

IBM DB2 Information Integrator



Release Notes

Version 8.1, Fix Pack 3

IBM DB2 Information Integrator



Release Notes

Version 8.1, Fix Pack 3

Before using this information and the product it supports, be sure to read the general information under “Notices” on page 57.

This document contains proprietary information of IBM. It is provided under a license agreement and Copyright law protects it. 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

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 2003. 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	v	Federated function and data source support issues	16
DB2 Information Integrator support on the Web	v	Data type restrictions	16
What's New for Version 8.1, Fix Pack 3 . . .	1	Unicode UTF-8 issues	19
Life sciences user-defined functions included . . .	1	Materialized query table issues	19
Getstats_nr tool	1	LIKE predicate issues	20
HMMER wrapper now supports CS and RF records in data sources	2	CHAR columns that contain DBCS data	21
Altering long data types to varchar data types . . .	2	BioRS wrapper support	21
Limiting the number of rows that are retrieved for Entrez queries	3	Entrez wrapper support	21
Known problems, limitations, and workarounds.	5	Microsoft SQL Server wrapper support	21
DB2 Information Integrator installation issues	5	ODBC wrapper support	23
Changing to a different edition of DB2 Information Integrator	5	Oracle wrapper support	24
DB2 Information Integrator installer fails when the directory name contains a blank character	7	Sybase wrapper support	25
Installing DB2 Information Integrator over versions of DB2 Universal Database or DB2 Connect higher than 8.1.2 (Linux only)	7	Extended Search wrapper restrictions	26
DB2 Information Integrator installer has English license text on Russian Linux installers	9	Teradata wrapper support	28
Changing library permissions	9	Accessing Microsoft Excel data sources	28
Setting up Oracle code page conversion	9	Using the CREATE TYPE MAPPING statement	28
Chinese code page GB 18030 updates	10	DB2LOOK command limitations generating federated DDL statements	29
Remote operations fail when they use the Oracle 9i client and NET8 wrapper on a federated server running Linux	10	The getstats and get_stats_nr tools	29
Installing life sciences user-defined functions	12	CREATE TABLE restrictions	30
New variable for Teradata in the db2dj.ini file	13	Error in left outer join queries to BLAST, Entrez, Extended Search, HMMER, and XML data sources	30
Updated information for DB2 Universal Database documentation	14	DB2 Control Center updates	31
Migration issues	14	Differences between federated online help and DB2 Control Center object names	31
Accessing DB2 family remote data sources after migration	14	db2updvd8 command	31
Migrating server mappings with unsupported settings for ODBC data sources	15	Federated client and server compatibility	32
		MQ user-defined function message string updates	32
		Corrections to the documentation	33
		DB2 Information Integrator Installation Guide	33
		DB2 Information Integrator Migration Guide	36
		DB2 Information Integrator Data Source Configuration Guide	50
		DB2 Information Integrator Federated Systems Guide	53
		DB2 Information Integrator Developer's Guide	54
		Notices	57
		Trademarks	59

About the Release Notes

The release notes contain the latest information about DB2 Information Integrator, Version 8.1, Fix Pack 3.

The release notes are divided into three parts. The first part highlights what is new in this release. The second part contains the details of the problems, limitations, and workarounds that were known at the time of publication and apply to the latest version of the product. Read these notes to become familiar with any known outstanding issues with this release of DB2 Information Integrator. The third part contains corrections and updates to previously released HTML, PDF, and printed documentation.

DB2 Information Integrator support on the Web

For more information about DB2 Information Integrator and the latest support issues, see the DB2 Information Integrator support Web page at: www.ibm.com/software/data/integration/db2ii/support.html.

What's New for Version 8.1, Fix Pack 3

The following information describes new functions and features for DB2® Information Integrator, Version 8.1, Fix Pack 3.

Life sciences user-defined functions included

Life sciences user-defined functions are basic and frequently used life sciences functions that simplify operations when you are querying life sciences data sources. Life sciences user-defined functions are now part of DB2 Information Integrator and are installed with DB2 Universal Database, Version 8 Fix Pack 3.

The life sciences user-defined functions include:

- A Back Translate function, which converts an amino acid sequence into a nucleotide sequence
- A GeneWise function, which compares a protein sequence to DNA
- Translate functions, which convert a nucleotide sequence into a peptide sequence
- Reverse functions, which reverse a nucleotide or amino acid sequence

The life sciences user-defined functions support pattern matching in nucleotide or amino acid sequences, and regular expression pattern matching, including over alignments. You can also translate nucleotide sequences to amino acid sequences when coding information is absent or might be in error.

For information about installing the life sciences user-defined functions, see “Installing life sciences user-defined functions” on page 12.

Getstats_nr tool

A new tool, **getstats_nr**, is now available for DB2 Information Integrator. **getstats_nr** enables you to gather statistical information and cost parameters for remote nonrelational data sources, and update the DB2 catalog. By running **getstats_nr** with typical query workloads, you can automate collecting and updating statistics and cost parameters.

For more information about the **getstats_nr** tool, see “The getstats and get_stats_nr tools” on page 29.

HMMER wrapper now supports CS and RF records in data sources

The HMMER wrapper has been enhanced to support data sources that contain Reference (RF) and Consensus Structure (CS) annotation flags.

Altering long data types to varchar data types

To enable insert and update functions for long data types, you can alter the long data types to the VARCHAR data type. Table 1 lists the long data type by data source that you can alter.

Table 1. Long data types by data source that can be altered to the varchar data type

Data source	Remote data type	Length	Local default data type	ALTER to VARCHAR
DRDA	long varchar	1–32672	CLOB	varchar
	long varchar for bit data	1–32672	BLOB	varchar for bit data
	clob	1–32672	CLOB	varchar
	blob	1–32672	BLOB	varchar for bit data
Oracle Net8	long	1–32672	CLOB	varchar
	long raw	1–32672	BLOB	varchar for bit data
Oracle SQLNET	long	1–32672	CLOB	varchar
	long raw	1–32672	BLOB	varchar for bit data
Informix	text	1–32672	CLOB	varchar
	byte	1–32672	BLOB	varchar for bit data
Sybase ctplib	text	1–32672	CLOB	varchar
	image	1–32672	BLOB	varchar for bit data
Mssql	text	1–32672	CLOB	varchar
	image	1–32672	BLOB	varchar for bit data
Teradata	char	32673–64000	CLOB	varchar
	varchar	32673–64000	CLOB	varchar
	byte	32673–64000	BLOB	varchar for bit data

Table 1. Long data types by data source that can be altered to the varchar data type (continued)

Data source	Remote data type	Length	Local default data type	ALTER to VARCHAR
	varbyte	32673–64000	BLOB	varchar for bit data

Limiting the number of rows that are retrieved for Entrez queries

A new server option, `MAX_ROWS`, can be used to limit the number of rows that are returned for a query that uses the Entrez wrapper.

Unlike the `FETCH FIRST N ROWS ONLY` clause of an SQL statement, which limits the number of rows that are returned to a user or an application, the `MAX_ROWS` server option enables you to limit the number of rows that can be retrieved from the NCBI Web site.

The `MAX_ROWS` option value is always used as an upper (maximum) limit to the number of rows that a query can retrieve. If a query attempts to retrieve more rows than what is specified in the `MAX_ROWS` option, the result set is truncated, and a warning message is issued.

You can set the `MAX_ROWS` server option when a server is created, or you can use the `ALTER SERVER` statement to change the option value.

The `MAX_ROWS` server option is not required. If you do not set the option, a default value is used. The specific default value that is used depends on your operating system. For Microsoft Windows operating systems, the default value is 2000 rows. For UNIX-based operating systems, the default value is 5000 rows.

You can specify only positive numbers and 0 (zero). When you set the option to be 0 (zero), you enable queries to retrieve an unlimited number of rows from the NCBI Web site. However, setting the `MAX_ROWS` server option to 0 (zero) or to a very high number can possibly impact your query performance.

Known problems, limitations, and workarounds

The following information describes the currently known limitations, problems, and workarounds for DB2® Information Integrator Version 8.1, Fix Pack 3. Any limitations and restrictions might or might not apply to other releases of the product.

DB2 Information Integrator installation issues

Changing to a different edition of DB2 Information Integrator

Before you change from one edition of DB2 Information Integrator to another, you must remove the existing DB2 Information Integrator license key and add the new DB2 Information Integrator license key to your system. If you move to DB2 Information Integrator Replication Edition, you must remove any nonrelational wrappers that are installed.

Each edition of DB2 Information Integrator has a separate license key except the DB2 Information Integrator Developer Edition. The DB2 Information Integrator Developer Edition registers the license key for the DB2 Information Integrator Advanced Edition. However, the licensing terms for each edition of DB2 Information Integrator differ depending on the edition that you install, regardless of the license key.

The license files for DB2 Information Integrator are:

DB2 Information Integrator Edition	License file name
DB2 Information Integrator Replication Edition	db2iire.lic
DB2 Information Integrator Standard Edition	db2iise.lic
DB2 Information Integrator Advanced Edition	db2iaae.lic
DB2 Information Integrator Developer Edition	db2iaie.lic

Prerequisites:

You must have administrator authority on the DB2 instance to remove the DB2 Information Integrator license key.

Procedure:

To change to a different edition of DB2 Information Integrator:

1. Log on to the system with a user ID that has administrator authority on the DB2 instance.
2. Close all open programs so that the DB2 Information Integrator installation program can update files as required.
3. From a command prompt, change directories to the directory where DB2 Universal Database is installed.

By default, DB2 Universal Database™ is installed in one of the following directories depending on your operating system:

- /usr/opt/db2_08_01 (AIX)
- /opt/IBM/db2/V8.1 (HP-UX, Linux, Solaris™ Operating Environment)
- \Program Files\IBM\SQLLIB (Windows)

4. Enter the following command to remove the old DB2 Information Integrator license key from your system:

```
db2licm -r db2ii
```

5. Start the DB2 Information Integrator launchpad. The length of time that the launchpad takes to start varies depending on your system configuration. It might take a while for the DB2 Information Integrator launchpad to open.

Windows® CD installations: For CD-based installations on Windows systems, insert the DB2 Information Integrator CD into your CD drive. The DB2 Information Integrator launchpad opens.

Windows network installations: For network installations on Windows systems, map to the network drive and the directory that you are installing DB2 Information Integrator from. Double-click **iiSetup.exe** to open the DB2 Information Integrator launchpad. The **iiSetup.exe** file is located in the root directory of the DB2 Information Integrator CD.

UNIX® CD or network installations:

- a. Mount the DB2 Information Integrator CD or navigate to the directory that you are installing DB2 Information Integrator from.
- b. At the prompt, enter the following command to start the DB2 Information Integrator installation wizard:

```
./iiSetup.bin
```

6. From the DB2 Information Integrator launchpad, click **Install Products**.
7. Read the license agreement carefully before you continue. The DB2 Information Integrator installation wizard will detect DB2 Universal Database on your system.
8. Optional: On the Product Selection page, select the wrappers that you want to install. If relational or nonrelational wrappers are already installed on your system, you do not need to install them again.
9. Click **Next**. Follow the prompts in the wizard to complete the installation.

See the *DB2 Command Reference* for information about the **db2licm** command.

For instructions on installing relational and nonrelational wrappers, see the topic "Fast track to installing DB2 Information Integrator and setting up a federated server and database" in the *IBM DB2 Information Integrator Installation Guide*.

DB2 Information Integrator installer fails when the directory name contains a blank character

If you copy the DB2 Information Integrator installer or DB2 installer to a directory path that contains a blank character in its name, the installation will fail. To avoid this problem, you take one of the following actions:

- Perform the DB2 installation using the CDs provided.
- Ensure that no blank characters are used in the directory path name in which the DB2 Information Integrator installer or DB2 installer is copied.

Installing DB2 Information Integrator over versions of DB2 Universal Database or DB2 Connect higher than 8.1.2 (Linux only)

If you have installed a version of DB2 Universal Database that is higher than 8.1.2, or a version of DB2 Connect™ that is higher than 8.1.2, and then install the DB2 Information Integrator relational wrappers or the DB2 Information Integrator nonrelational wrappers, the wrappers are not installed. A message appears in the DB2 Setup wizard (or in the log file during a silent installation procedure) that warns you about missing prerequisites.

The following example shows the message that appears when an attempt is made to install the DB2 Information Integrator nonrelational wrappers:

```
Missing Prerequisites
-----
DB2 Information Integrator Nonrelational Wrappers requires one of the following products to be installed:
DB2 UDB Enterprise Server Edition, DB2 Connect Enterprise Edition.
Please install one of these products before installing DB2 Information Integrator Nonrelational Wrappers
```

This message is incorrect, because DB2 Universal Database or DB2 Connect is already installed on your system. However, the DB2 Information Integrator wrappers do not get installed as a result of this warning message. Therefore, you must install the DB2 Information Integrator wrappers manually. You can use the `db2_install` script or the RedHat Package Manager to install the DB2 Information Integrator wrappers.

After you install the wrappers, you must reinstall the DB2 Version 8.1 FixPak that was originally used to update your DB2 Universal Database Version 8.1 or DB2 Connect Version 8.1 products. This reinstallation of the FixPak is necessary to ensure that all of your components are at the same FixPak level.

Finally, you must configure DB2 Information Integrator to access the remote data sources for the wrappers that you installed.

The rest of this topic describes how to manually install the DB2 Information Integrator wrappers. For information about installing DB2 Version 8.1 FixPaks, see the documentation that came with your FixPak. For information about configuring DB2 Information Integrator to access remote data sources, see "Installing DB2 Information Integrator and setting up a federated server and database" in the *DB2 Information Integrator Installation Guide*.

Installing wrappers by using the `db2_install` script

To install wrappers using the `db2_install` script:

1. Log in as a user with root authority.
2. Mount the DB2 Information Integrator CD.
3. Change your directory to `Linux/prod` directory on your CD, where *prod* is **rcon** for relational wrappers and **lsdc** for nonrelational wrappers.
4. Use the `./db2_install` command to start the `db2_install` script.
5. At the prompt, enter the product keyword. Enter `DB2.LSDC` to install nonrelational wrappers or `DB2.RCON` to install relational wrappers.

Installing wrappers by using RedHat Package Manager (RPM)

To install wrappers using the RedHat Package Manager (RPM):

1. Log in as a user with root authority.
2. Mount the DB2 Information Integrator CD.
3. Identify the components that you want to install.

Each DB2 Information Integrator product CD provides a file that lists the components that are available for installation. The component list is in a file called `ComponentList.htm`. Use this file to identify required, typical, and optional components.

The `ComponentList.htm` file is in the `Linux/prod/db2/linux` directory on your CD. *prod* is **rcon** for relational wrappers, and **lsdc** for nonrelational wrappers.

4. Change your directory to `Linux/prod/db2/linux` directory on your CD. *prod* is **rcon** for relational wrappers, and **lsdc** for nonrelational wrappers.
5. Install the component or components that you need by using the **rpm** command:

```
rpm -ivh component_name
```

component_name is the name of the component that you want to install.

You can specify multiple components with the **rpm** command.

For example, suppose you want to install the Signature for DB2 Information Integrator Nonrelational Wrappers component (filename: IBM_db2lssg81-8.1.0.16.i386.rpm) and the Structured Files Data Sources component (filename: IBM_db2lsstf81-8.1.0.16.i386.rpm). You can install both of these components at the same time by using the following command:

```
rpm -ivh IBM_db2lssg81-8.1.0.16.i386.rpm IBM_db2lsstf81-8.1.0.16.i386.rpm
```

DB2 Information Integrator installer has English license text on Russian Linux installers

A Java™ problem causes the license acceptance text on Russian Linux computers to be corrupted. In DB2 Information Integrator, Version 8.1, the Russian license text that is displayed by the installer on Linux computers is replaced with English text.

A new license directory has been added to the DB2 Information Integrator installer images that contains the license acceptance text for all supported languages, including Russian.

Changing library permissions

You do not need to use the **chmod** command to change the wrapper library permissions. The permissions are set automatically when you run a djxlink script such as djxlinkInformix.

The steps in the topic Checking the wrapper library file permissions (UNIX), are no longer necessary.

Setting up Oracle code page conversion

You can set the optional Oracle® environment variable, NLS_LANG, for Oracle code page conversion.

Each time that the Oracle wrapper connects to an Oracle data source, the wrapper determines which code page value to use for that connection. If the NLS_LANG environment variable is set in the db2dj.ini file on the federated server, then the wrapper uses the value in the db2dj.ini file. The db2dj.ini file contains configuration information about the Oracle client software that is installed on your federated server.

If the NLS_LANG variable is not set on the federated server, the wrapper determines the locale of the federated server and the code page of the federated database. The wrapper sets the NLS_LANG variable to the closest matching Oracle locale. If there is no closely matching locale, the NLS_LANG variable is set to American_America.US7ASCII.

See the documentation for your Oracle software for a list of valid locales.

To set the NLS_LANG Oracle environment variable:

1. Open the sqllib/cfg/db2dj.ini file in a text editor. If the file does not exist, create a file named db2dj.ini.

2. In the db2dj.ini file, set the NLS_LANG environment variable to the Oracle locale that you want to use:

```
NLS_LANG=<Oracle_client_locale_value>
```

<Oracle_client_locale_value> is the locale that you want to use.

For example, to set the locale to be Japanese, set NLS_LANG to:

```
NLS_LANG=Japanese_Japan.JA16SJIS
```

Chinese code page GB 18030 updates

If you access a data source that contains data that uses the Chinese code page GB 18030, your federated database must use the UTF-8 code page.

For Oracle data sources, the Oracle wrapper sets the Oracle NLS_LANG environment variable as follows:

```
NLS_LANG=Simplified_Chinese_China.UTF8
```

If you are using the Oracle 9i client, change the NLS_LANG setting in your sqllib/cfg/db2dj.ini file to Simplified_Chinese_China.AL32UTF8, so that the Oracle 9i client translates the GB 18030 data into Unicode correctly. For example:

```
NLS_LANG=Simplified_Chinese_China.AL32UTF8
```

For Informix™ data sources, the Informix wrapper sets Informix environment variables as follows:

```
CLIENT_LOCALE=zh_cn.UTF8
```

```
GL_USEGLU=1
```

You must add the following setting to your sqllib/cfg/db2dj.ini file so that the Informix client correctly translates the GB 18030 data to Unicode:

```
DB_LOCALE=zh_cn.GB18030-2000
```

If you want Informix to verify that the DB_LOCALE setting matches the actual locale of the Informix database, then you need to also add the following setting to your sqllib/cfg/db2dj.ini file:

```
DBNLS=1
```

Remote operations fail when they use the Oracle 9i client and NET8 wrapper on a federated server running Linux

If you use the following configuration, you must edit the genclntsh script and create the libclntsh.so file:

- The Oracle 9i client
- The Oracle NET8 wrapper

- A federated server that runs either the Linux, Solaris Operating Environment, or HP-UX operating system

Otherwise, your DB2 federated instance will fail when you attempt any remote operation that includes Oracle and the NET8 wrapper.

This topic gives the instructions for editing the genclntsh script on Linux. See the topics, "Editing the Oracle genclntsh script before the federated server is set up to access Oracle data sources" and "Editing the Oracle genclntsh script after the federated server is set up to access Oracle data sources," for instructions on editing the genclntsh script for the Solaris Operating Environment and the HP-UX operating system.

Prerequisites:

Back up the \$ORACLE_HOME/lib/libclntsh.so file and the \$ORACLE_HOME/bin/genclntsh file before you begin this procedure.

Procedure:

You can prevent remote operations from failing by editing the genclntsh script and creating the libclntsh.so file before or after you install DB2 Information Integrator.

To prevent remote operations from failing on Linux before you install DB2 Information Integrator:

1. Edit the \$ORACLE_HOME/bin/genclntsh script and add `-Wl,-Bsymbolic` to the link line. For example:

```
LD="gcc -shared -Wl,-relax -L${ORACLE_HOME}/lib -Wl,-Bsymbolic" # shared library link command
```

2. From a command prompt, run the genclntsh script to create the libclntsh.so file.

To prevent remote operations from failing on Linux after you install DB2 Information Integrator:

1. Edit the \$ORACLE_HOME/bin/genclntsh script and add `-Bsymbolic` to the link line. For example:

```
LD="ld -v -G -b +s -L${OLIB}-Bsymbolic " #shared library link command
```

2. From a command prompt, run the genclntsh script to create the libclntsh.so file.
3. From a command prompt, run the djxlinkOracle script to update the Oracle wrapper library.
4. Issue the db2iupdt command on each DB2 instance to enable federated access to the data sources.

Installing life sciences user-defined functions

The life sciences user-defined functions are basic and frequently used life sciences functions that simplify operations when you are querying life sciences data sources. The life sciences user-defined functions are part of DB2 Information Integrator and are installed with DB2 Universal Database, Version 8.1 Fix Pack 3.

The LSGeneWise user-defined function is available for AIX® 32-bit operating systems. The other life sciences user-defined functions are available for Windows and AIX 32-bit operating systems.

Before you use the LSGeneWise user-defined function, you must install the Wise2, Version 2.1.20 client.

The life sciences user-defined function library files are used when you register the functions during the configuration process. When you install the life sciences user-defined functions, the following library files are placed on your system:

Windows

sqllib/function/db2lsudfs.dll

AIX

- sqllib/function/libdb2lsudfs.a
- sqllib/function/libdb2lsgenewise.a

For the LSGeneWise user-defined function, the sqllib/lib/libdb2lsSTgenewise.a archive file is also installed on your system.

The following life sciences user-defined functions are installed with DB2 Universal Database, Version 8.1 Fix Pack 3:

- LSPep2AmbNuc
- LSPep2ProbNuc
- LSDefineParse
- LSPatternMatch
- LSPrositePattern
- LSGeneWise (AIX only)
- LSMultiMatch
- LSMultiMatch3
- LSBarcode
- LSTransAllFrames
- LSRevNuc
- LSRevPep

- LSRevComp
- LSNuc2Pep

Prerequisites:

You must install DB2 Information Integrator and the Life Sciences User-Defined Functions component of the nonrelational wrappers before you install DB2 Universal Database, Version 8.1 Fix Pack 3.

To install the life sciences user-defined functions:

1. Start the DB2 Information Integrator installation wizard.
2. Select the **Nonrelational wrappers** check box on the Product Selection window and click **Next**. Follow the instructions in the wizard to complete the installation. When you install the nonrelational wrappers, select the Life Sciences User-defined Functions component. You do not need to reinstall DB2 Universal Database or any wrappers that are already installed.
3. Install DB2 Universal Database, Version 8.1 Fix Pack 3.
4. If you plan to use the LSGeneWise user-defined function, run the `djxlinkLSGeneWise` script to create the `sqlib/function/libdb2lsgenewise.a` library.

See the *DB2 Information Integrator Addendum to the Data Source Configuration Guide: BioRS Wrapper and Life Sciences User-Defined Functions* for instructions on registering the life sciences user-defined functions. You can obtain this document on the Web at www.ibm.com/software/data/integration/db2ii/.

The life sciences user-defined functions are now ready to be registered. See the *DB2 Information Integrator Addendum to the Data Source Configuration Guide: BioRS Wrapper and Life Sciences User-Defined Functions* for instructions on registering the life sciences user-defined functions.

New variable for Teradata in the `db2dj.ini` file

The Teradata® wrapper has an additional variable, `TERADATA_CHARSET`, that you set in the `db2dj.ini` file. Set this variable to the suggested client character set. If you do not set this variable, DB2 Information Integrator detects the client character set based on the code page of the database. When you set the variable, DB2 Information Integrator uses its value as the client character set. The value in the `TERADATA_CHARSET` variable is not validated, but if it is not set to the correct value, the remote data source issues an error.

The following values are correct values for the `TERADATA_CHARSET` variable:

For UNIX systems:

- KanjiEUC_0U
- HANGULKSC5601_2R4
- SCHGB2312_1T0
- TCHBIG5_1R0
- UTF8

For Windows systems:

- KanjiSJIS_0S
- HANGULKSC5601_2R4
- SCHGB2312_1T0
- TCHBIG5_1R0
- UTF8

Updated information for DB2 Universal Database documentation

The file names in the following table supersede the file names in the "DB2 Universal Database documentation" topic in the *DB2 Information Integrator Installation Guide*.

Table 2. Updated file names for DB2 Universal Database documentation

Title and form number	PDF file name
<i>CLI Guide and Reference, Volume 1,</i> SC09-4849-00	db211x80
<i>CLI Guide and Reference, Volume 2,</i> SC09-4850-00	db212x80
<i>Installation and Configuration Supplement,</i> GC09-4837-00	db2iyx80

Migration issues

Accessing DB2 family remote data sources after migration

If you need to access a DB2 family remote data source after you migrate to DB2 Information Integrator from either DB2 Version 7 or from DataJoiner[®] Version 2.1.1, you must invoke the bind utility to reestablish a connection from the server to the remote data source. DB2 family remote data sources include databases that reside on DB2 for UNIX and Windows, DB2 for z/OS[™] and OS/390[®], DB2 for iSeries[™], and DB2 Server for VM and VSE.

If you do not reestablish the connection by using the bind utility, you receive a -805 error message.

You must invoke the bind utility to reestablish a connection from the server to the remote data source for each DB2 family remote database that you want to access. You need to perform this task only one time for each remote database that you want to access.

To use the bind utility to reestablish a connection to a remote data source after migration:

1. Connect to the remote data source.
2. Issue the following commands:

```
cd /home/instancename/sql1lib/bnd
db2 bind @db2cli.lst grant public
```

instancename is the name of the instance that you want to access.

3. Disconnect from the remote data source.

Migrating server mappings with unsupported settings for ODBC data sources

You might need to use the workarounds listed here to ensure that you can use your remote server mappings for ODBC data sources after you migrate to DB2 Information Integrator Version 8.

If you are migrating a DataJoiner server mapping with a server type or server version that DB2 Information Integrator does not support, you cannot use the server definition or change the server type or the server version on DB2 Information Integrator for ODBC data sources.

You can use the following workarounds if DB2 Information Integrator supports the version and type of your remote server, but your DataJoiner server mapping specifies a server type or version that does not match the server type or version of the remote data source.

- If you did not yet migrate your DataJoiner server mapping and only your server versions are incompatible with DB2 Information Integrator, change the DataJoiner server version by using the ALTER SERVER MAPPING statement. For example:

```
ALTER SERVER MAPPING FROM my_odbc_server SET VERSION 5
```

After you migrate the altered DataJoiner server mapping, you can use the server definition on DB2 Information Integrator.

Restriction: You cannot use the ALTER SERVER MAPPING statement to change the server type.

- If you already migrated your DataJoiner server mappings or if your DataJoiner server type is incompatible with DB2 Information Integrator, use the following steps to recreate the federated objects in the database:

1. Issue the **db2look** command.

For example:

```
db2look -d MYDB -wrapper odbc -e -o mydb.ddl
```

MYDB is the name of the database. odbc is the name of the wrapper.

This command generates an output file named mydb.ddl that contains the DDL statements for all the federated objects of the specified wrapper.

2. Edit the statements in the output file to recreate the federated objects.
 - a. Delete all CREATE WRAPPER statements.
 - b. Delete any CREATE SERVER statements for server definitions with correct versions and types.
 - c. Delete any CREATE NICKNAME, CREATE TYPE MAPPING, CREATE FUNCTION, CREATE USER MAPPING, and CREATE FUNCTION MAPPING statements for federated objects that belong to the server definitions in step 2b.

Step 2b and step 2c ensure that you do not delete and then recreate server definitions with correct versions and types.

If a small number of the server definitions have incorrect server versions or server types, then you can specify the `-server` option on the `db2look` command to generate DDL for only that server definition.

- d. Add a DROP SERVER statement before each CREATE SERVER statement to drop the server definition and all objects for that server.
 - e. Change each CREATE SERVER statement to specify the correct server type or the correct server version.
 - f. Specify the correct password for the REMOTE_PASSWORD option for each CREATE USER MAPPING statement.
3. Run the following command to issue the SQL statements in your output file and recreate the federated objects.

For example:

```
db2 -tvf mydb.ddl
```

Federated function and data source support issues

Data type restrictions

DB2 Information Integrator does not support some data types. You cannot create nicknames for data source objects (such as tables and views) that contain any of these unsupported data types. Additionally, DB2 Information Integrator does not allow insert, update, and delete operations on data source objects that contain certain data types.

Unsupported data types

You cannot create nicknames for data source objects that contain the data types listed in the following table:

Table 3. Unsupported data source data types

Data source	Unsupported data types
Extended Search	DECIMAL
Microsoft® SQL Server	SQL_VARIANT
Oracle (NET8 wrapper only)	NCHAR NVARCHAR2 TIMESTAMP (fractional_seconds_precision) WITH TIME ZONE TIMESTAMP (fractional_seconds_precision) WITH LOCAL TIME ZONE
Oracle (SQLNET wrapper only)	BLOB CLOB NCHAR NVARCHAR2 TIMESTAMP TIMESTAMP (fractional_seconds_precision) WITH TIME ZONE TIMESTAMP (fractional_seconds_precision) WITH LOCAL TIME ZONE
Sybase	unichar univarchar

For the data types that DB2 Information Integrator does not support, you can create a view at the data source that is based on the data source object and create a nickname for the view. The view must not contain any columns that use the unsupported data types. Alternatively, you can create a view that is based on the data source object and change the unsupported data types to supported data types.

Insert, update, and delete restrictions on data types

DB2 Information Integrator does not allow insert, update, and delete operations on data source objects that contain certain data types. You cannot perform write operations on data source objects that contain the data types listed in the following table:

Table 4. Write operation restrictions on data types

Data source	Unsupported data types
DB2 family (DRDA)	BLOB CLOB DBCLOB
Informix	BLOB CLOB

Table 4. Write operation restrictions on data types (continued)

Data source	Unsupported data types
Microsoft SQL Server	image ntext text SQL_VARIANT
ODBC	SQL_LONGBINARY (length > 255) SQL_LONGVARCHAR (length > 255) SQL_WLONGVARCHAR (length > 255)
Oracle (NET8 wrapper only)	INTERVAL DAY (day_precision) TO SECOND (fractional_seconds_precision) INTERVAL YEAR (year_precision) TO MONTH LONG LONG RAW NCHAR NVARCHAR2 TIMESTAMP (fractional_seconds_precision) WITH TIMEZONE TIMESTAMP (fractional_seconds_precision) WITH LOCAL TIME ZONE Exceptions: You can perform select, insert, update, and delete operations on LONG data types if you alter the column to the VARCHAR data type. You can perform select, insert, update, and delete operations on LONG RAW data types if you alter the column to the VARCHAR FOR BIT DATA data type.
Oracle (SQLNET wrapper only)	BLOB CLOB INTERVAL DAY (day_precision) TO SECOND (fractional_seconds_precision) INTERVAL YEAR (year_precision) TO MONTH NCHAR NVARCHAR2
Sybase (CTLIB wrapper only)	image text unichar univarchar
Sybase (DBLIB wrapper only)	All data types. Write operations are not supported by the DBLIB wrapper.
Teradata	char (32673–64000) varchar (32673–64000) byte (32673–64000) varbyte (32673–64000) graphic (16337–32000) vargraphic (16337–32000)

Mapping to the federated GRAPHIC and VARGRAPHIC data types

Only the DB2 family and Teradata data sources can override the default data type mappings to map remote data types to the federated GRAPHIC and VARGRAPHIC data types. You use either the CREATE TYPE MAPPING

statement or the ALTER NICKNAME statement to override the default data type mappings, depending on the circumstances in which you want the mapping to apply.

Unicode UTF-8 issues

UTF-8 code page data expansion

If the federated database uses the UTF-8 code page and the data source client converts to the UTF-8 code page, the conversion might result in data expansion. For example, a 1-byte character at the remote data source might be stored as 2 bytes in the federated UTF-8 database system catalog. Ensure that the local columns in the catalog are wide enough to contain any expanded data. If the columns are not wide enough, use the ALTER NICKNAME statement to make the columns wider.

Chinese code page GB 18030

If you access a data source that contains data that uses the Chinese code page GB 18030, your federated database must use the UTF-8 code page.

For Oracle data sources, the Oracle wrapper sets the Oracle NLS_LANG environment variable as follows:

```
NLS_LANG=Simplified Chinese_China.UTF8
```

For Informix data sources, the Informix wrapper sets Informix environment variables as follows:

```
CLIENT_LOCALE=zh_cn.UTF8  
GL_USEGLU=1
```

You must add the following setting to your sqllib/cfg/db2dj.ini file so that the Informix client correctly translates the GB 18030 data to Unicode:

```
DB_LOCALE=zh_cn.GB18030-2000
```

If you want Informix to verify that the DB_LOCALE setting matches the actual locale of the Informix database, you must add the following setting to your sqllib/cfg/db2dj.ini file:

```
DBNLS=1
```

Materialized query table issues

These issues apply only to DB2 Information Integrator configurations that use the DB2 Universal Database, Version 8, Enterprise Server Edition component.

Limitation

Query optimization is supported only for aggregate materialized query tables.

Creating materialized query tables that reference nicknames

A materialized query table that references nicknames cannot be created if the CREATE TABLE statement is issued from a non-catalog node. If you attempt to create such a materialized query table, error SQL020208 with reason code 3 is issued.

Refreshing materialized query tables that reference nicknames

If you try to refresh a materialized query table that references a nickname, the REFRESH TABLE command returns error SQL0668 with reason code 6. To avoid this problem, define a materialized query table with the MAINTAINED BY USER option. Then, issue an INSERT statement with a subquery to populate the materialized query table with the proper query results.

Workaround example:

To create a materialized query table called ast2 that references the nickname loopback_t2:

1. Define the ast2 materialized query table by using the following statement:

```
create table ast2 as (select c2, count(*) as cnt from loopback_t2 where c2 > 10 group by c2)
data initially deferred refresh deferred enable query optimization MAINTAINED BY USER;
```

2. Take the ast2 materialized query table out of the check pending state by using the following statement:

```
set integrity for ast2 all immediate unchecked;
```

3. Populate the ast2 materialized query table with valid query data by using the following statement:

```
insert into ast2 (select c2, count(*) from loopback_t2 where c2 > 10 group by c2 );
```

A workaround does not exist for materialized query tables that were defined with the REPLICATED keyword (replicated materialized query tables).

LIKE predicate issues

Examples of issues with the LIKE predicate are described in the following topics. A potential workaround for these problems is to set the PUSHDOWN server option to 'N'. Setting the PUSHDOWN server option to 'N' forces the LIKE portion of an SQL statement to be processed at the federated server. This portion of the SQL statement is not pushed down to the data source. Setting this server option might result in decreased performance.

LIKE predicate applied to a CHAR column

For Informix, Microsoft SQL Server, Oracle, and Sybase data sources, the LIKE predicate applied to a CHAR column is not pushed down to the data source. These data sources use different blank padding rules than DB2. For example, if a CHAR(10) column contains 'a', then the predicate char_col LIKE 'a' will be false in DB2 but true in the other data sources.

Additionally, for Microsoft SQL Server data sources, the LIKE predicate performs case-insensitive string comparisons, which cannot be pushed down to the data source.

You can improve performance by creating a function mapping for the LIKE(CHAR,...) function so that the LIKE predicate is pushed down to the data source. For example:

```
CREATE FUNCTION MAPPING FOR
  SYSIBM.LIKE(SYSIBM.CHAR,SYSIBM.VARCHAR(1)) SERVER infx_server;
```

However, using a function mapping can cause LIKE predicates to return different results than DB2 returns.

CHAR columns that contain DBCS data

Predicates that compare CHAR columns that contain DBCS data might return incorrect results when the federated server and the data source use different blank padding characters. You might be able to avoid the incorrect results by changing the local column type in the federated database system catalog from CHAR to VARCHAR.

BioRS wrapper support

The BioRS wrapper provides access to the BioRS query and retrieval system developed by Biomax.

Complete information about configuring the BioRS wrapper is available in the *DB2 Information Integrator Addendum to the Data Source Configuration Guide: BioRS Wrapper and Life Sciences User-Defined Functions*. You can obtain this document on the Web at www.ibm.com/software/data/integration/db2ii/.

Entrez wrapper support

The Entrez wrapper provides access to the PubMed and Nucleotide data sources in networks that use firewalls without proxies. If a proxy exists, the Entrez wrapper might not be able to access PubMed and Nucleotide data sources.

Microsoft SQL Server wrapper support

MSSQL wrapper restrictions on COMMIT statements as CURSOR_CLOSE_ON_COMMIT set YES

The MSSQL wrapper does not support DECLARE CURSOR WITH HOLD semantics. When a cursor is declared on Microsoft SQL Server, any COMMIT statements that are issued while the cursor is opened will result in the cursor being closed on the server. The next fetch or close cursor statement will also produce errors. The following statement returns the invalid cursor state:

```
Declare cursor (with hold semantics not allowed in SQL Server)
Fetch rows (SQLExtendedFetch)
Intermittent commits issued by user application causes SQL Server to closes cursors
Fetch continues successfully until the end of the last block fetch
Close Cursor returns SQLState: 24000 [Microsoft][ODBC SQL Server Driver]Invalid cursor state
```

Microsoft SQL Server insert, update, and delete restriction on tables with triggers

When you issue an insert, update, or delete command on a Microsoft SQL Server table through a nickname, you will receive one of the following error messages:

- SQLState: 24000 [Microsoft][ODBC SQL Server Driver]Invalid cursor state
- SQLState: HY000 [Microsoft][ODBC SQL Server Driver]Connection is busy with results for another hstmt

These error messages are a result of the wrapper's manual commit mode trying to reuse the connections or statement handles that have pending Microsoft SQL Server operations related to insert, update, and delete triggers defined on the Microsoft SQL Server tables. No workaround currently exists for inserting, updating, and deleting through nicknames on tables with triggers that are defined at the source.

Unicode restriction

The Microsoft SQL Server wrappers do not support Unicode. You cannot use the UTF-8 code page on a federated database with the Microsoft SQL Server wrappers.

Code page conversion requirements

On Windows federated servers, the DB2 federated database code page must match the default code page of the current operating system locale when you use the Microsoft SQL Server wrapper. The code pages must match because the ODBC driver for Microsoft SQL Server performs character set conversion based on the current operating system locale rather than on the DB2 code page.

On Linux and UNIX federated servers, the DB2 federated database code page must match the AppCodePage setting in the `odbc.ini` configuration file when you use the Microsoft SQL Server wrapper. The code pages must match because the DataDirect Technologies Connect ODBC driver for Microsoft SQL Server performs character set conversion based on the AppCodePage setting rather than on the DB2 federated database code page.

You set the DB2 code page when you create the federated database.

Example of code page conversion configuration (UNIX and Linux):

Suppose that you use the DataDirect Technologies Connect ODBC driver to access Microsoft SQL Server data sources. You want to use a Japanese code page on the federated server. In the `db2dj.ini` file, you must add the following settings:

```
LC_MESSAGES=Ja_JP
LANG=Ja_JP
LC_ALL=Ja_JP
```

In the reference information provided with the DataDirect Technologies Connect ODBC driver, you determine that the proper code page is:

11 = Microsoft CP 932 DBCS

In the `odbc.ini` file, you must set the `AppCodePage` to 11. An example of the `odbc.ini` file settings is:

```
[japan2000]
Driver=/opt/odbc/lib/ivmsss18.so
Description=MS SQL Server 2000
Database=jtest
Address=9.xx.xxx.xxx,1433
AppCodePage=11
AutoTranslate=yes
```

If an `AppCodePage` value is not set in the `odbc.ini` file, the ODBC driver uses the default English code page.

ODBC wrapper support

Using the **PUSHDOWN** server option with ODBC data sources

If queries fail due to remote source syntax errors, set the **PUSHDOWN** server option value to 'N', or remove the **PUSHDOWN** server option entirely.

ODBC wrapper **DB2_ONE_REQUEST_PER_CONNECTION** restriction

For ODBC drivers that allow only one active statement on a connection, you must set the **DB2_ONE_REQUEST_PER_CONNECTION** server option to 'Y' for the data source. When you set the **DB2_ONE_REQUEST_PER_CONNECTION** server option, you must also set the **DB2_IUD_ENABLE** server option to 'N'.

To set these server option, use the **ALTER SERVER** statement:

```
ALTER SERVER odbserv OPTIONS
  (ADD DB2_ONE_REQUEST_PER_CONNECTION 'Y', ADD DB2_IUD_ENABLE 'N')
```

The ODBC wrapper does not allow **INSERT**, **UPDATE**, or **DELETE** statements when the **DB2_ONE_REQUEST_PER_CONNECTION** option is set to 'Y'.

To determine if your ODBC driver allows only one active statement on a connection, use the **SQLGetInfo** function of the ODBC driver and check the returned value for the **SQL_ACTIVE_STATEMENTS** or **SQL_MAX_CONCURRENT_ACTIVITIES** InfoTypes. If the returned value is 1, then your ODBC driver allows only one active statement on a connection.

Unicode restriction

The ODBC wrapper does not support Unicode. You cannot use the UTF-8 code page on a federated database with the ODBC wrapper.

Code page conversion requirements

Code page conversions are performed by the ODBC driver or ODBC Driver Manager. The ODBC wrapper does not perform any coded character set conversions.

Using the ODBC wrapper instead of an existing data source wrapper

The following problems occur when you use the ODBC wrapper that is to access a data source instead of the wrapper designed to access that specific data source:

DB2 for Linux, UNIX, and Windows data sources

When you use the ODBC wrapper to access DB2 for Linux, UNIX, and Windows data sources, the DB2 federated database processes end abnormally. Use the DRDA[®] wrapper to access DB2 for Linux, UNIX, and Windows data sources.

Informix data sources

Do not attempt to use the ODBC wrapper to access Informix data sources. You cannot use the ODBC wrapper to create nicknames for Informix data source objects. You cannot use the ODBC wrapper to create Informix tables by using a pass-through session or by using transparent DDL. To access Informix data sources, use the Informix wrapper.

ODBC data sources with indexes

When you create a nickname on a remote table that contains indexes, the ODBC wrapper does not record the index information in the federated database system catalog. You must create index specifications for the table by using the CREATE INDEX statement with the SPECIFICATION ONLY clause.

ODBC wrapper restriction

Positioned UPDATE and DELETE statements and certain searched UPDATE and DELETE statements on a nickname fail if a unique index on non-nullable, non-character (fixed length or varying length) columns does not exist on the nickname or on a corresponding remote table for the nickname. The failure error is SQL30090 reason code 21.

Oracle wrapper support**djxlinkOracle script errors**

If you use the djxlinkOracle script on AIX to link the SQLNET wrapper and the AIX Base Application Development Math Library is not installed, then the script will fail with linkage editor errors.

You can determine if the library is installed by issuing the following AIX command:

```
lsllpp -l bos.adt.libm
```

To avoid linkage editor errors, either install the AIX Base Application Development Math Library or edit the `djxlinkOracle` script and remove all occurrences of option `-lm` from the `ld` (linkage editor) commands.

NET8 64-bit wrapper support

The 64-bit Oracle NET8 wrapper on UNIX federated servers uses the Oracle 9i client library `libclntsh.<suffix>`, where `<suffix>` is a suffix that is determined by the operating system. This library is in the `$ORACLE_HOME/lib` directory. To ensure that this library is installed, you might need to install the Oracle 9i client using a server installation. You can then use the custom option to remove any server-specific options.

Sybase wrapper support

Required version of Sybase Adaptive Server Enterprise

If you are using the Sybase wrappers with Sybase Adaptive Server Enterprise 11.9, you must use version 11.9.2.6 or later. If you do not have one of these versions installed, you must install the latest Emergency Bug Fix (EBF) from Sybase on your Sybase server.

Calculations on SMALLINT columns

An SQL statement that includes a calculation on a `SMALLINT` column might cause an arithmetic overflow error. You can avoid this error if you explicitly define the column as an `INTEGER` data type.

This problem occurs with both the `DBLIB` and `CTLIB` wrappers.

Changing VARCHAR columns to CLOB or BLOB columns

If you change the local column type from `VARCHAR` to a `CLOB` or `BLOB` data type, the `CTLIB` wrapper cannot perform `SELECT` statements on the altered column.

Changing to BIGINT columns

On Windows federated servers, if you change the local column type to `BIGINT`, incorrect results are returned for that column.

This problem occurs only with the `CTLIB` wrapper. The `DBLIB` wrapper is not affected by this problem.

Column names must be unique

The `DBLIB` wrapper cannot create a nickname for a table that contains duplicate column names. For example, the `DBLIB` wrapper considers the column names `abc`, `ABC`, `Abc` to be identical.

To create a nickname for a Sybase table that contains the same column names using different cases, you must either change the remote column names to unique names or use the CTLIB wrapper instead of the DBLIB wrapper.

DBCS for federated server name

The DBLIB wrapper cannot use DBCS for a federated server name. To use DBCS for a federated server name, you must use the CTLIB wrapper.

Using SELECT statements on LOB columns (DBLIB wrapper only)

To select a LOB column, the data source table must have a unique index and a TIMESTAMP column.

This requirement does not apply to the CTLIB wrapper.

Changing DECIMAL or NUMERIC columns to INTEGER columns

If you change the local column type from DECIMAL or NUMERIC to INTEGER, the DBLIB wrapper cannot process a SELECT statement on that column. The select operation fails with an SQL0303N error.

Extended Search wrapper restrictions

Incorrect date format in Extended Search queries

When you use an SQL expression of the Extended Search custom function ESWRAPPER.ES_SEARCH, you must use the YYYYMMDD date format. For example, for December 31, 2002, the date is 20021231.

The following example shows a sample of a query that uses an incorrect date format. The query is then rewritten using the correct date format:

Incorrect format

```
SELECT OWNER, DOC_CONTENT
FROM MedDocs
WHERE ESWRAPPER.ES_Search(DOC_RANK, '( ( TOKEN:EXACT "MARS") AND
( "TITLE" IN "Artificial Liver") AND
("DATE" >= "01/01/2001") ) )' ) = 1
```

Correct format

```
SELECT OWNER, DOC_CONTENT
FROM MedDocs
WHERE ESWRAPPER.ES_Search(DOC_RANK, '( ( TOKEN:EXACT "MARS") AND
( "TITLE" IN "Artificial Liver") AND
("DATE" >= "20010101") ) )' ) = 1
```

SQL0901N error returned for an SQL query that uses an Extended Search nickname

The following error message is returned when an SQL query that uses an Extended Search nickname consists only of predicates in the WHERE clause that cannot be handled by the Extended Search wrapper.

SQL0901N The SQL statement failed because of a non-severe system error. Subsequent SQL statements can be processed. (Reason "sqlno_crule_remote_pushdow[200]:rc(-2144272270) Error generatin".) SQLSTATE=58004

If the SQL query has at least one predicate that can be handled by the Extended Search wrapper, the query is accepted and run.

The following examples show queries that return the SQL0901N error:

A query without any predicate:

```
SELECT E.COLUMN  
FROM   ES_NICKNAME as E
```

SQL queries that use ANY, ALL, SOME, EXIST predicates: Such predicates require an additional predicate that can be handled by the Extended Search wrapper.

```
SELECT E.COLUMN  
FROM   ES_NICKNAME as E  
WHERE  E.COLUMN = ALL (SELECT COL FROM TABLE)
```

The NOT IN predicate is not supported using FULL SELECT:

```
SELECT E.COLUMN  
FROM   ES_NICKNAME as E  
WHERE  E.COLUMN NOT IN (SELECT COL FROM TABLE)
```

The IN predicate is not supported using FULL SELECT:

```
SELECT E.COLUMN  
FROM   ES_NICKNAME as E  
WHERE  E.COLUMN IN (SELECT COL FROM TABLE)
```

A fixed column predicate:

```
SELECT E.COLUMN  
FROM   ES_NICKNAME as E  
WHERE  DOC_ID = 'ABC'
```

The following example shows a query that does not return the SQL0901N error:

```
SELECT E.COLUMN  
FROM   ES_NICKNAME as E  
WHERE  E.COLUMN IN (SELECT COL FROM TABLE) AND  
       E.COLUMN = 'ABC'
```

In this query, the predicate E.COLUMN IN (SELECT COL FROM TABLE) is rejected by the Extended Search wrapper, but the predicate E.COLUMN = 'ABC' can be handled by the Extended Search wrapper.

FIELD_DATATYPE SMALLINT supported data types

The data type information in the table field, FIELD_DATATYPE SMALLINT, is incorrect. The correct integer values that represent the actual data type of the field value are DATE, VARCHAR, DOUBLE, and INTEGER.

Unsupported data types

The DECIMAL (integer) data type is not supported by the Extended Search wrapper.

Teradata wrapper support

DB2 Universal Database for AIX only: GRAPHIC and VARGRAPHIC Teradata data types are not allowed for EUC-JP databases

Teradata uses the EUC-JP encoding in GRAPHIC and VARGRAPHIC data types. If the federated database uses the EUC-JP code page, you cannot create a nickname on a Teradata table that contains GRAPHIC or VARGRAPHIC columns. You also cannot use transparent DDL to create a Teradata table that contains GRAPHIC or VARGRAPHIC columns. DB2 Universal Database uses UCS-2 encoding because some EUC-JP characters are 3 bytes.

If you try to create a nickname on a Teradata table with GRAPHIC or VARGRAPHIC columns or use transparent DDL to create a Teradata table with GRAPHIC or VARGRAPHIC columns, you will get error message SQL3324.

GRAPHIC and VARGRAPHIC Teradata data types are not allowed for UTF-8 databases

Teradata does not support GRAPHIC and VARGRAPHIC data types for the UTF-8 character set. If the federated database uses the UTF-8 code page, you cannot create a nickname on a Teradata table that contains GRAPHIC or VARGRAPHIC columns. You cannot use transparent DDL to create a Teradata table that contains GRAPHIC or VARGRAPHIC columns. See the Teradata documentation for more information.

Accessing Microsoft Excel data sources

You can access Microsoft Excel spreadsheets by using either the Excel wrapper or the ODBC wrapper. Additional details on the advantages and disadvantages of each access method, the configuration requirements, and the limitations of each method is available on the Web at www.ibm.com/software/data/integration/db2ii/support.

Using the CREATE TYPE MAPPING statement

If you specify the REMOTE keyword in the CREATE TYPE MAPPING statement, incorrect server information is stored in the federated database catalog table. For example, issuing the following DDL statement results in ODBC REMOTE stored as the server type in the catalog:

```
CREATE TYPE MAPPING TM1 FROM LOCAL TYPE SYSIBM.INTEGER
  TO SERVER TYPE ODBC REMOTE TYPE SQL_INTEGER
```

To ensure that the correct server information is stored in the catalog, do not specify the REMOTE keyword in the CREATE TYPE MAPPING statement. For example:

```
CREATE TYPE MAPPING TM1 FROM LOCAL TYPE SYSIBM.INTEGER
  TO SERVER TYPE ODBC TYPE SQL_INTEGER
```

DB2LOOK command limitations generating federated DDL statements

The DB2LOOK command generates DDL statements from the metadata in the DB2 catalogs. After the DB2LOOK command detects a CREATE TYPE MAPPING for LOB data, DB2LOOK does not generate any subsequent CREATE TYPE MAPPING statements. You must manually create the data type mapping statements that DB2LOOK omits.

The getstats and get_stats_nr tools

Because the **runstats** utility is not supported for nicknames in this release, you can download the **getstats** or **get_stats_nr** tools, which are available from the DB2 Information Integrator Web site at www.ibm.com/software/data/integration/db2ii. Both of these tools gather key statistical information that is used by the DB2 query optimizer. The information that is gathered includes table cardinality, column cardinality, and index firstkeycard and fullkeycard. The **getstats** and **get_stats_nr** tools run queries on the remote data source to gather the statistics. These queries might use a large amount of resources at the remote data source.

The **getstats** and **get_stats_nr** tools are temporary solutions that allow you to update nickname statistics without dropping and recreating the nickname. Generally, use the **getstats** tool for relational data sources. Use the **get_stats_nr** tool for nonrelational data sources.

getstats

You can use the **getstats** tool with DRDA, Oracle, Sybase, IBM[®] Informix, Microsoft SQL Server, Teradata, ODBC, table-structured file, Excel, XML and BioRS data sources. The **getstats** tool has limitations for some data sources.

XML The **getstats** tool can collect the statistics of the root nickname only when that nickname was not created with the DOCUMENT option set to FILE.

Table-structured files

The **getstats** tool can run only if the nickname was not created with the DOCUMENT option.

ODBC

The **getstats** tool can run only if the data sources that are accessed by the ODBC wrapper support functions such as count(*), count(distinct), min(), and max.

The **getstats** tool runs on Windows, AIX, Solaris Operating Environment, Linux, and HP-UX systems.

get_stats_nr

You can use the **get_stats_nr** with flat file, Excel, BioRS, Documentum, BLAST, HMMR, Entrez, XML, and Extended Search data sources. Like the **getstats** tool, the **get_stats_nr** tool gathers basic statistical information for data sources. The **get_stats_nr** tool also gathers cost parameters for the remote nonrelational data sources, and updates the DB2 catalog.

The **get_stats_nr** tool runs only on Windows and AIX.

The **getstats** and **get_stats_nr** sample tools are provided with no warranty, formal or implied. The **getstats** and **get_stats_nr** tools are provided for demonstration purposes with no guarantees, obligations, or commitments from the IBM company.

You can download the **getstats** and **get_stats_nr** tools from the following Web site: www.ibm.com/software/data/integration/db2ii/support.html.

CREATE TABLE restrictions

If you specify an invalid option on the CREATE TABLE statement, DB2 ignores the option and does not issue an error message. For example, in the following statement, DB2 ignores the option, `bad_option`:

```
CREATE TABLE my_table(c1 INT)
  OPTIONS(remote_server 'MY_SERVER', remote_schema 'J15USER1', bad_option 'for fun');
```

The following options are valid CREATE TABLE options:

- REMOTE_SERVER
- REMOTE_SCHEMA
- REMOTE_TABNAME
- SQL_SUFFIX (DRDA wrapper only)

Error in left outer join queries to BLAST, Entrez, Extended Search, HMMER, and XML data sources

An error might be returned on some left outer join queries that include nicknames for BLAST, Entrez, Extended Search, HMMER, and XML data sources. These errors occur when a nickname for one of these nonrelational data sources is on the inner portion of the left outer join. For example:

```
SELECT * FROM db2_table LEFT OUTER JOIN non_relational_nickname
ON db2_table.col = non_relational_nickname.col
WHERE <optional_predicates>
db2_table can be a local table or a nickname.
```

The error message returned is -901 "Input plan empty: cannot build lolepop."

At this time, there is no workaround for this problem.

DB2 Control Center updates

Differences between federated online help and DB2 Control Center object names

The window names that are documented in the online help are not current for some of the DB2 Control Center windows.

Table 5. Changed window names

Window name in online help	Actual window name
Wrapper properties	Properties
Discover Server	Discover
Server Discovery	Discover
Discovery	Discover
Multi-value Select	Values

The control names that are documented in the online help are not current for some of the controls on windows in the DB2 Control Center.

Table 6. Changed control names

Window	Control name	Actual control name
Discover (XML data sources)	Specify data source document type	Specify data source

db2updv8 command

If you are using a federated database with DB2 Universal Database Version 8.1.2 that was created using DB2 Universal Database Version 8.1 or DB2 Universal Database Version 8.1 Fix Pack 1, you must update your federated database using the db2updv8 command.

Syntax:

```
▶▶—db2updv8—-d—database-name—┐—————┘▶▶  
└-u—userid—-p—password—┘
```

An example of this command is:

```
db2updv8 -d testdb -u dbadmin -p dbpasswd
```

If you do not update your database to DB2 Universal Database Version 8.1.2, you will get error message SQL0444N when you try to create a wrapper from the Federated Database Objects folder in the Control Center.

Federated client and server compatibility

If you use DB2 Universal Database Version 7 clients with a DB2 Universal Database Version 8.1.2 server, you must upgrade your clients to DB2 Universal Database Version 8.1.2 to use the new federated functionality in the DB2 Control Center.

If you use DB2 Universal Database Version 8.1.2 clients and a DB2 Universal Database Version 7 server, you must upgrade your server to DB2 Universal Database Version 8.1.2 to use the new federated functionality in the DB2 Control Center.

If you use DB2 Universal Database Version 8.1 or DB2 Universal Database Version 8.1 Fix Pack 1 clients with a DB2 Universal Database Version 8.1.2 server, only some federated functionality is available in the DB2 Control Center. To use all of the federated functionality in the DB2 Control Center, you must upgrade your clients to DB2 Universal Database Version 8.1.2.

MQ user-defined function message string updates

In the following messages, '2pc' should not be part of the message. The value '2pc', for two-phase commit, is not available in this release. The valid values are '0pc', '1pc', and 'all'. The messages they should read:

```
=====
Usage: disable_MQFunctions -n dbName -u uID -p password
                        [-v 0pc | 1pc | all]
=====";
=====
Usage: enable_MQFunctions -n dbName -u uID -p password
                        [-q queuemanager] [-v 0pc|1pc|all]
                        [-novalidate] [-interactive] [-force]
=====";
Only a value of 'all', '0pc', or '1pc' is allowed
for the -v option.";
```

Corrections to the documentation

The following information describes corrections to previously released HTML, PDF, and printed documentation. Updated versions of the affected topics will be available in a future version of the DB2 Information Integrator documentation.

All of the navigational references to the DB2 Information Center in this section assume that you are using the DB2 Information Integrator navigation tree as the starting point. You choose the navigation tree that you want to use from the upper left corner of the DB2 Information Center when it is displayed in a browser.

DB2 Information Integrator Installation Guide

Topic title: Supported operating systems for DB2 Information Integrator

Location in the DB2 Information Center

Installing DB2 Information Integrator → Planning to install DB2 Information Integrator Version 8. This topic was last updated for DB2 Information Integrator Version 8.1.

Location in PDF and printed book

Book title: *DB2 Information Integrator Installation Guide*

Chapter title: Planning to install DB2 Information Integrator

Correction

The supported operating systems shown in the following table supersede the list of supported operating systems shown in this topic:

Table 7. Supported operating systems for DB2 Information Integrator

Operating system	32-bit support	64-bit support
Windows NT	✓	
Windows 2000	✓	
Windows XP	✓	
Solaris Operating Environment 7, 8, 9	✓	✓
AIX 4.3	✓	
AIX 5	✓	✓

Table 7. Supported operating systems for DB2 Information Integrator (continued)

HP-UX, Version 11i	✓	✓
RedHat Linux, Version 7. See the DB2 for Linux Web page for the supported level of the Linux kernel: http://www.ibm.com/software/data/db2/linux/validate	✓	

DB2 Enterprise Server Edition and DB2 Connect Enterprise Edition are supported for development and testing on Windows XP Professional Edition, Windows 2000 Professional Edition, and Windows NT Workstation.

Topic title: Checking the data source environment variables

Location in the DB2 Information Center

Installing DB2 Information Integrator -> Installing DB2 Information Integrator and setting up a federated database and server -> Post-installation tasks-> Checking the data source environment variables, section "Manually setting the Informix environment variables, " subsection "Setting up Informix code page conversion." This topic was last updated for DB2 Information Integrator Version 8.1.

Location in PDF and printed book

Book title: *DB2 Information Integrator Installation Guide*

Chapter title: Installing DB2 Information Integrator and setting up a federated server and database

Correction

This topic currently contains the following text:

Each time you issue a SQL statement that accesses data from an Informix data source, the Informix wrapper determines which code page value to use for that connection. If the Informix environment variable CLIENT_LOCALE is set in the db2dj.ini file on the federated server, then the wrapper uses the value in the db2dj.ini file.

The text should be replaced with:

Each time that the Informix wrapper connects to an Informix data source, the wrapper determines which code page value to use for that connection.

Topic title: Checking the data source environment variables

Location in the DB2 Information Center

Installing DB2 Information Integrator → Installing DB2 Information Integrator and setting up a federated database and server → Post-installation tasks. This topic was last updated for DB2 Information Integrator Version 8.1.

Location in PDF and printed book

Book title: *DB2 Information Integrator Installation Guide*

Chapter title: Installing DB2 Information Integrator and setting up a federated server and database

Correction

In addition to the data source environment variables that are listed, the following optional environment variables are valid for Oracle data sources:

- NLS_LANG
- ORACLE_BASE

Topic title: Checking the data source environment variables

Location in the DB2 Information Center

Installing DB2 Information Integrator → Installing DB2 Information Integrator and setting up a federated database and server → Post-installation tasks. This topic was last updated for DB2 Information Integrator Version 8.1.

Location in PDF and printed book

Book title: *DB2 Information Integrator Installation Guide*

Chapter title: Installing DB2 Information Integrator and setting up a federated server and database

Correction

The following restriction on the `db2dj.ini` file has been removed:

The environment variable value cannot contain white space.

The following restriction is incorrect:

Every line must end with an ASCII carriage return (0x0D) or line feed (0x0A) character.

The restriction should read:

The maximum length of any line in the file is 1021 bytes. Data beyond that length will be ignored.

Topic title: Hardware and software requirements for a federated system

Location in the DB2 Information Center

Installing DB2 Information Integrator → Planning to install

DB2 Information Integrator Version 8. This topic was last updated for DB2 Universal Database Version 8.1.2.

Location in PDF and printed book

Book title: *DB2 Information Integrator Installation Guide*

Chapter title: Planning to install DB2 Information Integrator

Correction

BioRS 5.0.14 is supported for use with the BioRS wrapper.

DB2 Information Integrator Migration Guide

Topic title: Migration roadmap

Location in the DB2 Information Center

Migrating to DB2 Information Integrator -> Planning to migrate to DB2 Information Integrator Version 8. This topic was last updated for DB2 Information Integrator Version 8.1.

Location in PDF and printed book

Book title: *DB2 Information Integrator Migration Guide*

Chapter title: DB2 Information Integrator migration overview

Correction

Add the following text to the beginning of this topic:

Migration converts older versions of instances and databases to DB2 Information Integrator. You can migrate from DataJoiner, Version 2.1.1 or from DB2 Universal Database for UNIX, Windows, and OS/2 Version 7. You can migrate your system to DB2 Information Integrator by issuing migration commands, such as **db2ckmig**, **db2imigr**, and **db2 migrate database** and by installing DB2 Information Integrator and DB2 for Linux, UNIX, and Windows, Version 8. You do not need to copy files or create links, unless you are explicitly instructed to perform these tasks.

Table 1 of this topic is not complete and now includes the following changes:

- In Step 1, the actions also include understanding the restrictions that are associated with migrating from DataJoiner and DB2.
- In Step 3, the correct action is to install the proper software. The sequence that you follow to install the software depends on the version of the product that you are migrating from and the operating system that you use.

- In Step 4, the correct action is to migrate your instances and databases. To migrate instances and databases, issue the **db2imigr** command (on UNIX operating systems only) and the **db2 migrate database** command. To migrate the databases to new instances, issue the **db2 restore db** command.

Topic title: Preparing to migrate to DB2 Information Integrator

Location in the DB2 Information Center

Migrating to DB2 Information Integrator -> Planning to migrate to DB2 Information Integrator Version 8 -> Before you migrate. This topic was last updated for DB2 Information Integrator Version 8.1.

Location in PDF and printed book

Book title: *DB2 Information Integrator Migration Guide*

Chapter title: Before you migrate

Correction

Add a new section titled "Copying the Sybase interfaces file" with the following text:

Before you migrate to DB2 Information Integrator, make a copy of the interfaces file. This file is not automatically migrated to DB2 Information Integrator, and you must restore it after you complete the migration steps. The interfaces file is typically located in the \$HOME/sqllib directory.

Topic title: Installation sequence for migrating instances and databases

Location in the DB2 Information Center

Migrating to DB2 Information Integrator -> Migrating Windows servers. This topic was last updated for DB2 Information Integrator Version 8.1.

Location in PDF and printed book

Book title: *DB2 Information Integrator Migration Guide*

Chapter title: Migrating to DB2 Information Integrator

Correction

Add the following text to the beginning of the "Installation sequence for migrating instances and databases on UNIX operating systems" section:

You must complete all the pre-migration tasks except verify the database before you install the software. After you install DB2 Enterprise Server Edition and DB2 Information

Integrator, follow the steps in the topic "Verifying that your databases are ready for migration" before you proceed with the migration steps.

Topic title: Space considerations for DB2 migration

Location in the DB2 Information Center

Migrating to DB2 Information Integrator -> Planning to migrate to DB2 Information Integrator Version 8. This topic was last updated for DB2 Information Integrator Version 8.1.

Location in PDF and printed book

Book title: *DB2 Information Integrator Migration Guide*

Chapter title: Migration considerations

Correction

This topic discusses the use of the `logfilsiz`, `logprimary`, and `logsecond` configuration parameters and instructs you to double the values of these parameters. However, when you migrate from DataJoiner, the log file space that you need depends on the size of your database. The total size for all log files must be twice the backup file size of the database.

Use the following formula to determine if you have sufficient log file space on your DataJoiner system:

$$(4096 \times \text{logfilsiz} \times (\text{logprimary} + \text{logsecond})) > 2 \times \text{(DataJoiner database backup file size)}$$

4096 is the page size of the log file. `logfilsiz`, `logprimary`, and `logsecond` are database configuration parameters.

For example, suppose that the backup file size for the DataJoiner database is 117.5 MB. Before you migrate to DB2, the total size for all log files needs to be $2 \times 117.5 = 235$ MB. The `logfilsiz` parameter is 5000. The DataJoiner database has 6 `logprimary` files and 4 `logsecond` files. Using these numbers, the total size for all log files is:

$$(4096 \times 5000 \times (6 + 4)) = 204.8 \text{ MB}$$

Because the total file size for all log files is less than twice the backup file size of the DataJoiner database, the migration fails. You can avoid this error by increasing the `logfilsiz` parameter or by creating additional log files before you migrate.

Topic title: Migrating databases

Location in the DB2 Information Center

Migrating to DB2 Information Integrator -> Migrating Windows servers. This topic was last updated for DB2 Information Integrator Version 8.1.

Migrating to DB2 Information Integrator → Migrating UNIX servers. This topic was last updated for DB2 Information Integrator Version 8.1.

Location in PDF and printed book

Book title: *DB2 Information Integrator Migration Guide*

Chapter title: Migrating to DB2 Information Integrator

Correction

When you migrate DataJoiner databases and the log file size is too small, the following error appears:

SQL1704N Database migration failed. Reason code "3".

You must increase the log file size and then issue the **db2 migrate database** command again.

The total size for all log files must be twice the DataJoiner backup file size of the database. Use the following formula to determine if you have sufficient log file space:

$$(4096 \times \text{logfilsiz} \times (\text{logprimary} + \text{logsecond})) > 2 \times (\text{DataJoiner database backup file size})$$

4096 is the page size of the log file. logfilsiz, logprimary, and logsecond are database configuration parameters. After the migration is complete, reset the logfilsiz, logprimary, and logsecond parameters.

Topic title: Verifying that your databases are ready for migration

Location in the DB2 Information Center

Migrating to DB2 Information Integrator → Before you migrate. This topic was last updated for DB2 Information Integrator Version 8.1.

Location in PDF and printed book

Book title: *DB2 Information Integrator Migration Guide*

Chapter title: Before you migrate

Correction

The **db2ckmig** command also checks that abstract data types do not exist in a database and that objects such as tables and view do not use either the SYSCAT or SYSSTAT schema in a database.

Add the following steps to the Procedure section:

1. Log on to the instance that you are migrating from.

- If you are migrating from a previous version of DB2 Universal Database, log on as the DB2 instance owner.
 - If you are migrating from DataJoiner, log on as the DataJoiner instance owner.
2. Issue the **db2stop** command to stop the instance.
 3. On Windows servers, insert the DB2 Universal Database Version 8 product CD in the CD-ROM.
 4. From a DB2 command line prompt, change to the appropriate directory:
 - On AIX, the path for this command is `/usr/opt/db2_08_01/bin/db2ckmig`.
 - On Linux, HP-UX, and the Solaris Operating Environment, the path for this command is `/opt/IBM/db2/V8.1/bin/db2ckmig`.
 - On Windows operating systems, change to the `\db2\Windows\utilities` directory on the CD.
 5. Issue the **db2ckmig** command with the `/1` parameter to verify that the databases that are owned by the current instance are ready to be migrated and to generate a log file.

If the **db2ckmig** command does not find errors, you receive the following message:

```
db2ckmig was successful. Database(s) can be migrated.
```
 6. If the **db2ckmig** command finds errors, correct the problems before you migrate.
 7. Check the log file to ensure that it contains the following text:


```
Version of DB2CKMIG being run: VERSION 8
```

Topic title: Migrating DB2 (Windows)

Location in the DB2 Information Center

Migrating to DB2 Information Integrator -> Migrating Windows servers. This topic was last updated for DB2 Information Integrator Version 8.1.

Location in PDF and printed book

Book title: *DB2 Information Integrator Migration Guide*

Chapter title: Migrating to DB2 Information Integrator

Correction

Add the following text after step 7:

On Windows operating systems, all existing instances are migrated during the server installation.

Topic title: Migrating DB2 (UNIX)

Location in the DB2 Information Center

Migrating to DB2 Information Integrator -> Migrating UNIX servers. This topic was last updated for DB2 Information Integrator Version 8.1.

Location in PDF and printed book

Book title: *DB2 Information Integrator Migration Guide*

Chapter title: Migrating to DB2 Information Integrator

Correction

This topic explains how to use the **db2imigr** command to migrate an instance. If you are migrating from DataJoiner and need to modify the configuration settings, change the ownership of the db2dj.ini file from root to the owner of the instance after you issue the **db2imigr** command. The db2dj.ini file is in the sqllib/cfg directory.

Topic title: Migrating to DB2 Version 8 on UNIX 64-bit from DataJoiner on UNIX 32-bit

Location in the DB2 Information Center

Migrating to DB2 Information Integrator -> Migrating from DataJoiner on UNIX 32-bit to DB2 Version 8 on UNIX 64-bit. This topic was last updated for DB2 Information Integrator Version 8.1.

Location in PDF and printed book

Book title: *DB2 Information Integrator Migration Guide*

Chapter title: Migrating to DB2 Information Integrator

Correction

The steps in the Procedure section are now:

1. Install DB2 Universal Database Version 8 and DB2 Information Integrator on your DataJoiner system. Do not remove the previous version.
2. Stop the DataJoiner Version 2.1.1 instance.
3. Run the **/usr/opt/db2_08_01/bin/db2ckmig** command as the instance owner on the DataJoiner 2.1.1 32-bit system to ensure that the database can be migrated.
4. Backup the DataJoiner Version 2.1.1 database.
5. Install DB2 Universal Database Version 8 and DB2 Information Integrator Version 8 on your 64-bit system.

6. Create the 32-bit instance by using either the installer or by using the **db2icrt -w 32 -u <instance_name>** command.
7. Migrate the database by using the **db2 restore database** command.
8. If you receive the following error, the database was restored, but the database was not migrated.

SQL2519N The database was restored but the restored database was not migrated to the current release. Error "-1704" with tokens "3" is returned. SQLSTATE=57011

You must complete the following steps to migrate the database:

- a. Increase the total size for all log files. When you migrate from DataJoiner, the total size for all log files must be twice the backup file size of the database. Use the following formula to determine if you have sufficient log file space on your DataJoiner system:

$$(4096 \times \text{logfilsiz} \times (\text{logprimary} + \text{logsecond})) > 2 \times (\text{DataJoiner database backup file size})$$
 4096 is the page size of the log file. logfilsiz, logprimary, and logsecond are database configuration parameters.
- b. Migrate the database by using the **db2 migrate database** command.
- c. If the log file size is still not large enough, the following error appears:

SQL1704N Database migration failed. Reason code "3".

You must increase the log file size, and issue the **db2 migrate database** command again.
- d. After the migration is complete, reset the logfilsiz, logprimary, and logsecond parameters.
9. Update the instance to a DB2 Universal Database Version 8 64-bit instance. To update the DataJoiner Version 2.1.1 32-bit instance to a DB2 Universal Database Version 8 64-bit instance, use the **db2iupdt** command, and specify the **-w** parameter with a value of **64**.
10. Restart the instance.

Topic title: After you migrate to DB2 Information Integrator

Location in the DB2 Information Center

Migrating to DB2 Information Integrator -> Accessing data sources after migration. This topic was last updated for DB2 Information Integrator Version 8.1.

Location in PDF and printed book

Book title: *DB2 Information Integrator Migration Guide*

Chapter title: After you migrate

Correction

Add a new section titled “Restoring DataJoiner configuration settings” with the following text:

Before you migrated to DB2 Information Integrator Version 8, you copied the DB2 DataJoiner configuration settings. After you migrate the instances and databases, you must restore the configuration settings.

Locate the `djenv.log` file that you created before migrating to DB2 Information Integrator. Compare the variables listed in the `djenv.log` file with the variables in the `$HOME/sql1lib/cfg/db2dj.ini` file. If necessary, edit the `db2dj.ini` file to add any missing variables.

Correction

Add a new section titled “Restoring DB2 Version 7 federated systems configuration settings” with the following text:

Before you migrated to DB2 Information Integrator Version 8, you copied the configuration settings in the `db2dj.ini` file. After you migrate the instances and the databases, you must restore the configuration settings.

Locate the `db2dj.ini` file that you copied before migrating to DB2 Information Integrator Version 8. Restore that copy of the `db2dj.ini` file to the `$HOME/sql1lib/cfg/` directory.

Correction

Add a new section titled “Dropping mappings for a specific server type” with the following text:

In DB2 DataJoiner, you can create data type mappings and function mappings that are associated with a particular type of server. For example, you can create a mapping that applies to all Oracle servers or to all Sybase servers.

After you migrate to DB2 Information Integrator Version 8, you might need to drop a wrapper and create the wrapper again. When you drop the wrapper, other migrated objects that are dependent upon the wrapper, such as the mappings, are also dropped. If you attempt to drop the mapping after you drop the wrapper, you receive an error. The error is:

DB21034E The command was processed as an SQL statement because it was not a valid Command Line Processor command. During SQL processing it returned: SQL0901N The SQL statement failed because of a non-severe system error. Subsequent SQL statements can be processed.
(Reason "missing type mapping from server pd".) SQLSTATE=58004

You can avoid this error by taking one of the following actions:

- Drop the mapping before you drop the wrapper.
- Drop the mapping at any time, if you specify the mapping name in the DROP statement with trailing spaces so that the total length of the name is equal to eighteen characters.

Topic title: Accessing DB2 family data sources after migration

Location in the DB2 Information Center

Migrating to DB2 Information Integrator -> Accessing data sources after migration. This topic was last updated for DB2 Information Integrator Version 8.1.

Location in PDF and printed book

Book title: *DB2 Information Integrator Migration Guide*

Chapter title: After you migrate

Correction

The Prerequisites section is missing. The prerequisites for this topic are:

Record the server names and the database alias names for the databases that you accessed before you migrated to DB2 Information Integrator. See the topic "Recording database information before migrating."

Step 3 in the procedure now states that you must catalog the node by issuing the **CATALOG TCPIP NODE** command if you migrated only the database.

Step 4 in the procedure now states that you must catalog the database, if you migrated instances and the database or if you migrated only the database. Issue the **CATALOG DATABASE** command to catalog and store remote database information in the federated system database directory. The name of the node that you specify must match the node name that you added in the **CATALOG TCPIP NODE** command. If the length of the remote database name is greater than eight characters, issue the **CATALOG DCS DATABASE** command. This command creates a DCS directory entry. The length of the database names for DB2 Universal Database for z/OS and DB2 Universal Database for iSeries data sources are typically greater than eight characters.

Topic title: Accessing Informix data sources after migration

Location in the DB2 Information Center

Migrating to DB2 Information Integrator -> Accessing data sources after migration. This topic was last updated for DB2 Information Integrator Version 8.1.

Location in PDF and printed book

Book title: *DB2 Information Integrator Migration Guide*

Chapter title: After you migrate

Correction

Step 4b in the procedure contains an incorrect command name. The correct command is **db2iupdt**.

Topic title: Accessing Microsoft SQL Server data sources after migration

Location in the DB2 Information Center

Migrating to DB2 Information Integrator -> Accessing data sources after migration. This topic was last updated for DB2 Information Integrator Version 8.1.

Location in PDF and printed book

Book title: *DB2 Information Integrator Migration Guide*

Chapter title: After you migrate

Correction

Step 5b in the procedure contains an incorrect command name. The correct command is **db2iupdt**.

Topic title: Accessing ODBC data sources after migration

Location in the DB2 Information Center

Migrating to DB2 Information Integrator -> Accessing data sources after migration. This topic was last updated for DB2 Information Integrator Version 8.1.

Location in PDF and printed book

Book title: *DB2 Information Integrator Migration Guide*

Chapter title: After you migrate

Correction

After the procedure, add the following text:

If you use transparent DDL to create remote tables on ODBC data sources through DataJoiner, you might receive errors with SELECT statements after you migrate to DB2 Information Integrator Version 8. The DataJoiner DATE data type is mapped to the ODBC DATETIME data type when you create a remote ODBC table through DataJoiner.

You can avoid these errors after you migrate to DB2 Information Integrator Version 8 by changing the local data type in the federated database systems catalog from DATE to TIMESTAMP.

Topic title: Accessing Oracle data sources after migration

Location in the DB2 Information Center

Migrating to DB2 Information Integrator -> Accessing data sources after migration. This topic was last updated for DB2 Information Integrator Version 8.1.

Location in PDF and printed book

Book title: *DB2 Information Integrator Migration Guide*

Chapter title: After you migrate

Correction

Step 5b in the procedure contains an incorrect command name. The correct command is **db2iupdt**.

Topic title: Accessing Sybase data sources after migration

Location in the DB2 Information Center

Migrating to DB2 Information Integrator -> Accessing data sources after migration. This topic was last updated for DB2 Information Integrator Version 8.1.

Location in PDF and printed book

Book title: *DB2 Information Integrator Migration Guide*

Chapter title: After you migrate

Correction

Step 5b in the procedure contains an incorrect command name. The correct command is **db2iupdt**.

Correction

Before you migrated to DB2 Information Integrator Version 8, you copied the interfaces file. In the procedure, add a new step before step 1: In the federated database instance, restore the interfaces file back into the \$HOME/sqllib directory.

Topic title: Accessing Teradata data sources after migration

Location in the DB2 Information Center

Migrating to DB2 Information Integrator -> Accessing data sources after migration. This topic was last updated for DB2 Information Integrator Version 8.1.

Location in PDF and printed book

Book title: *DB2 Information Integrator Migration Guide*

Chapter title: After you migrate

Correction

Step 3b in the procedure contains an incorrect command name. The correct command is **db2iupdt**.

Topic title: DataJoiner migration restrictions

Location in the DB2 Information Center

Migrating to DB2 Information Integrator -> Planning to migrate to DB2 Information Integrator Version 8. This topic was last updated for DB2 Information Integrator Version 8.1.

Location in PDF and printed book

Book title: *DB2 Information Integrator Migration Guide*

Chapter title: DB2 Information Integrator migration overview

Correction

Add a new section titled "Sybase server options" with the following text:

DB2 Information Integrator, Version 8, does not support the Sybase server options APPLY_BUFFER_SIZE and APPLY_PACKET_SIZE.

Topic title: DataJoiner migration restrictions

Location in the DB2 Information Center

Migrating to DB2 Information Integrator -> Planning to migrate to DB2 Information Integrator Version 8. This topic was last updated for DB2 Information Integrator Version 8.1.

Location in PDF and printed book

Book title: *DB2 Information Integrator Migration Guide*

Chapter title: DB2 Information Integrator migration overview

Correction

The topic describes certain LONG and LOB data type replication restrictions for migrating from DataJoiner to DB2 Information Integrator. These restrictions state that you might not be able to replicate large data objects.

However, you can alter certain data type columns that contain large data objects to VARCHAR data types to enable replication.

When you migrate from DataJoiner to DB2 Information Integrator Version 8, the large object data types of your remote data sources are mapped to local CLOB or BLOB data

types. With the exception of the Oracle NET8 wrapper, federated servers cannot write to DB2 CLOB or BLOB data types in nicknames, and these data types cannot be replicated.

If you need to enable replication of these data types, you must alter these local CLOB and BLOB data type columns to VARCHAR data types to give DB2 replication insert and update privileges and to enable replication of these columns.

You can now enable replication of these large data objects for the following data sources:

DB2 Universal Database

The DRDA wrapper maps remote LONG VARCHAR and LONG VARCHAR FOR BIT DATA data types by default to local CLOB and BLOB data types. You can alter these local CLOB and BLOB data type columns to VARCHAR and VARCHAR FOR BIT DATA data types with lengths up to 32672 characters.

Oracle The NET8 and SQLNET wrappers map remote LONG and LONG RAW data types by default to local CLOB and BLOB data types. You can now alter these local CLOB and BLOB data type columns to VARCHAR and VARCHAR FOR BIT DATA data types with lengths up to 32672 characters.

Informix

The INFORMIX wrapper maps remote TEXT and BYTE data types by default to local CLOB and BLOB data types. You can alter these local CLOB and BLOB data type columns to VARCHAR and VARCHAR FOR BIT DATA data types with lengths up to 32672 characters.

Sybase

The CTLIB wrapper maps remote TEXT and IMAGE data types by default to local CLOB and BLOB data types. You can alter these local CLOB and BLOB data type columns to VARCHAR and VARCHAR FOR BIT DATA data types with lengths up to 32672 characters. You cannot alter large data objects if you are using the DBLIB wrapper.

Microsoft SQL Server

The DJXMSSQL3 wrapper and the MSSQLODBC3 wrapper map remote TEXT and IMAGE data types to local CLOB and BLOB data types. You can alter these local CLOB and BLOB data type columns to

VARCHAR and VARCHAR FOR BIT DATA data types with lengths up to 32672 characters.

TERADATA

Remote CHAR and VARCHAR data types with lengths that are greater than 32672 character are mapped by default to local CLOB data types. Remote BYTE and VARBYTE data types with lengths that are greater than 32672 character are mapped by default to local BLOB data types. You can alter these local CLOB and BLOB data type columns to VARCHAR and VARCHAR FOR BIT DATA data types with lengths up to 32672 characters.

For example, suppose that you want to use a remote table that is named EMP on your DB2 remote data source. The EMP table contains a column named EMP_INFO with a remote data type of LONG VARCHAR. You then create a nickname named EMPLOYEE that references the DB2 table EMP_INFO. The nickname that you create contains a local column called INFO that references the EMP_INFO column on the remote table.

When you create a nickname for this remote table, the DRDA wrapper maps the data type of the EMP_INFO column by default to a local CLOB data type.

To enable replication of the local data in the INFO column, alter the data type of this column from CLOB to VARCHAR. For example:

```
ALTER NICKNAME EMPLOYEE
ALTER COLUMN INFO
LOCAL TYPE VARCHAR(32672)
```

Topic title: DB2 Version 7 federated system migration restrictions

Location in the DB2 Information Center

Migrating to DB2 Information Integrator -> Planning to migrate to DB2 Information Integrator Version 8. This topic was last updated for DB2 Information Integrator Version 8.1.

Location in PDF and printed book

Book title: *DB2 Information Integrator Migration Guide*

Chapter title: DB2 Information Integrator migration overview

Correction

Add a new section titled "Sybase server options" with the following text:

DB2 Information Integrator, Version 8, does not support the Sybase server options APPLY_BUFFER_SIZE and APPLY_PACKET_SIZE.

For instructions on migrating UNIX or Windows databases to a new computer, see the topic "Migrating a database to a different computer" on the DB2 Information Integrator Web site at www.ibm.com/software/data/integration/db2ii/.

DB2 Information Integrator Data Source Configuration Guide

Topic title: PubMed and Nucleotide schema tables

Location in the DB2 Information Center

Configuring federated systems and data sources→
Configuring data sources→ Configuring Entrez data sources.
This topic was last updated for DB2 Information Integrator Version 8.1.

Location in PDF and printed book

Book title: *DB2 Information Integrator Data Source Configuration Guide*

Chapter title: Configuring access to Entrez data source

Correction

A row in the Nucleotide DBSeq nickname table in the "Nucleotide schema" section is incorrect. The data type in the Sequence row should be VARCHAR (32000) instead of CLOB.

Topic title: User mapping options for federated servers

Location in the DB2 Information Center

Reference information→ Federated systems reference. This topic was last updated for DB2 Information Integrator Version 8.1.

Location in PDF and printed book

Book title: *DB2 Information Integrator Data Source Configuration Guide*

Appendix title: User mapping options for federated systems

Correction

The name of the accounting option should be ACCOUNTING instead of ACCOUNTING_STRING.

Topic title: Adding Oracle data sources to federated servers

Location in the DB2 Information Center

Configuring federated systems and data sources→
Configuring data sources → Configuring Oracle data sources.
This topic was last updated for DB2 Information Integrator
Version 8.1.

Location in PDF and printed book

Book title: *DB2 Information Integrator Data Source
Configuration Guide*

Chapter title: Configuring access to Oracle data sources

Correction

In the Prerequisites section, add NLS_LANG to the list of
variables that should be set before you configure the federated
server to access Oracle data sources.

Topic title: CREATE NICKNAME statement - Example for Documentum wrapper

Location in the DB2 Information Center

Configuring federated systems and data sources→
Configuring data sources → Configuring Documentum data
sources. This topic was last updated for DB2 Information
Integrator Version 8.1.

Location in PDF and printed book

Book title: *DB2 Information Integrator Data Source
Configuration Guide*

Chapter title: Configuring access to Documentum data
sources

Correction

The maximum length for the Documentum string data types
object_name and title is not 255. The maximum length
depends on the operating system that the federated server
uses. On Windows federated servers, the maximum length is
260. On UNIX federated servers, the maximum length is 1024.

Topic title: CREATE NICKNAME statement - Examples for XML wrapper

Location in the DB2 Information Center

Samples and examples→ Data source configuration examples
→XML data source examples. This topic was last updated for
DB2 Information Integrator Version 8.1.

Location in PDF and printed book

Book title: *DB2 Information Integrator Data Source
Configuration Guide*

Chapter title: Configuring access to XML data sources

Correction

In the example for the payment nickname view, the column name p.amount in the select clause should be p.number.

In the example for the item nickname view, the columns named it.quantity and it.name in the select clause should be i.quantity and i.name.

Topic title: Registering nicknames for Excel data sources

Location in the DB2 Information Center

Configuring federated systems and data sources→
Configuring data sources → Configuring Excel data sources.
This topic was last updated for DB2 Information Integrator Version 8.1.

Location in PDF and printed book

Book title: *DB2 Information Integrator Data Source Configuration Guide*

Chapter title: Configuring access to Excel data sources

Correction

The date range restriction no longer applies. The Excel wrapper supports the same date range as the date range that is supported by the Microsoft Excel application.

Topic title: Extended Search vertical tables

Location in the DB2 Information Center

Configuring federated systems and data sources→
Configuring data sources → Configuring Extended Search data sources. This topic was last updated for DB2 Information Integrator Version 8.1.

Location in PDF and printed book

Book title: *DB2 Information Integrator Data Source Configuration Guide*

Chapter title: Configuring access to Extended Search data sources

Correction

In the table on Column Name row for FIELD_DATATYPE, remove DECIMAL from the list of data types.

Topic title: Messages for the Extended Search wrapper

Location in the DB2 Information Center

Configuring federated systems and data sources→

Configuring data sources -> Configuring Extended Search data sources. This topic was last updated for DB2 Information Integrator Version 8.

Location in PDF and printed book

Book title: *DB2 Information Integrator Data Source Configuration Guide*

Chapter title: Configuring access to Extended Search data sources

Correction

The following error message will be returned when a query that uses an Extended Search nickname contains predicates in the WHERE clause that cannot be handled by the Extended Search Wrapper.

SQL0901N The SQL statement failed because of a non-severe system error. Subsequent SQL statements can be processed.
(Reason "sqlno_crulte_remote_pushdow[Z00]:rc(-2144272270) Error generatin".) SQLSTATE=58004

Examples of queries that will return the SQL0901N error:

- Queries without a predicate
- Queries that use the predicates ANY, ALL, SOME, or EXIST
- Queries that use the predicates IN or NOT IN with a FULL SELECT statement
- Queries that use a fixed column predicate, such as WHERE DOC_ID = 'ABC'

DB2 Information Integrator Federated Systems Guide

Topic title: User mapping options for federated servers

Location in the DB2 Information Center

Reference information-> Federated systems reference. This topic was last updated for DB2 Information Integrator Version 8.1.

Location in PDF and printed book

Book title: *DB2 Information Integrator Federated Systems Guide*

Appendix title: User mapping options for federated systems

Correction

The name of the accounting option should be ACCOUNTING instead of ACCOUNTING_STRING.

Topic title: Default forward data type mappings (Sybase)

Location in the DB2 Information Center

Reference information→ Federated systems reference. This topic was last updated for DB2 Information Integrator Version 8.1.

Location in PDF and printed book

Book title: *DB2 Information Integrator Federated Systems Guide*

Appendix title: Default forward data type mappings

Correction

Sybase automatically converts the CHAR NULL data type to the VARCHAR data type. See the VARCHAR data type row for the default data type mapping of the CHAR NULL data type.

Sybase automatically converts the NCHAR NULL data type to NVARCHAR. See the NVARCHAR data type row for the default forward data type mapping of the NCHAR NULL data type.

Topic title: Default forward data type mappings (Informix)**Location in the DB2 Information Center**

Reference information→ Federated systems reference. This topic was last updated for DB2 Information Integrator Version 8.1.

Location in PDF and printed book

Book title: *DB2 Information Integrator Federated Systems Guide*

Appendix title: Default forward data type mappings

Correction

The REMOTE_UPPER_LEN value for the remote type of DECIMAL that is mapped to the federated type DOUBLE has changed from 32 to 130.

DB2 Information Integrator Developer's Guide**Topic title: Installing DB2 WebSphere MQ Series****Location in the DB2 Information Center**

Developing applications→ Developing applications that use WebSphere MQ Series messaging. This topic was last updated for DB2 Information Integrator Version 8.1.

Location in PDF and printed book

Book title: *DB2 Information Integrator Developer's Guide*

Chapter title: Developing database applications that exploit WebSphere Message Queue functions

Corrections

Add Application Messaging Interface, Version 1.2 or later, as a prerequisite for the MQ user-defined functions. Add DB2 UDB XML Extender as a prerequisite for the MQ user-defined functions.

Topic title: WebSphere Application Server for administering Web applications

Location in the DB2 Information Center

Developing applications-> Developing Web services and Web applications-> Deploying and testing your Web application. This topic was last updated for DB2 Information Integrator Version 8.1.

Location in PDF and printed book

Book title: *DB2 Information Integrator Developer's Guide*

Chapter title: Deploying and testing your Web application

Corrections

Change the references from WebSphere Application Server Express to application server for DB2. Add installation information for the application server. Add pre-installation requirement for the DB2 application server.

Topic title: Performance and tuning planning - federated systems and materialized query tables

Location in the DB2 Information Center

Product overview-> Information integration — overview-> Extending your data warehouse — a solution example-> Warehouse example — Cottonwood Distributors, Inc.-> Deploying the application — warehouse example-> to Performance and tuning planning. This topic was last updated for DB2 Information Integrator Version 8.1.

Location in PDF and printed book

Book title: *DB2 Information Integrator Developer's Guide*

Chapter title: Overview of information integration solutions

Corrections

Remove all references to REFRESH IMMEDIATE relating to materialized query tables.

Notices

This information was developed for products and services offered in the U.S.A. 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.

Licenses 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 Corporation
J46A/G4
555 Bailey Avenue
San Jose, CA 95141-1003
U.S.A.

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:

IBM
AIX
CICS
DataJoiner
DB2
DB2 Connect
DB2 Universal Database
DRDA
Informix
iSeries
OS/390
z/OS

The following terms are trademarks or registered trademarks of other companies:

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

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

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

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



Printed in USA