U+2212 (Minus sign)	U+FF0D (Fullwidth hyphen-minus)

For example, the character EM dash with the CCSID 5039 code point of X'815C' is converted to the Unicode code point U+2014 when using the IBM conversion table, but is converted to U+2015 when using the Microsoft conversion table. This can create potential problems for Microsoft ODBC applications because they would treat U+2014 as an invalid code point. To avoid these potential problems, DB2 UDB provides the alternate Microsoft conversion table from CCSID 5039 to Unicode, in addition to the default IBM conversion table. You need to replace the default IBM conversion table with the alternate Microsoft conversion table. Note that the default IBM conversion table from Unicode to CCSID 5039 matches the Microsoft version.

Replacing the Unicode conversion tables for coded character set (CCSID) 5039 with the Microsoft conversion tables

When you convert from CCSID 5039 to Unicode, the DB2 UDB default code page conversion table is used. If you want to use a different version of the conversion table, such as the Microsoft version, you must manually replace the default conversion table (.cnv) file.

Prerequisites:

Before replacing the existing code page conversion table file in the `sqllib/conv` directory, you should back up the file in case you want to change it back. On UNIX and Linux, the `sqllib/conv` directory is linked to the install path of DB2 UDB.

Restrictions:

For this to be effective, every DB2 UDB client that connects to the same database must have its conversion table changed. Otherwise the different clients might store the same character using different code points.

Procedure:

To replace the DB2 UDB default conversion table for converting from CCSID 5039 to Unicode, follow these steps:

1. Copy `sqllib/conv/ms/5039ucs2.cnv` to `sqllib/conv/5039ucs2.cnv`

6 2. Restart DB2 UDB.

6 **Alternative Unicode conversion tables for the coded character set**
6 **identifier (CCSID) 954**

6 The IBM coded character set identifier (CCSID) for the Japanese EUC code
6 page is registered as CCSID 954. CCSID 954 is a common encoding for
6 Japanese UNIX and Linux platforms. When using Microsoft ODBC
6 applications to connect to a DB2 UDB database of CCSID 954, you might
6 encounter a potential problem when converting data from CCSID 954 to
6 Unicode. The potential problem is due to differences between IBM's code
6 page conversion table and Microsoft's code page conversion table. The IBM
6 conversion table conforms to the character names as specified in the Japanese
6 Industry Standard (JIS) JISX0208, JISX0212, and JISX0221.

6 The following characters, when converted from CCSID 954 to Unicode, will
6 result in different code points depending on whether the IBM or Microsoft
6 conversion table is used.

6 *Table 3. CCSID 954 to Unicode code point conversion*

EUC-JP code point (character name)	IBM primary code point (Unicode name)	Microsoft primary code point (Unicode name)
X'A1BD' (EM dash)	U+2014 (EM Dash)	U+2015 (Horizontal Bar)
X'A1C1' (Wave dash)	U+301C (Wave Dash)	U+FF5E (Fullwidth Tilde)
X'A1C2' (Double vertical line)	U+2016 (Double vertical line)	U+2225 (Parallel To)
X'A1DD' (Minus sign)	U+2212 (Minus sign)	U+FF0D (Fullwidth hyphen-minus)
X'8FA2C3' (Broken bar)	U+00A6 (Broken bar)	U+FFE4 (Fullwidth broken bar)

6 For example, the character EM dash with the CCSID 954 code point of
6 X'A1BD' is converted to the Unicode code point U+2014 when using the IBM
6 conversion table, but is converted to U+2015 when using the Microsoft
6 conversion table. Due to this difference of conversion mapping, you might
6 have two different code points for the same character in a DB2 UDB Unicode
6 database, or in a graphic column of a DB2 UDB 954 database. This can create
6 potential problems for Microsoft ODBC applications because they would treat
6 U+2014 as an invalid code point. To avoid these potential problems, DB2 UDB
6 provides the alternate Microsoft conversion table from CCSID 954 to Unicode,
6 in addition to the default IBM conversion table. You need to replace the
6 default IBM conversion table with the alternate Microsoft conversion table.
6 Note that the default IBM conversion table from Unicode to CCSID 954
6 matches the Microsoft version.

Replacing the Unicode conversion tables for coded character set (CCSID) 954 with the Microsoft conversion tables

When you convert from CCSID 954 to Unicode, the DB2 UDB default code page conversion table is used. If you want to use a different version of the conversion table, such as the Microsoft version, you must manually replace the default conversion table (.cnv) file.

Prerequisites:

Before replacing the existing code page conversion table file in the `sqllib/conv` directory, you should back up the file in case you want to change it back. On UNIX and Linux, the `sqllib/conv` directory is linked to the install path of DB2 UDB.

Restrictions:

For this to be effective, every DB2 UDB client that connects to the same CCSID 954 database must have its conversion table changed. If your client is Japanese Windows, whose ANSI code page is Shift-JIS (CCSID 943), you will also need to change the DB2 default conversion tables between CCSID 943 and Unicode to the Microsoft version. Otherwise, the different clients might store the same character using different code points.

Procedure:

To replace the DB2 UDB default conversion table for converting from CCSID 954 to Unicode, follow these steps:

1. Copy `sqllib/conv/ms/0954ucs2.cnv` to `sqllib/conv/0954ucs2.cnv`
2. Restart DB2 UDB.

To replace the DB2 UDB default conversion tables for converting between CCSID 943 and Unicode, follow these steps:

1. Copy `sqllib/conv/ms/0943ucs2.cnv` to `sqllib/conv/0943ucs2.cnv`
2. Copy `sqllib/conv/ms/ucs20943.cnv` to `sqllib/conv/ucs20943.cnv`
3. Restart DB2 UDB.

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.

DB2 UDB for AIX 4.3.3 and 5.1 and above require AIX C++ Version 6 Runtime

DB2 UDB Version 8.1.4 and above for AIX Versions 4.3.3 and 5.1 and above require installation of the AIX C++ Version 6 Runtime libraries on your AIX

5 system. Download the March 2003 C++ Runtime PTF update from the
5 following Web site:

5 [http://www-
5 1.ibm.com/support/docview.wss?rs=0&q=xlC.rte&uid=swg24004427&loc=en_US&cs=utf-
5 8&cc=us&lang=en](http://www-1.ibm.com/support/docview.wss?rs=0&q=xlC.rte&uid=swg24004427&loc=en_US&cs=utf-8&cc=us&lang=en)

5 Follow the installation instructions given at the previously listed Web site
5 before using DB2 Version 8.1.4 or later.

5 **Accessing Java shared libraries (Linux)**

5 To run Java stored procedures or user-defined functions, the Linux run-time
5 linker must be able to access certain Java shared libraries, and DB2 UDB must
5 be able to load both these libraries and the Java virtual machine. Since the
5 program that does this loading runs with `setuid` privileges, it looks for the
5 dependent libraries only in the `/usr/lib` directory.

5 **Note:** This topic indicates essential shared libraries that you need to link to.
5 Depending on the applications you are building and running, you
5 might need to link to additional shared libraries.

6 Create symbolic links in the `/usr/lib` directory to point to the Java shared
6 libraries.

5 For the IBM Developer Kit 1.3, you need symbolic links to `libjava.so`,
5 `libjvm.so`, and `libhpi.so`. You can create the symbolic links by running the
5 following commands as root:

```
5 cd /usr/lib  
5 ln -fs JAVAHOME/jre/bin/libjava.so .  
5 ln -fs JAVAHOME/jre/bin/classic/libjvm.so .  
5 ln -fs JAVAHOME/jre/bin/libhpi.so .
```

5 where *JAVAHOME* is the base directory for the SDK. If DB2 UDB cannot find
5 these libraries, you will get a -4301 error when trying to run a Java routine,
5 and there will be messages in the administration notification log about
5 libraries not found.

6 **Note:** An alternative method is to add the Java shared libraries to the
6 `/etc/ld.so.conf` file instead of creating links in the `/usr/lib` directory.
6 If you do this, you must run the `ldconfig` command as root after
6 changing the `/etc/ld.so.conf` file, otherwise it will not work, as the
6 call to the routine will not complete. This alternative method might still
6 not work in specific instances, also resulting in the routine not
6 completing. If the alternative method does not work, create the links in
6 the `/usr/lib` directory as instructed previously.

5

Backup and restore (Linux 390)

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

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 run only from the DB2 command line processor.
- You must fully qualify the XML and trace file names for the **-x** and **-t** parameters.

2

Enabling view docking when accessing the Development Center with Hummingbird Exceed

2

2

2

2

2

2

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

2

To enable the XTEST extension:

2

2

2

2

2

2

2

2

2

2

2

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

Installation, migration, upgrade, and configuration information

Installation information

3 **Installing DB2 UDB (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.

2 **Installing DB2 UDB (Linux)**

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

2 However, in this case, the installation leaves the JDK_PATH database
2 configuration parameter pointing to the Java 1.3 path, /opt/IBMJava2-13/. As
2 a result, none of the Java-dependant functionality, including the installation of
2 the DB2 Tools Catalog, will work.

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

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

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

5 **Installing additional Asian fonts (Linux)**

5 IBM offers additional font packages for Linux that contain additional
5 double-byte character set (DBCS) support for Asian characters. These font
5 packages are necessary with some versions of Linux that install only the fonts
5 required to display the country-specific or region-specific characters.

5 If you run the **db2setup** command and find missing characters in the DB2
5 setup wizard interface, it is likely that your Linux system does not have all
5 the necessary fonts installed. To enable the **db2setup** command to properly
5 refer to the fonts embedded in the installation CD, perform the following task:

5 1. Enter the following command:

5 export JAVA_FONTS=/<cdrom>/db2/<linux_platform>/java/jre/lib/fonts

5 where <cdrom> is the location of the install image and <linux_platform>
3 is a directory name with a *Linux* prefix.

5 2. Re-run the **db2setup** command.

5 If you notice missing characters when using the DB2 GUI tools after
5 installation, install the necessary fonts provided with the DB2 product. These
5 fonts can be found in the fonts directory on either of the following CDs:

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

5 In this directory, there are two typefaces available: Times New Roman
5 WorldType and Monotype Sans Duospace WorldType. For each typeface, there
5 is a country-specific or region-specific font. The following table lists the eight
5 fonts provided in compressed format in the fonts directory.

5 *Table 4. File names for additional Asian fonts*

Font typeface	Font file name	Country/Region
Times New Roman WT J	tnrwt_j.zip	Japan and other countries/regions
Times New Roman WT K	tnrwt_k.zip	Korea
Times New Roman WT SC	tnrwt_s.zip	China (Simplified Chinese)
Times New Roman WT TC	tnrwt_t.zip	Taiwan (Traditional Chinese)
Monotype Sans Duospace WT J	mtsansdj.zip	Japan and other countries/regions
Monotype Sans Duospace WT K	mtsansdk.zip	Korea
Monotype Sans Duospace WT SC	mtsansds.zip	China (Simplified Chinese)
Monotype Sans Duospace WT TC	mtsansdt.zip	Taiwan (Traditional Chinese)

5 **Note:** These fonts do not replace the system fonts. These fonts are to be used
5 in conjunction with or for use with DB2 Universal Database. You
5 cannot engage in the general or unrestricted sale or distribution of
5 these fonts.

5 To install a font:

- 5 1. Unzip the font package.

2. Copy the font package to the /opt/IBMJava2-131/jre/lib/fonts directory. You need to create the directory if it does not already exist.
3. Enter the following command:

```
export JAVA_FONTS=/opt/IBMJava2-131/jre/lib/fonts
```

As a minimum, you need to install one font of each typeface for your country or region. If you are in China, Korea, or Taiwan, use the country-specific or region-specific versions; otherwise, use the Japanese version of the fonts. If you have space on your system, install all eight fonts.

Restrictions to adding products using the db2setup command

Once a DB2 product has been installed, additional DB2 products can be added. If you use the **db2setup** command to add products the following recommendations and restrictions apply.

Recommendations:

Both the installed product and the product being added should be at the same code level. For example, DB2 ESE Server Version 8 FixPak 5 is already installed and you want to add the DB2 Information Integrator product. In this case, DB2 Information Integrator should also be at the Version 8 FixPak 5 level.

Restrictions:

- If the DB2 fixpak level is higher than the fixpak level of the product being added, the combination is allowed. However, since the fixpak level of the product being added is lower than the DB2 fixpak level, the DB2 fixpak level must be reapplied after installing the additional product. See the appropriate FixPak Readme for instructions to reapply the fixpak.
- If the DB2 fixpak level is lower than the fixpak level of the product being added, an error is generated. The product being added cannot be at a higher fixpak level than DB2. In this case, you must first bring DB2 to the appropriate level, then install the additional product. See the appropriate FixPak Readme for instructions.

The following table summarizes the db2setup combinations:

Table 5. db2setup combinations

DB2 fixpak level	Additional product fixpak level	Is this combination permitted?
Version 8 FixPak 3	Version 8 FixPak 3	Yes. This is recommended.

Table 5. db2setup combinations (continued)

DB2 fixpak level	Additional product fixpak level	Is this combination permitted?
Version 8 FixPak 3	Version 8 GA	Yes, but Version 8 FixPak 3 must be reapplied. See the appropriate FixPak Readme for instructions to reapply the fixpak by going to the DB2 Support Web site.
Version 8 FixPak 3	Version 8 FixPak 5	No. DB2 must first be brought to the higher fix pack level (in this example, Version 8 FixPak 5) before installing the additional product. See the appropriate FixPak Readme for instructions to install the required Version 8 FixPak by going to the DB2 Support Web site.

The DB2 Support Web site address is
<http://www.ibm.com/software/data/db2/udb/winos2unix/support>

Downloading and uncompressing Spatial Extender FixPaks

Starting with DB2 Version 8 for Windows Fixpak 3, IBM is providing product-specific FixPaks instead of one general FixPak. This change affects only DB2 Version 8 products on Windows platforms.

If you have DB2 Version 8 and DB2 Spatial Extender for Windows installed on the same system, you must download the DB2 FixPak image and the Spatial Extender FixPak image, and you must uncompress each image under a common parent directory. All images must be uncompressed for the GUI installation or silent installation to proceed.

For complete FixPak installation instructions, see the DB2 Version 8 FixPak 6 Readme.

Binding Query Patroller packages after applying fixpaks

If you have Query Patroller installed, after applying a fixpak and performing all post-fixpak installation tasks, perform the following steps:

1. Log in as a user with DBADM authority.
2. Change to the proper directory:

- INSTPATH/bnd (Linux and UNIX)
- INSTPATH\bnd (Windows)

where INSTPATH is the DB2 UDB instance path.

3. Run the following commands:

```
db2 connect to dbname
db2 bind @qpserver.lst blocking all grant public
```

Installing Query Patroller at FixPak 3 level or later

Query Patroller is a query management system used to control the flow of queries against your DB2 UDB database. In DB2 UDB Version 8.1.2, DB2 Query Patroller became a stand-alone product. It is no longer a component of the DB2 Warehouse Manager.

If you have DB2 Version 8 FixPak 3 or later installed and you install the DB2 Query Patroller base or GA version, you must reapply DB2 UDB FixPak 3 or later. Otherwise, the Query Patroller changes are not applied to the DB2 UDB FixPak level you are at.

If you are installing the Query Patroller server, you must also update the DB2 instances after reapplying the DB2 UDB FixPak level. The instances must be restarted after they are updated.

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

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

The DB2 Setup wizard can install the DB2 HTML documentation only 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 drive.
2. Copy the following directory to your computer:
 - `/cdrom/program files/IBM/SQLLIB/doc/htmlcd/language`

3
3

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 installation, not to the documentation that you copied manually.

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.

Installing MDAC files for national language versions of DB2 UDB

If you do not install the national language version of Microsoft Data Access Components (MDAC) 2.7 prior to installing the national language version of DB2 Version 8.1, 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 in a language other than English. To fix this problem, install the "MDAC 2.7 RTM - Refresh" bundle from the Microsoft Web site at <http://msdn.microsoft.com/data/downloads/updates/default.aspx>. Choose the language that you want to install, download the required executable file, and run it. This installs the translated ODBC Data Source Administrator files.

Migration information

1
1
1
1
1
1
1
1
1
1

Migrating DB2 Universal Database when using DataJoiner or replication

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

Migrating a DB2 Version 8 Windows 32-bit database to Windows 64-bit

This topic lists the steps for migrating your 32-bit DB2 Version 8 database on a 32-bit computer to a 64-bit database on a 64-bit Windows operating system.

Prerequisites:

- A 64-bit version of DB2 Version 8 must be installed on your 64-bit computer.
- Ensure that your 32-bit Windows system is running DB2 Version 8.

Procedure:

To migrate to DB2 Version 8 on Windows 64-bit:

1. Backup your DB2 Version 8 databases on your 32-bit Windows system.
2. Restore your DB2 Version 8 backup (created in step #1) on your 64-bit Windows system.

Note: In addition to migrating DB2 UDB from 32-bit systems to 64-bit systems, the following migrations are also possible:

- Migrating between versions of Windows
- Migrating between versions of DB2 UDB
- Migrating everything at once
- Migrating back to 32-bit

Detailed information is available in the following IBM Redbook: Scaling DB2 UDB on Windows Server 2003. This Redbook can be found at the following URL:

<http://publib-b.boulder.ibm.com/Redbooks.nsf/RedbookAbstracts/sg247019.html>

Migrating XML Extender from previous versions

If you are using an earlier version of DB2 XML Extender, you must migrate each database that is enabled for XML Extender before you use an existing XML-enabled database with the updated XML Extender release. Each new fix pack contains all of the previous fix packs' updates.

To migrate an XML enabled database and XML enabled columns, complete the following steps.

1. From the DB2 Command Line, enter:

```
db2 connect to database_name
db2 bind dxxinstall/bnd7@dxxMigv.lst
db2 bind dxxinstall/bnd/@dxxbind.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*

4 **Database migration (HP-UX on IA64)**

4 Database migration is not supported for DB2 UDB for HP-UX on IA64
4 throughout the Version 8.x releases.

4 Restoring a Version 7 DB2 backup image to a Version 8 instance is not
4 supported on DB2 UDB for HP-UX on IA64.

Known limitations, problems, and workarounds

The following are the currently known limitations, problems, and workarounds for DB2® Universal Database Version 8.1 FixPak 6. The information in this section applies only to the Version 8.1 FixPak 6 release of DB2 Universal Database™, and its supporting products. Any limitations and restrictions might or might not apply to other releases of the product.

Limitations

Development Center limitations for 64-bit operating systems

6 Debugging of Java stored procedures against a 64-bit server is not supported
6 by the Development Center. Debugging SQL stored procedures is supported
2 on 64-bit Windows operating systems. Developing OLE DB or XML User
2 Defined Functions is not supported for 64-bit servers.

Development Center (Linux)

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).

Debugging stored procedures with double quotation marks

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

Path settings needed 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 directories for your developer kit versions are written to your \$HOME/IBM/DB2DC/DB2DC.settings file when the Development Center starts for the first time. You can copy these directories into your \$USER.settings file and modify them with a Unicode editor, or you can create symbolic links to your developer kit directories in the default locations.

Development Center limitations to run and debug Java stored procedures simultaneously

6 The DB2 Universal Database Development Center does not support running
6 and debugging Java stored procedures simultaneously. In the Development
6 Center, you can either run multiple Java stored procedures at the same time
6 or debug a single Java stored procedure; you cannot run a Java stored
6 procedure while debugging another Java stored procedure. By default, the

KEEPFENCED database manager configuration keyword is set to KEEPFENCED=YES as required for debugging SQL stored procedures. When the keyword KEEPFENCED has the default value YES, the routine process is kept active and JVM port conflicts will occur. Java stored procedures will cause a JVM exception and will fail in any of the following circumstances when the default KEEPFENCED=YES database manager configuration setting is used:

- Building a Java stored procedure in the Development Center then debugging it
- One user running a Java stored procedure and another user debugging a Java stored procedure while the initial Java stored procedure is still running
- One user debugging a Java stored procedure and another user running a Java stored procedure while the initial Java stored procedure is being debugged

To work around this limitation, ensure that the KEEPFENCED database manager configuration keyword is set to KEEPFENCED=NO by running the following commands:

```
db2 update dbm cfg using KEEPFENCED NO
db2stop
db2start
```

When KEEPFENCED is set to NO, the db2fmp process will shut down when a Java stored procedure call is finished, and DB2 UDB will start a new db2fmp process to handle the next routine call. This ensures that there will be no existing JVM in debug mode when you start to debug a Java stored procedure.

The KEEPFENCED=YES setting is required to build SQL stored procedures for debugging and to debug SQL stored procedures. When KEEPFENCED=NO, you can still build and execute SQL stored procedures but you cannot debug them.

Using LOB locators on a DB2 UDB for z/OS or OS/390 server

If you intend to use LOB locators to access DBCLOB columns in DB2 UDB tables on a DB2 UDB for z/OS or OS/390 server, run the com.ibm.db2.jcc.DB2LobTableCreator utility on each of those servers to create tables that are needed for fetching LOB locators.

DB2LobTableCreator syntax:

```
►—java—java com.ibm.db2.jcc.DB2LobTableCreator—————►
►—url—jdbc:db2://server—[ :port ]—/database—user—user-ID—————►
```


6
6



6
6

DB2LobTableCreator option descriptions:

6

-url

6

Specifies the data source at which DB2LobTableCreator is to run. The variable parts of the -url value are:

6

jdbc:db2:

6

Indicates that the connection is to a server in the DB2 UDB family.

6

server

6

The domain name or IP address of the database server.

6

port

6

The TCP/IP server port number that is assigned to the database server. This is an integer between 0 and 65535. The default is 446.

6

database

6

A name for the database server.

6

database is the DB2 UDB location name that is defined during installation. All characters in this value must be uppercase characters. You can determine the location name by executing the following SQL statement on the server:

6

```
SELECT CURRENT SERVER FROM SYSIBM.SYSDUMMY1;
```

6

-user

6

Specifies the user ID under which DB2LobTableCreator is to run. This user must have authority to create tables in the DSNATPDB database.

6

-password

6

Specifies the password for the user ID.

6

-help

6

Specifies that the DB2LobTableCreator utility describes each of the options that it supports. If any other options are specified with -help, they are ignored.

6

Failures when reporting a successful login during a connect (AIX)

5

When using OS authentication on AIX, DB2 UDB will attempt to report a successful login to AIX upon successful authentication during a connect. Prior to Version 8 FixPak 5, if DB2 UDB was unable to report a successful login, the connection would fail despite the user being authenticated. Starting with Version 8 FixPak 5, the connection is allowed to continue and the failure will be logged in the db2diag.log file.

5

5

5

5

5

5

JDBC driver limitation (HP-UX)

4

The IBM DB2 Universal JDBC Driver cannot connect to databases that were created using the HP default character set, roman8. All SQLJ and JDBC

4

4

4 applications that use the universal JDBC driver must connect to a database
4 created with a different character set. If your LANG is set to "C" or to a
4 "roman8" locale, you must change it to the corresponding ISO locale. For
4 instance, if your LANG is set to de_DE.roman8, it must be changed to
4 de_DE.iso88591 using the following command:

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

4 To run the DB2 SQLJ and JDBC sample programs with the Universal JDBC
4 driver, you can create the SAMPLE database with the commands in the
4 following example, which uses the ISO locale for US English:

```
4 export LANG=en_US.iso88591  
4 db2 terminate  
4 db2sampl
```

4 **Note:** If the SAMPLE database exists, you must drop it before executing the
4 **db2sampl** command.

4 **SNA support limitations in Version 8**

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

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

Tools catalog database creation not supported (Linux AMD64)

The creation of the tools catalog database under a 64-bit DB2 UDB instance on Linux (AMD64) is not supported. Do not attempt to create the tools catalog under a 64-bit instance using any of the following methods:

- Installing DB2 UDB
- Updating a 64-bit instance using the **db2setup** command
- Issuing the **CREATE TOOLS CATALOG CLP** command after the installation is finished.

The creation of the tools catalog database under a 32-bit instance on Linux (AMD64) is supported as of Version 8.1.4.

Tools catalog database creation not supported (AIX, Solaris Operating Environments, and HP-UX)

The creation of the tools catalog is not supported during the installation of DB2 UDB against 64-bit instances on the hybrid 64-bit platforms. The hybrid platforms are:

- AIX
- Solaris Operating Environments
- HP-UX
- Other platforms that support both 32-bit and 64-bit instances

If you want to create a tools catalog against a 64-bit instance, this can be done after the installation of DB2 UDB either through the Command Line Processor using the **CREATE TOOLS CATALOG CLP** command or using the Control Center. You will also need to install a 64-bit IBM Developer Kit for Java for this operation. Refer to the DB2 Administration Server section of the *DB2 Administration Guide* for more details.

Workgroup Server Edition 64-bit server instance limitations

DB2 UDB Workgroup Server Edition and DB2 UDB Workgroup Server Unlimited Edition are not licensed for 64-bit server instances. With these products, you can create:

- 32-bit server instances
- 32-bit or 64-bit client instances.

Asynchronous execution of CLI

Asynchronous execution with CLI is not available.

NUM_LOG_SPAN configuration parameter on a partitioned database

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

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

Limitation to the dasdrop command in multiple fix pack 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 have only 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 situation 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 applies only 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 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 Information Catalog Center, which is dependent on the Manage Information Catalog wizard component of DB2 Warehouse Manager.

Japanese names of remote objects

Schema, table, and column names of remote sources in Japanese cannot contain certain characters. Unicode mapping differences could 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

Parameter restrictions:

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.

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.

Replicating and accessing Client Connect warehouse sources using the warehouse agent

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 database partition 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 database partition.

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.

Data Warehouse Center restriction for import and export

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

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

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

Initialize the warehouse control database after installing version 8.1.2 or later

If you apply a fix pack or install a new release of DB2 UDB or the Data Warehouse Center, you must migrate the existing control database to update the objects that it contains. Also, after you install version 8.1.2 or later, you must re-initialize the warehouse control database.

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. Starting with 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.

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 Task Center 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 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 and use a display palette of at least 32 colors.

Do not partition Information Catalog Center tables

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

1. Open a DB2 Command Line Processor and issue these commands:

- a. CREATE DATABASE PARTITION GROUP *partition_group_name*
ON DBPARTITIONNUM *partition_number*
- b. CREATE REGULAR TABLESPACE *tablespace_name*
IN DATABASE PARTITION GROUP *partition_group_name*
MANAGED BY SYSTEM USING ('*cname*')

where *partition_group_name* is the same in both commands.

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

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

Query Patroller limitations when DYN_QUERY_MGMT is disabled

Query Patroller cannot perform the following actions if the database configuration parameter DYN_QUERY_MGMT is disabled:

- Release queries from a held state
- Make a running or queued query run in the background when the query is in the foreground

If you attempt to release a query from held state, or change a foreground query to a background query when DYN_QUERY_MGMT is set to DISABLE, an error message will be displayed and the state of the query will not change. If held queries are scheduled to run and DYN_QUERY_MGMT is disabled at the time they start running, an error message is written to the qpdiag.log file and the queries are left in held state.

Query Patroller result tables now use DB2QPRT schema

Starting with FixPak 5, all new result tables are created in the schema DB2QPRT rather than in the schema of the submitter.

DROPIN privilege on the DB2QPRT schema is granted to operators whose profiles were created prior to installation of FixPak 5 and had either:

- The MONITORING privilege with edit authority
- The HISTORICAL ANALYSIS privilege with edit authority

DROPIN privilege on the DB2QPRT schema is granted the first time Query Patroller creates a result table in this schema.

5 Operators who are given MONITORING privilege with edit authority or
5 HISTORICAL ANALYSIS privilege with edit authority after the installation of
5 FixPak 5 are also granted DROPIN privilege on the DB2QPRT schema upon
5 creation or update of their profiles.

5 **ALTER WRAPPER SQL statement not supported**

5 You cannot alter a wrapper. The ALTER WRAPPER statement is not
5 supported.

2 **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.

3 **System monitor snapshot UDF restrictions**

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

3 SQL1427N An instance attachment does not exist.

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

- 4 • GET MONITOR SWITCHES
- 4 • UPDATE MONITOR SWITCHES
- 4 • RESET MONITOR

4 This limitation occurs because such commands use an INSTANCE ATTACH,
4 while the snapshot UDFs make use of DATABASE CONNECTs.

Known problems and workarounds

DB2 license policy for DB2 Universal Database Workgroup Server Edition

Although the *DB2 Quick Beginnings for Servers* guide 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.

IBM Developer Kit for Java 1.3.1 required on Linux (x86, 32-bit)

The IBM Developer Kit for Java 1.3.1 Service Release 4 is required for DB2 UDB to use the DB2 Control Center or to create and run Java applications, including stored procedures and user-defined functions. Only the IBM Developer Kit for Java is supported.

The IBM Developer Kit for Java is installed whenever a component that requires Java is installed. However, if the installer detects that the IBM Developer Kit for Java 1.3.1 is already installed, it will not install the SDK again. The IBM Developer Kit for Java 1.3.1 is installed in its own directory and will not overwrite any previous levels of the IBM Developer Kit for Java.

Restrictions:

The installation of the IBM Developer Kit for Java will be attempted only if you use one of the following DB2 UDB installation methods:

- GUI installation program (db2setup)
- Response file installation (db2setup -r response_file)

Procedure:

To manually install the IBM Developer Kit for Java, run the following command from the /cdrom/db2/linux/Java-1.3.1 directory:

```
rpm -ivh IBMJava2-SDK-1.3.1-4.0.i386.rpm
```

This command installs the IBM Developer Kit for Java to the /opt/IBMJava2-131 directory.

To verify that IBM Developer Kit for Java is installed, run the following command from the UNIX[®] shell prompt:

```
<path>/jre/bin/java -version
```

where <path> represents the path where Java is installed. For example, if your installation directory path is /opt/IBMJava2-131/, the command is:

```
/opt/IBMJava2-131/jre/bin/java -version
```

You should receive output similar to the following:

```
java version "1.3.1"  
Java(TM) 2 Runtime Environment, Standard Edition (build 1.3.1)  
Classic VM (build 1.3.1, J2RE 1.3.1 IBM build  
cxia32131-20030329 (JIT enabled: jitc))
```

The IBM Developer Kit for Java is also available from the IBM developerWorks Web site at <http://www.ibm.com/developerworks/java/jdk/index.html>.

The simplified Chinese locale (AIX)

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

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

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

5 Red Hat Version 8 and later (including Red Hat Enterprise Linux [RHEL]
5 versions 2.1 and 3) has changed the default code set for simplified Chinese
5 from GBK (code page 1386) to GB18030 (code page 5488 or 1392).

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

3 To work around this limitation, you have two options:

- 3 • You can override the Red Hat default code set from GB18030 to GBK and
3 the territory from US to China (whose territory ID is CN and territory code
3 is 86).
- 3 • 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, 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.

Merant Driver Manager incompatibility (UNIX)

There are incompatibilities with Unicode support when the Merant Driver Manager accesses the DB2 ODBC driver on UNIX. These incompatibilities cause the Merant Driver Manager to use Unicode even if the application did not request Unicode use. This situation can lead to problems with components such as the Data Warehouse Center, Information Catalog Manager, and MQSI, which require the Merant Driver Manager to support non-IBM data sources. You can use an alternate DB2 ODBC driver library without Unicode support enabled until a permanent solution is available.

An alternative DB2 ODBC driver library without Unicode support enabled is included with DB2 UDB Version 8.1 for AIX, HP-UX, and Solaris Operating Environment. To use this alternative library, you must create a copy of it, giving the copy the original DB2 ODBC driver library's name.

Note: The alternative (_36) library contains the Unicode functions required by the DB2 JDBC driver. Using this library allows JDBC applications, including WebSphere Application Server, to work successfully with DB2 UDB.

To switch to the non-Unicode ODBC library on AIX, HP-UX, or the Solaris Operating Environment, see the following instructions. Because this is a manual process, you must carry it out every time you update your product, including after the application of successive FixPak or modification level.

Procedure:

AIX

To create the necessary library on AIX:

1. As the instance owner, shut down all database instances using the **db2stop force** command.

- 5 2. Using the DB2 administration server (DAS) instance ID, shut
5 down the DAS instance using the **db2admin stop force** command.
- 5 3. Back up the original db2.o file in the /usr/lpp/db2_81/lib
5 directory.
- 5 4. Using root authority, issue the **slibclean** command.
- 5 5. Copy file db2_36.o to backup file db2.o, ensuring that ownership
5 and permissions remain consistent. Use the following commands:
5

```
cp db2_36.o db2.o
```


5

```
-r--r--r-- bin:bin for db2.o
```

5 To switch back to the original object, follow the same procedure using
5 the backup file instead of the db2_36.o file.

5 Solaris Operating Environment

5 To create the necessary library in a Solaris Operating Environment:

- 5 1. As the instance owner, shut down all database instances using the
5 **db2stop force** command.
- 5 2. Using the DB2 administration server (DAS) instance ID, shut
5 down the DAS instance using the **db2admin stop force** command.
- 5 3. Back up the original libdb2.so.1 file in the /opt/IBMDB2/V8.1/lib
5 directory.
- 5 4. Copy file libdb2_36.so.1 to backup file libdb2.so.1, ensuring that
5 ownership and permissions remain consistent. Use the following
5 commands:
5

```
cp libdb2_36.so.1 libdb2.so.1
```


5

```
-r-xr-xr-x bin:bin libdb2.so.1
```
- 5 5. Issue the **db2iupdt <instance>** command for each database
5 instance and the **dasupdt <das_instance>** command for the DAS
5 instance.

5 To switch back to the original object, follow the same procedure using
5 the backup file instead of the libdb2_36.so.1 file.

5 HP-UX

5 To create the necessary library on HP-UX:

- 5 1. Shut down all database instances using the **db2stop force**
5 command.
- 5 2. Shut down the DB2 administration server (DAS) instance using the
5 **db2admin stop force** command.
- 5 3. Back up the original libdb2.sl file in the /opt/IBMDB2/V8.1/lib
5 directory.

- 5 4. Copy file libdb2_36.sl to backup file libdb2.sl, ensuring that the
5 ownership and permissions remain consistent. Use the following
5 command to ensure the consistency:
5 cp libdb2_36.sl libdb2.sl
5 -r-xr-xr-x bin:bin for libdb2.sl
- 5 5. Issue the **db2iupdt <instance>** command for each database
5 instance and the **dasupdt <das_instance>** command for the DAS
5 instance.

5 To switch back to the original object, follow the same procedure using
5 the backup file instead of the libdb2_36.sl file.

5 Other UNIX Operating Systems

5 Contact IBM Support if you require assistance with DB2 UDB and the
5 Merant Driver Manager on other UNIX operating systems.

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 the following command:

```
db2 backup <d1fm_db><path>
```

where:

- <d1fm_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**.

4 **SQLFLAG (STD) precompiler option error**

4 If the SQLFLAG(STD) precompile option is enabled, it will cause the
4 following error: Abend C6 occurred while running Precompile program
4 DSNHPC

4 Remove the SQLFLAG (STD) precompile option when using the Development
4 Center to create SQL stored procedures to run on DB2 for z/OS, Version 8.

5 **DB2 UDB Path for SQL Routine Compile Command (Windows)**

5 The DB2_SQLROUTINE_COMPILE_COMMAND does not normally need to
5 be set. However, when it is set on Windows, there can be a problem caused
5 by the %DB2PATH% variable containing a space between "Program" and
5 "Files" in the default path "C:\Program Files\IBM\SQLLIB". This problem can
5 occur even though the entire command value is in quotation marks.

5 The workaround uses the short form for the %DB2PATH% value. In the
5 default case, this is "C:\PROGRA~1\IBM\SQLLIB". You can either set
5 %DB2PATH% to the short form of the path where DB2 UDB is installed:

```
5 set db2path=C:\PROGRA~1\IBM\SQLLIB
```

5 and run the DB2_SQLROUTINE_COMPILE_COMMAND with the default
5 value:

```
5 db2set DB2_SQLROUTINE_COMPILE_COMMAND="c1 -Ox -W2 -TC -D_X86_=1 -MD  
5 -I%DB2PATH%\include SQLROUTINE_FILENAME.c /link -dll  
5 -def:SQLROUTINE_FILENAME.def /out:SQLROUTINE_FILENAME.dll  
5 %DB2PATH%\lib\db2api.lib"
```

5 Or you can change the value of the command itself by substituting the short
5 form of the path value for %DB2PATH%:

```
5 db2set DB2_SQLROUTINE_COMPILE_COMMAND="c1 -Ox -W2 -TC -D_X86_=1 -MD  
5 -IC:\PROGRA~1\IBM\SQLLIB\include SQLROUTINE_FILENAME.c /link -dll  
5 -def:SQLROUTINE_FILENAME.def /out:SQLROUTINE_FILENAME.dll  
5 C:\PROGRA~1\IBM\SQLLIB\lib\db2api.lib"
```

5 The actual path used must be where DB2 UDB is installed. For example, if
5 DB2 UDB is installed in the same path on the D: drive, the setting would be
5 "D:\PROGRA~1\IBM\SQLLIB".

1 **Documentation search might fail unless all documentation categories are** 1 **installed (AIX)**

1 If you do not install all of the categories of documentation on the DB2 HTML
1 documentation CD, an attempt to search against "All topics" might fail. An
1 InvalidParameterException might be reported in the Java console of your
1 browser, and you will receive no search results.

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

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

1 **Documentation search problem with Java 2 JRE1.4.0**

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

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

- 1 • Upgrade your browser's JRE version to 1.4.1, available at
1 <http://java.sun.com/j2se/1.4.1/download.html>

- 1 • Downgrade your browser's JRE version to 1.3.x, available at
- 1 <http://www.ibm.com/software/data/db2/udb/ad/v8/java/>

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:

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

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

Table 6. 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 come with Indic fonts that can be used with DB2 GUI tools:

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

GUI tools are not supported for zSeries servers (Linux)

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

Information Catalog Center log file not generated when importing tag language files

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

When running db2icmimport from a command line:

- If output files were not generated (.xml, .out, .err, .log), there is probably an error on the command line. Verify that the first five arguments, which are UserId, Password, Database, Catalog, and Tagfile, are correct. View the syntax by entering the **db2icmimport** command. If this does not solve the problem, modify **db2icmimport** to capture the output of the **db2javit** command by using the **-g** option to save the output to a file. For example:

```
db2javit -j:com.ibm.db2.common.icm.tag.IcmImport -w: -i:  
-o:"-Xmx128m -Xms32m" -g:"d:\temp\myimport.trc" ...
```
- If a log file was not generated, there is usually a parsing error. Look at the .xml file and the .out file. If you can, insert a **":COMMIT.CHKPID(DEBUG)"** command at the beginning of the tag language file. This command generates debug report messages and checks the .xml and .out files for parsing errors.
- After parsing, errors should appear in the .log file. When the debug report is being generated, look at the .log file and the .out file for information.
- Always check the .err file to see if there is a run-time error.

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

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

Binding Query Patroller packages

If the Query Patroller packages are not bound after applying a fixpak, a user without DBADM authority or proper Query Patroller privileges can encounter the following error when using the Query Patroller Center or Query Patroller command line:

```
SQL0001N - Binding or precompilation did not complete successfully.
```

6 If you are using the Query Patroller Center, the SQL0001N error is logged in
6 qpdiag.log file. If you are using the Query Patroller command line, the
6 SQL0001N is returned to the console

6 Auto-bind code exists to initiate automatic binding. However, the automatic
6 binding fails when the connecting user does not have the necessary privileges
6 to execute all statements in the Query Patroller packages. A symptom of this
6 problem is missing folders in the Query Patroller Center.

6 To avoid this problem, the qpserver.lst packages should be bound manually
6 by a user with DBADM authority or necessary privileges after applying a
6 fixpak.

1 **Secure environments (Windows)**

1 You might 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 shown in the following information:

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

1 **Problem**

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

Second workaround

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

XML Extender sample programs renamed

Some XML Extender sample programs might have the same name as other installed programs. Accidentally invoking another program with the same name as an XML Extender sample program can damage your XML files. The following list shows the old XML Extender sample program names as well as new replacement program names that are less likely to cause conflicts. Make sure you use the new sample program names instead of the old ones to prevent damaging your XML files.

Table 7. Replacement sample programs for XML Extender (Windows)

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

Table 8. Replacement sample programs for XML Extender (UNIX)

Old Program (Do not use)	New Program (Use)
insertx	dxxisrt
retrieve	dxxretr
retrieve2	dxxretr2
retrievec	dxxretrc
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 previously 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 dxshrd.exe. On UNIX platforms, you need only 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 previous table.

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

You can now decompose documents that contain non-unique attributes 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>
    <ItemNo>yyyy-yyyy</ItemNo>
    <Quantity>4</Quantity>
    <UnitPrice>24.99</UnitPrice>
  </OrderDetail>
</Order>
```

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

```
<element_node name="Order">
  <RDB_node>
    <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 col-->
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 col-->
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 col-->
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 col-->
4         <attribute_node name="ID">
4             <RDB_node>
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 col-->
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      <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>

```

4 The contents of the tables would look like the following sample after the
4 preceding document is decomposed:

4 ORDER_TAB:

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

4 DETAIL_TAB:

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

5 **Note:** To map multiple elements and attributes to the same column of the
5 same table, define an alias for the table and use the alias in the DAD
5 <table> element of one of the mappings.

4 Differences between SNA and TCP/IP when using DB2 Connect

5 When connecting to an OS/390 system using SNA, the host VTAM layer
5 automatically flows a commit when a new connection is made. The automatic
5 commit allows the host side thread state to be inactive, and the thread
5 immediately becomes inactive.

5 However, when connecting to an OS/390 system using TCP/IP, there is no
5 automatic commit. The application itself must flow an explicit commit after
5 the connection to allow the thread to be inactive on the host. Without the
5 explicit commit, the thread is subject to an idle thread time out.

5 The suggested workaround is to rewrite the application so that it will perform
5 an explicit commit if the connection goes idle after the connection.

Documentation updates

Administration: Performance tuning

6 **Default value for DB2_ENABLE_BUFDPD registry variable**

6 The default value for the DB2_ENABLE_BUFDPD registry variable has been
6 changed to ON.

Application development: Call Level Interface (CLI)

6 Three new keywords have been added to the Call Level Interface (CLI):

- 6 • MapDateDescribe
- 6 • MapTimeDescribe
- 6 • MapTimestampDescribe

6 **MapDateDescribe CLI/ODBC configuration keyword**

6 **Keyword description:**

6 Controls the SQL data type returned when DATE columns and
6 parameter markers are described.

6 **db2cli.ini keyword syntax:**

6 MapDateDescribe = 0 | 1 | 2

6 **Default setting:**

6 The default SQL data type for DATE data is returned: SQL_DATE for
6 ODBC 2.0 or SQL_TYPE_DATE for ODBC 3.0.

6 **Usage notes:**

6

6 To control the SQL data type that is returned when DATE columns and
6 parameter markers are described, set MapDateDescribe as follows:

- 6 • 0 - to return the default SQL data type: SQL_DATE for ODBC 2.0 or
6 SQL_TYPE_DATE for ODBC 3.0
- 6 • 1 - to return the SQL_CHAR SQL data type
- 6 • 2 - to return the SQL_WCHAR SQL data type

6 Only the following DB2 CLI functions are affected by setting
6 MapDateDescribe:

- 6 • SQLDescribeCol()
- 6 • SQLDescribeParam()
- 6 • SQLGetDescField()

6 • SQLGetDescRec()

6 The DB2 CLI catalog functions are not affected by this keyword.

6 **MapTimeDescribe CLI/ODBC configuration keyword**

6 **Keyword description:**

6 Controls the SQL data type returned when TIME columns and
6 parameter markers are described.

6 **db2cli.ini keyword syntax:**

6 MapTimeDescribe = 0 | 1 | 2

6 **Default setting:**

6 The default SQL data type for TIME data is returned: SQL_TIME for
6 ODBC 2.0 or SQL_TYPE_TIME for ODBC 3.0

6 **Usage notes:**

6

6 To control the SQL data type that is returned when TIME columns and
6 parameter markers are described, set MapTimeDescribe as follows:

- 6 • 0 - to return the default SQL data type: SQL_TIME for ODBC 2.0 or
6 SQL_TYPE_TIME for ODBC 3.0
- 6 • 1 - to return the SQL_CHAR SQL data type
- 6 • 2 - to return the SQL_WCHAR SQL data type

6 Only the following DB2 CLI functions are affected by setting
6 MapTimeStampDescribe:

- 6 • SQLDescribeCol()
- 6 • SQLDescribeParam()
- 6 • SQLGetDescField()
- 6 • SQLGetDescRec()

6 The DB2 CLI catalog functions are not affected by this keyword.

6 **MapTimeStampDescribe CLI/ODBC configuration keyword**

6 **Keyword description:**

6 Controls the SQL data type returned when TIMESTAMP columns and
6 parameter markers are described.

6 **db2cli.ini keyword syntax:**

6 MapTimeStampDescribe = 0 | 1 | 2

6 **Default setting:**

6 The default SQL data type for TIMESTAMP data is returned:
6 SQL_TIMESTAMP for ODBC 2.0 or SQL_TYPE_TIMESTAMP for
6 ODBC 3.0.

6 **Usage notes:**
6

6 To control the SQL data type that is returned when `TIMESTAMP` columns and
6 parameter markers are described, set `MapTimestampDescribe` as follows:

- 6 • 0 - to return the default SQL data type: `SQL_TIMESTAMP` for ODBC 2.0 or
6 `SQL_TYPE_TIMESTAMP` for ODBC 3.0
- 6 • 1 - to return the `SQL_CHAR` SQL data type
- 6 • 2 - to return the `SQL_WCHAR` SQL data type

6 Only the following DB2 CLI functions are affected by setting
6 `MapTimeStampDescribe`:

- 6 • `SQLDescribeCol()`
- 6 • `SQLDescribeParam()`
- 6 • `SQLGetDescField()`
- 6 • `SQLGetDescRec()`

6 The DB2 CLI catalog functions are not affected by this keyword.

Command Reference

6 **db2inidb - Initialize a mirrored database command**

6 Do not issue the `db2 connect to database` command prior to issuing the
6 `db2inidb database as mirror` command.

6 Attempting to connect to a split mirror database prior to initializing it erases
6 the log files needed for rollforward later.

6 The `connect` sets your database back to the state it was in when you
6 suspended the database. If the database is marked as consistent at the time of
6 the suspend, DB2 UDB concludes there is no need for crash recovery and
6 empties the logs for future use. If this situation occurs, attempting to
6 rollforward causes a `SQL4970` error.

Data Warehouse Center

6 **Updates to the Business Intelligence Tutorial**

6 **Verifying that the `DWCTBC` and `TBC_MD` databases are registered with
6 ODBC:**

6 In Version 8, the control database, `TBC_MD` that is used in the tutorial, does
6 not need to be a system ODBC data source. However, the target database or
6 database source `DWCTBC` must be a system ODBC data source.

6 **Opening the Define Warehouse Source notebook:**

6 The procedure for opening the Define Warehouse Source notebook for the
6 Tutorial Relational Source has changed.

6 To open the Define Warehouse Source notebook for the Tutorial Relational
6 Source:

- 6 1. From the Data Warehouse Center window, right-click the **Warehouse**
6 **Sources** folder.
- 6 2. Click **Define** → **ODBC** → **DB2** → **DB2 Family**.

6 The Define Warehouse Source notebook opens.

6 **Opening the Define Warehouse Target notebook:**

6 The procedure for opening the Define Warehouse Target notebook has
6 changed.

6 To open the Define Warehouse Target notebook:

- 6 1. From the Data Warehouse Center window, right-click the **Warehouse**
6 **Targets** folder.
- 6 2. Click **Define** → **ODBC** → **DB2** → **DB2 Family**.

6 The Define Warehouse Target notebook opens.

6 **Setting the purge limit for warehouse log files**

6 The log file holds records until a designated count limit is reached. The
6 default count limit is 1000 records. Typically, each job that you run creates 12
6 to 15 log records. Set the purge limit to a number that meets your needs by
6 updating the **Purge log when total records equal** field on the Server tab of
6 the Warehouse Properties page.

4 **Data Warehouse Center support for CURSOR load**

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

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

GUI tools

Control Center plug-in support

The Control Center now supports custom folders. Custom folders can contain user-selected system or database objects. Creating Control Center plug-ins specifically for a custom folder is not supported, but plug-ins can be created

for the object contained in custom folders. For more information about Control Center plug-ins, see the topic Introducing the plug-in architecture for the Control Center.

Message Reference

Version 8.1.4 Information Center message topics

Version 8.1.4 of the DB2 Information Center does not contain the new and modified message topics. The new and modified message topics are available on the IBM Web site:

<http://publib.boulder.ibm.com/infocenter/db2help>

ADM message updates

ADM5530E is incorrectly categorized as an error. It should be categorized as a warning message. The message is logged as an error in the Event Log on Windows and in the Notification Log on UNIX. ADM5530E should be treated as a warning message.

SQL message additions

SQL20271W

SQL20271W The name at ordinal position "<number>" in the statement, with name "<column-or-param-name>", was truncated.

Explanation:

At least one name in the described statement was truncated. The first name that was truncated is identified by ordinal position "<number>" and name "<column-or-param-name>".

If you are performing a describe output of a prepared query, the ordinal position is relative to the select list column of the query.

If you are performing a describe output of a CALL statement, the ordinal position is relative to the OUT or INOUT parameters of the procedure to which the CALL resolved.

If you are performing a describe input of a CALL statement, the ordinal position is relative to the IN or INOUT parameters of the procedure to which the CALL resolved.

The column name or parameter name was either too long, or became too long after code page conversion.

User Response:

For a column name, if the exact name of the column is significant, change the table, view, or nickname so the column has a shorter name, or use a client whose code page does not cause expansion of the column name beyond the supported maximum length. For a

parameter name, if the exact name of the parameter is significant, change the procedure so the parameter has a shorter name, or use a client whose code page does not cause expansion of the parameter name beyond the supported maximum length.

sqlcode: +20271

sqlstate: 01665 The column name or parameter name was truncated.

DBI message updates

DBI1060E Invalid package name <"pkg-name">

Explanation:

An incorrect name has been entered. The package either does not exist or the name has been entered incorrectly.

User Response:

Check to see if the name of the given package exists on the distribution media. If so, examine the name to see if it has been misspelled. All package names should be in lowercase.

DBI1001I

Usage:

```
db2icrt [-a AuthType]
         [-p PortName]
         [-s InstType]
         [-w WordWidth]
         -u FencedID InstName
```

Explanation:

An invalid argument was entered for the **db2icrt** command. Valid arguments for this command are:

- h | -?** displays the usage information.
- d** turns debug mode on.
- a** AuthType is the authentication type (SERVER, CLIENT, or SERVER_ENCRYPT) for the instance.
- p** PortName is the port name or port number to be used by this instance.
- s** InstType is the type of instance to be created (wse, ese, or client).
- u** FencedID is the name of the user under which fenced UDFs and fenced stored procedures will be run. This flag is not required if only a DB2 UDB client is installed.
- w** WordWidth is the width, in bits, of the instance to be

5 created (31, 32 or 64). You must have the requisite
5 version of DB2 UDB installed (31-bit, 32-bit or 64-bit)
5 to be able to select the appropriate width. The default
5 width value is the lowest bit width supported by the
5 current version of DB2 UDB, platform and instance
5 type.

5 InstName is the name of the instance.

5 **User Response:**

5 For more information about this command, refer to the *Installation and*
5 *Configuration Supplement* manual. Enter the command again with the
5 proper options and arguments.

5 **DBI1170E The -w flag only accepts 31, 32 or 64 as possible input.**

5 **Explanation:**

5 When using the **db2icrt** or **db2iupdt** command, you can
5 specify a value of 31, 32 or 64 for the optional flag -w. You
5 specify -w 64 for **db2icrt** when creating a 64 bit instance. You
5 can also specify -w 64 for **db2iupdt** when updating a 31 or 32
5 bit instance to become a 64 bit instance. The -w flag is not
5 needed otherwise. The required bit width to be updated to
5 must be supported for the current version of DB2 UDB,
5 platform, and instance type.

5 **User Response:**

5 For more information about this command, refer to the
5 *Administration Guide: Implementation* manual. Enter the
5 command again with the proper options and arguments.

5 **DBI1956E**

5 **Usage:**

5 db2ilist [-w 31|32|64] [-p] [-a] [inst_name]

5 **Explanation:**

5 An incorrect argument was entered for the **db2ilist** command. Valid
5 arguments for this command are:

5 **-h** displays the usage information.

5 **-w 31|32|64** lists the 31 or 32 or 64 bit instances. The -w option can
5 be used with the -p option, and is superseded by the
5 -a option.

5 **-p** lists the DB2 UDB install path that an instance is
5 running from. The -p option can be used with the -a
5 option, and it is superseded by the -a option.

5 **-a** returns all relevant information including the DB2
5 UDB install path associated with an instance, as well
5 as its bitwidth information (32 or 64).

5 **Note:** The returned information for 32 means 31-bit
5 for DB2 on Linux (S/390, zSeries)

5 **inst_name** returns the information for the specified instance. If no
5 instance is named, **db2ilist** returns information on all
5 instances of the current DB2 UDB release.

5 **User Response:**
5 Enter the command again as follows:
5 db2ilist [-w 31|32|64] [-p] [-a] [inst_name]

Query Patroller

Create Explain tables before running Query Patroller Historical Data Generator

When running the Historical Data Generator for Query Patroller, if the Explain tables do not already exist, the generator will create them for you. However, it is strongly recommended that you create the Explain tables before running the Historical Data Generator. When you create the Explain tables be sure you create them on the same partition. Actively creating the Explain tables on the same partition improves the performance of the Explain facility. This improvement increases the performance of the Historical Data Generator.

Checking Query Patroller log files for historical analysis

If the **Explain Run** column of the Query Activity over Time (Historical Analysis) report shows a status of **Ran unsuccessfully** for a query, historical data has not been generated for that query. Therefore, the query will not appear in any historical analysis reports or graphs. As documented in version 8, to determine why the query was unsuccessful, you can examine the qpuser.log file.

In addition to examining the qpuser.log file, you should examine the qpdiag.log file.

Dynamic query class updates

Certain query class operations no longer require Query Patroller to be stopped and restarted to take effect.

In the table that follows, an active query is a query whose status is Running or Queued.

Table 9. Conditions for query class changes to take effect

Nature of change	Conditions for change to take effect
Addition, removal, or update of a query class.	If there are no active queries, changes take effect immediately.
An update to a query class that involves only a change to the Maximum number of queries .	Takes effect immediately, even if there are active queries.
An update to a query class that involves only a change to the Maximum cost of a query .	<p>If there are active queries, the update takes effect when either:</p> <ul style="list-style-type: none"> • Query Patroller is stopped and restarted. • There are no more active queries. <p>Note: When there is a change pending against Maximum cost of a query, subsequent query class updates of any kind will not take effect until one of the two previous conditions is met.</p>
Addition or removal of a query class.	<p>If there are active queries, the addition or removal takes effect when either:</p> <ul style="list-style-type: none"> • Query Patroller is stopped and restarted. • There are no more active queries.

Nested query behavior

Nested queries cannot be queued. Instead, a nested query will run immediately if it exceeds a threshold that would normally cause it to be queued.

Abnormal shutdown of the Historical Data Generator

If you run the Historical Data generator and shut it down in an abnormal way, you will receive an error the next time you attempt to run the Historical Data generator. Examples of abnormal shutdown include:

- DB2 UDB stops unexpectedly
- Issuing a `db2stop force` command
- Issuing a `killdb2` command

When the Historical Data generator shuts down abnormally, you must issue the following command before attempting to rerun the Historical Data generator:

```
qp -d database generate historical_data stop
```

where *database* identifies the database that the command is being run against.

6 Limitations by SQL statement type

6 Contrary to previous documentation, the queries with the following
6 statements can be queued:

- 6 • Queries containing static SQL with host variables
- 6 • Queries containing an identity value function (IDENTITY_VAL_LOCAL) or
6 a sequence value function, such as NEXT VALUE FOR or PREVIOUS
6 VALUE FOR

System Monitor

6 **SYSMON_GROUP database manager configuration parameter replaces 6 DB2_SNAPSHOT_NOAUTH registry variable**

6 Users who are part of the SYSMON database manager level group have the
6 authority to gain access to database system monitor data.

6 The SYSMON authority group replaces the DB2_SNAPSHOT_NOAUTH
6 registry variable as the means to enable users without system administration
6 or system control authorities to access database system monitor data. System
6 monitor data is accessed using the snapshot monitor APIs, CLP commands, or
6 SQL table functions.

6 You can set the SYSMON authority group using the SYSMON_GROUP
6 database manager configuration parameter.

SQL Reference

5 **SQL limits**

5 The maximum NPAGES in a buffer pool for 32-bit releases has been
5 expanded to 1 048 576.

XML Extender

6 **XMLVARCHAR UDT size limit removed**

6 You can now define the size of the XMLVARCHAR user-defined type prior to
6 enabling the database. Previously, the XML Extender failed when the database
6 was enabled and a previously defined XMLVARCHAR existed. Perform the
6 following steps to change the size of the XMLVARCHAR UDT of an already
6 enabled database:

- 6 1. Back up all data in the XML Extender-enabled database.
- 6 2. Drop all XML collection tables or XML column side tables.
- 6 3. Disable the database with the **dxadm disable_db** command.
- 6 4. Create the XMLVARCHAR user-defined type.

- 6 5. Enable the database with the `dxadm enable_db` command.
- 6 6. Recreate and reload the tables.

6 **New environment variable: DB2DXX_MIN_TMPFILE_SIZE**

6 DB2 XML Extender may place large documents in temporary files to avoid
6 using too much memory during processing. On systems with large amounts
6 of physical memory, it is possible to avoid moving documents to temporary
6 files, reducing the amount of input/output activity. The environment variable
6 DB2DXX_MIN_TMPFILE_SIZE instructs XML Extender to use memory
6 buffers, rather than temporary files, for processing documents smaller than
6 the specified value. The variable is applicable only on the server, not on a
6 client. If multiple physical nodes participate in a multi-node partition, the
6 variable can be set differently on each node, accurately reflecting the amount
6 of memory installed on each node. If the environment variable is not set,
6 documents larger than 128KB will automatically be placed into temporary
6 files during processing, while documents smaller than 128K will be processed
6 in memory.

Appendix A. DB2 UDB FixPak CD directory structure

Windows operating systems

The files on the FixPak CD are located as follows:

Table 10. Windows files

Files	Location
DB2 product files:	x:\db2
Installation Notes:	x:\doc\<<language>\install.txt
Installation Notes (HTML):	x:\doc\<<language>\install.htm
License files:	x:\db2\license
Release Notes:	x:\doc\<<language>\release.txt
Release Notes (HTML):	x:\doc\<<language>\db2ir\index.htm

UNIX operating systems

The files on the FixPak CD are located as follows

Table 11. UNIX files

Files	Location
DB2 product files:	/cdrom/db2
Installation Notes:	/cdrom/doc/<language>/install.txt
Installation Notes (HTML):	/cdrom/doc/<language>/install.htm
License files:	/cdrom/db2/license
Release Notes:	/cdrom/doc/<language>/release.txt
Release Notes (HTML):	/cdrom/doc/<language>/db2ir/index.htm

where:

- x: refers to your CD drive (Windows)
- /cdrom refers to your mount point (UNIX)
- <language> refers to the language directory, consisting of a five-character code that corresponds to one of the languages in the following list:

Table 12. Directory names and their corresponding languages

Directory	Language
ar_AA	Arabic
bg_BG	Bulgarian
cs_CZ	Czech
da_DK	Danish
de_DE	German
el_GR	Greek
en_US	English
es_ES	Spanish
fi_FI	Finnish
fr_FR	French
hr_HR	Croatian
hu_HU	Hungarian
it_IT	Italian
iw_IL	Hebrew
ja_JP	Japanese
ko_KR	Korean
nl_NL	Dutch
no_NO	Norwegian
pl_PL	Polish
pt_BR	Brazilian Portuguese
pt_PT	Portuguese
ro_RO	Romanian
ru_RU	Russian
sk_SK	Slovak
sl_SI	Slovenian
sv_SE	Swedish
tr_TR	Turkish
zh_CN	Simplified Chinese
zh_TW	Traditional Chinese

Notes:

1. The directory names might appear in uppercase or lowercase letters, depending on your operating system.

-
2. All of the directories in the preceding list might not appear on this CD because all language directories are not available on all CDs.

Appendix B. Contacting IBM

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

- 1-800-IBM-SERV (1-800-426-7378) for customer service
- 1-888-426-4343 to learn about available service options
- 1-800-IBM-4YOU (426-4968) for DB2 marketing and sales

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

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

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

Product information

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

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

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

- 1-800-IBM-CALL (1-800-426-2255) to order products or to obtain general information.
- 1-800-879-2755 to order publications.

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

Appendix C. Notices

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

COPYRIGHT LICENSE:

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

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

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

Trademarks

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

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

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

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

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

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

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

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



Printed in USA