IBM DB2 Information Integrator

Release Notes

Version 8.1 Fix Pack 9

IBM DB2 Information Integrator

Release Notes

Version 8.1 Fix Pack 9

Before using this information and the product it supports, be sure to read the general information under "Notices" on page 89.

Contents

About the release notes

The release notes contain the latest information about DB2® Information Integrator Version 8.1 Fix Pack 9.

The release notes are in three parts. The first part contains high-level information about the new functions and features of the product. 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. 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 main DB2 Information Integrator support Web page at www.ibm.com/software/data/integration/db2ii/support.html.

Support for DB2 Universal Database™ for Linux®, UNIX®, and Windows®

For information about known issues with DB2 Universal Database™ for Linux, UNIX, and Windows Version 8 products, see the DB2 Technical Support Web page at www.ibm.com/software/data/db2/udb/winos2unix/support.

New in this release

The following features and changes are new in fix pack 9:

IBM DB2 Information Integrator software portfolio rebranded as IBM WebSphere® Information Integrator

IBM DB2 Information Integrator is now WebSphere Information Integrator. The WebSphere brand represents integration and openness which are also integral to the Information Integrator software portfolio and its role in providing organizations with real-time, integrated access to disparate and distributed information across IBM and non-IBM sources. Find more information at www.ibm.com/software/data/integration/launch.html. All the fix pack 9 documentation refers to the DB2 Information Integrator name; all Web site information refers to the new WebSphere branding.

New feature: KEGG API Version 3.2 now supported

The KEGG API that the DB2 Information Integrator KEGG user-defined functions use has changed. The new KEGG API, version 3.2, is not compatible with KEGG API Version 2.3. Version 2.3 will not be supported by the KEGG API after 2005/03/31. For more information about the new API, see the KEGG Version 3.2 API Web site www.genome.jp/kegg/soap/doc/keggapi_manual.html

The KEGG user-defined functions that are provided with DB2 Information Integrator have been updated to support the new API.

Summary of changes to the KEGG user defined functions:

> Significant changes were made to the KEGG user defined functions as a result of the version 3.2 API:
> - Three new KEGG user-defined functions are included in this fix pack.
> - Four of the KEGG user-defined functions are no longer supported.
> - Four of the names for SSDB functions have changed.
> - All of the names of the function arguments have changed.
> - Some of the descriptions for the function arguments have changed.
> - Five new function arguments are added to the KEGG user-defined functions.
> - For SSDB table functions, several new types of rows are returned.

> For additional information see the "New documentation" on page 39 section.

New feature: IMPORT and EXPORT commands supported for nicknames

IMPORT command
> You can use the IMPORT command to import data into a nickname. The IMPORT command is supported with the DB2 (DRDA), Informix®, Microsoft® SQL Server, Oracle, Sybase, and Teradata data sources.

> **Limitation**: Importing data into a nickname for a DB2 Universal Database for VM remote table does not work correctly for binary data (FOR BIT DATA) columns.

If you upgrade an existing federated database to fix pack 9, you need to update your database running the db2updv8 utility. The db2updv8 utility must be run on each federated database or the IMPORT command will fail.

The db2updv8 utility creates two stored procedures that are required when you use the IMPORT command with nicknames. If you create a federated database after you apply fix pack 9, the new database will include the new stored procedures and you do not need to run the db2updv8 utility.

The following rules apply when you use the IMPORT command to import data into a nickname:

- The remote object on which the nickname is defined must be a table. You cannot import data into a nickname that is defined on a view or synonym.
- The supported file types are IXF, ASC and DEL.
- The ALLOW WRITE ACCESS clause must be specified. This clause invokes the online import mode. The ALLOW WRITE ACCESS clause allows concurrent applications read and write access to the import target table.
- You cannot use the COMMITCOUNT AUTOMATIC mode with nicknames.
- The COMMITCOUNT *n* must be specified with *n* being a valid nonzero number.
- Only the INSERT and INSERT_UPDATE operations are supported with nicknames.
- The Datalink, LOBs, and generated columns column types are not supported with nicknames.
- The following file type modifiers are not supported with nicknames:
  - dldelfiletype
  - generatedignore
  - generatedmissing
  - identityignore
  - identitymissing
  - indexixf
  - indexschema
  - lobsinfile
  - nodefaults
  - no_type_idfiletype
  - usedefaults
- Hierarchy (typed table) is not supported with nicknames.

If you submit an IMPORT command that does not adhere to these rules, you will receive SQL27999N. For example:

```
SQL27999N The requested IMPORT operation into a remote target
(nickname) cannot be performed.  Reason code = "reason_code".
```

EXPORT command

You can use the EXPORT command to export data from a nickname. The EXPORT command is supported for all wrappers.

You do not need to run the db2updv8 utility to use the EXPORT command with nicknames.

The following restrictions apply when you use the EXPORT command to export data from a nickname:

- The description of the target table that is necessary to perform the Import CREATE operation will not be saved in the IXF file format. Use the db2look utility to collect the information that you need to recreate the table.
- Only IXF and DEL file types are supported.

Example of importing data from an IXF filetype

This example uses the INSERT option to import data from an IXF filetype:

```
IMPORT FROM import_file_1.ixf OF IXF
   ALLOW WRITE ACCESS
   COMMITCOUNT 20
   INSERT INTO NICKNAME_1;
```

Example of the importing data from a DEL filetype

This example uses the INSERT_UPDATE option to import data from a DEL filetype:

```
IMPORT FROM import_file_1.del OF DEL
   ALLOW WRITE ACCESS
   COMMITCOUNT 50
   INSERT_UPDATE INTO NICKNAME_1;
```

Example of importing data from an ASC filetype by using the INSERT option:

This example uses the INSERT option to import data from an ASC filetype. The example includes the STRIPTBLANKS file modifier to truncate any trailing blank spaces in the data. The METHOD L parameter specifies the start and end of the column numbers.

```
IMPORT FROM import_file_1.asc OF ASC MODIFIED BY STRIPTBLANKS
   METHOD L(1 6, 8 32, 34 44, 46 48)
   ALLOW WRITE ACCESS
   COMMITCOUNT 20
   INSERT INTO NICKNAME_1;
```

New feature: Oracle and Sybase targets now supported in Q replication

You can now use the Q replication function that is available in DB2 Information Integrator Replication Edition to replicate DB2 data to tables in Oracle and Sybase relational database management systems. This new feature uses the federated functionality provided with DB2 Information Integrator, and the low-latency, high-throughput ability of Q replication. You can find a white paper that fully describes the new functions, and a tutorial for replicating to Oracle targets, under "Latest articles" at the IBM developerWorks Web site:www.ibm.com/developerworks/db2/zones/db2ii/

Known problems, limitations, and workarounds

The following information describes the currently known limitations, problems, and workarounds for DB2® Information Integrator Version 8.2. Any limitations and restrictions might or might not apply to other releases of the product.

DB2 Information Integrator installation

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

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

- Install DB2 Information Integrator installer using the provided CDs.
  UNIX systems only: Ensure there are no space characters in the mount point name.
- Ensure that no blank characters are used in the directory path name in which the DB2 Information Integrator installer or the DB2 installer is copied.

The DB2 Information Integrator installer fails when the DB2 Administration Server and database instance user names are identical

The DB2 Administration Server provides support services for DB2 Universal Database tools such as the Control Center and Configuration Assistant. Every computer that uses DB2 Universal Database also contains a DB2 Administration Server.

If DB2 Universal Database is not installed on the computer on which you are installing DB2 Information Integrator, the DB2 Information Integrator Installer will automatically install a DB2 Administration Server. The DB2 Information Integrator Installer prompts you to specify names for the DB2 Administration Server user and for the related database instance user. You must specify unique names for the DB2 Administration Server user and for the related database instance user, or the installation will fail.

To avoid this installation failure, specify a database instance user name that is different from the current DB2 Administration Server user name during the DB2 Information Integrator installation process.

The DB2 Information Integrator installer fails on Red Hat Enterprise Linux Version 3.0

If you are trying to install DB2 Information Integrator on a computer that runs the Red Hat Enterprise Linux Version 3.0 operating system, the DB2 Information Integrator Installer will fail.

To avoid this problem, enter the following command before you start the DB2 Information Integrator Installer:

```
export LD_ASSUME_KERNEL=2.2.5
```

Required task after installing DB2 Information Integrator upgrades or fix packs

After you install a DB2 Information Integrator upgrade or fix pack, you must run the djxlink script for each of the wrappers that you installed and the db2iupdt command for each DB2 instance. See the DB2 Information Integrator support site at www.ibm.com/software/data/integration/db2ii/support.html for instructions on installing DB2 Information Integrator fix packs and running the djxlink script and the db2iupdt command.

DB2 Information Integrator installer Asian fonts limitation (UNIX systems only)

IBM offers additional font packages for UNIX that contain additional double-byte character set support for Asian characters. The font packages that IBM offers are necessary with some versions of UNIX that install only the fonts that are necessary to display the country-specific or region-specific characters.

If characters are missing when you run the iisetup program, the necessary fonts are likely not installed on your UNIX system.

To enable the **iisetup** program to properly reference the fonts embedded in the installation CD:

1. Copy the tnrmt30.ttf font file to a directory on your system. The file is in the root directory of the DB2 Information Integrator installation software.
2. Export the JAVA_FONT environment variable by running the following command:

   ```
   export JAVA_FONT=<path to font>
   ```

   <path to font> is the path to the directory where you copied the font file.
3. Run the **iisetup** program.
4. Remove the font file from the hard disk.

The fonts that IBM supplies do not replace the system fonts. Use the IBM-supplied fonts with DB2 Information Integrator. You cannot engage in the general or unrestricted sale or distribution of these fonts.

Migration

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 Universal Database Version 7 or from DataJoiner® Version 2.1.1, you must start the bind utility to reestablish a connection from the server to the remote data source. DB2 family remote data sources include databases on DB2 Universal Database for UNIX, Linux, and Windows, DB2 Universal Database for z/OS™ and OS/390®, DB2 Universal Database for iSeries™, and DB2 Server for VM and VSE.

If you do not reestablish the connection by using the bind utility, you receive an SQL0805N error message.

You must use 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 reestablish a connection to a remote data source after migration:
1. Access the remote data source with the DB2 CONNECT command.
2. Access the required DB2 Information Integrator directory.

   On UNIX systems, issue the following command:

   ```
   cd db2_instance_directory/sqllib/bnd
   ```

   db2_instance_directory is the directory that contains the database instance that you want to access.

   On Windows systems, issue the following command:

   ```
   cd db2_install_directory\sqllib\bnd
   ```

   db2_install_directory is the directory where DB2 Information Integrator is installed.
3. Access the bind utility by issuing the following command:

   ```
   db2 bind @db2cli.lst grant public
   ```
4. Disconnect from the remote data source.

Federated function and data source support

Lock escalation when using Q Replication to copy a high volume of updates to Sybase targets

When using Q Replication to copy a high volume of updates to Sybase targets, lock escalation can occur at the federated target server. In this case, it may be necessary to change the locking scheme of target tables to lock data rows. It may also be necessary to increase the maximum locks per server.

Performance impact on Teradata nicknames that have multicolumn indexes

Before DB2 Information Integrator Version 8.2 Fix Pack 9, nicknames that were created on Teradata tables with multicolumn indexes did not correctly store the column order for those indexes. With Fix Pack 9, the create nickname process will retrieve the index columns in the correct order. Because the optimizer will have more accurate index information, query performance might improve for those queries that can use the index. To take advantage of this potential performance improvement, drop and recreate any Teradata nicknames that have multicolumn indexes after you install DB2 Information Integrator Version 8.2 Fix Pack 9. After you drop and recreate the Teradata nicknames, performance might improve for some queries because the optimizer will receive accurate index information from the Teradata server.

SQL0901N occurs when creating data type mapping

If you create a data type mapping for something other than a specific server, SQL0901N can occur. For example, if you create a mapping for a server type, using the CREATE SERVER statement to define a server, the CREATE SERVER statement may cause the following error:

```
SQL0901N The SQL statement failed because of a non-severe system error.
Subsequent SQL statements can be processed. (Reason "sqlqgGetRemoteTypeInfo:
Failed to create server object".) SQLSTATE=58004
```

Possible workaround:
Before any data type mapping definition, define the server with the CREATE SERVER statement, and then define the type mappings using the CREATE TYPE MAPPING statement.

DB2 for VM nicknames

Inserting or importing binary (FOR BIT DATA) columns is not supported for DB2 for VM nicknames.

Replacing the KEGG user-defined functions

If you installed a version of the KEGG user-defined functions before DB2 Information Integrator Version 8.2 Fix Pack 9, you must replace the version 2.3 KEGG user-defined functions with the version 3.2 KEGG user-defined functions.

Procedure:

To replace the KEGG user-defined functions:

1. Install fix pack 9. The KEGG user-defined functions are installed on the federated server when you select **Life sciences user-defined functions** under the Nonrelational Wrappers component in the DB2 Information Integrator installation wizard.
2. Run the disable_KEGGFunctions command and include the [-ver2.3] flag.
   - On federated servers that run Windows, this command is in the `sqllib\bin` directory.
   - On federated servers that run UNIX, this command is in the `sqllib/bin` directory.

   Syntax:
   ```
   disable_KEGGFunctions -n dbName -u userID -p password [-ver2.3]
   ```

   -n *dbName*
   The name of the federated database that you want to disable the functions from.

   -u *userID*
   A valid user ID for the federated database.

   -p *password*
   A valid password for the user ID.

   [-*ver2.3*]
   Removes all of the version 2.3 KEGG user-defined functions from the *dbName* that you specify.

3. Run the enable_KEGGFunctions command to register the updated KEGG user-defined functions.
   - On federated servers that run Windows, this command is in the `sqllib\bin` directory
   - On federated servers that run UNIX, this command is in the `sqllib/bin` directory

   Syntax:
   ```
   enable_KEGGFunctions -n dbName -u userID -p password
   ```

-n *dbName*
> The name of the federated database that you are registering the functions in.

-u *userID*
> A valid user ID for the federated database.

-p *password*
> A valid password for the user ID.

Data type restrictions for replication to Oracle and Sybase targets (Q replication only)

The following data type limitations apply when you are using Q replication to replicate data to Oracle and Sybase targets.

Required configuration to replicate GRAPHIC, VARGRAPHIC, and DBCLOB data types

To replicate GRAPHIC, VARGRAPHIC, or DBCLOB data types, your Oracle server and client must be Version 9 or newer. Your server mapping must also be Version 9 or newer.

Replication of LONG VARGRAPHIC data types not supported

Replication of LONG VARGRAPHIC data types to Oracle and Sybase targets is not supported in fix pack 9.

Performance impact when altering nickname column types

When you alter the column type of a nickname, you might impact performance. Under certain conditions, II can no longer push down SQL that references the column or II might need to do additional data conversion. To prevent a possible performance degradation, do not alter the nickname column type.

Unexpected results when altering nickname column length

Do not alter a nickname column length to a value smaller than the original length. If the data doesn't fit, unexpected results or errors can occur.

Requirement for federated server computers (AIX® 64-bit systems only)

To use an AIX 64-bit computer as a federated server, you must install AIX APAR fix number IY53887 on that computer.

If you do not install AIX APAR fix number IY53887, your wrappers might not load correctly when you use a CREATE SERVER statement or CREATE NICKNAME statement. If your wrappers do not load correctly, you receive an SQL0901N error that indicates NULL WRAPPER.

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 1. Unsupported data source data types

Data source	Unsupported data types
Extended Search	DECIMAL
Microsoft® SQL Server	SQL_VARIANT
Oracle (NET8 wrapper only)	TIMESTAMP (fractional_seconds_precision) WITH TIME ZONE TIMESTAMP (fractional_seconds_precision) WITH LOCAL TIME ZONE INTERVAL YEAR (year_precision) TO MONTH INTERVAL DAY (day_precision) TO SECOND (fractional_seconds_precision) UROWID BFILE XMLTYPE URI Datatypes SPATIAL Datatypes

Oracle (SQLNET wrapper only)	BLOB
	CLOB
	NCHAR
	NVARCHAR2
	TIMESTAMP
	TIMESTAMP (fractional_seconds_precision) WITH TIME ZONE
	TIMESTAMP (fractional_seconds_precision)
	    WITH LOCAL TIME ZONE
	NCLOB
	INTERVAL YEAR (year_precision) TO MONTH
	INTERVAL DAY (day_precision) TO SECOND
	    (fractional_seconds_precision)
	UROWID
	BFILE
	XMLTYPE
	URI Datatypes
	SPATIAL Datatypes

Sybase (CTLIB wrapper only)	DATE TIME
Sybase (DBLIB wrapper only)	DATE TIME 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. Or 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.

Where upper and lower length limits are provided in the following table, character data types specify the length in bytes. Graphic data types specify the length in double bytes.

Table 2. Write operation restrictions on data types

Data source	Unsupported data types
DB2 family (DRDA®)	BLOB CLOB DBCLOB
Informix®	BLOB CLOB BYTE** TEXT* * You can perform insert, update, and delete operations on TEXT data types if you alter the nickname column to the VARCHAR data type. ** You can perform insert, update, and delete operations on BYTE data types if you alter the nickname column to the VARCHAR FOR BIT DATA data type.
Microsoft SQL Server	IMAGE** NTEXT* TEXT* SQL_VARIANT * You can perform insert, update, and delete operations on text and ntext data types if you alter the nickname column to the VARCHAR or VARGRAPHIC data type. ** You can perform insert, update, and delete operations on image data types if you alter the nickname column to the VARCHAR FOR BIT DATA data type.
ODBC	SQL_LONGBINARY (length > 32672) SQL_LONGVARCHAR (length > 32672) SQL_WLONGVARCHAR (length > 32672)
Oracle (NET8 wrapper only)	INTERVAL DAY (day_precision) TO SECOND (fractional_seconds_precision) INTERVAL YEAR (year_precision) TO MONTH LONG* LONG RAW** TIMESTAMP (fractional_seconds_precision) WITH TIMEZONE TIMESTAMP (fractional_seconds_precision) WITH LOCAL TIME ZONE INTERVAL YEAR (year_precision) TO MONTH INTERVAL DAY (day_precision) TO SECOND (fractional_seconds_precision) UROWID BFILE XMLTYPE URI Datatypes SPATIAL Datatypes * You can perform insert, update, and delete operations on LONG data types if you alter the nickname column to the VARCHAR data type. ** You can perform insert, update, and delete operations on LONG RAW data types if you alter the nickname column to the VARCHAR FOR BIT DATA data type.

Table 2. Write operation restrictions on data types (continued)

Data source	Unsupported data types
Oracle (SQLNET wrapper only)	BLOB CLOB INTERVAL DAY (day_precision) TO SECOND (fractional_seconds_precision) INTERVAL YEAR (year_precision) TO MONTH NCHAR NVARCHAR2 NCLOB TIMESTAMP (fractional_seconds_precision) TIMESTAMP (fractional_seconds_precision) WITH TIMEZONE TIMESTAMP (fractional_seconds_precision) WITH LOCAL TIMEZONENCLOB INTERVAL YEAR (year_precision) TO MONTH INTERVAL DAY (day_precision) TO SECOND (fractional_seconds_precision) UROWID BFILE XMLTYPE URI Datatypes SPATIAL Datatypes
Sybase (CTLIB wrapper only)	DATE TIME TEXT* IMAGE** * You can perform insert, update, and delete operations on text data type if you alter the nickname column to the VARCHAR data type. ** You can perform insert, update, and delete operations on image data types if you alter the nickname column to the VARCHAR FOR BIT DATA data type.
Sybase (DBLIB wrapper only)	All data types. Write operations are not supported by the DBLIB wrapper.
Teradata	CHAR (32673 to 64000) VARCHAR (32673 to 64000) BYTE (32673 to 64000) VARBYTE (32673 to 64000) GRAPHIC (16337 to 32000) VARGRAPHIC (16337 to 32000)

Mapping to the federated GRAPHIC and VARGRAPHIC data types

If the federated database does not use the UTF-8 code page (Unicode), 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.

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 process might result in data expansion. For example, if you select a character that is stored as 1 byte at the remote data source, the data source might return 2 bytes after converting the selected character to UTF-8. 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.

Oracle national character set data (NCHAR, NVARCHAR2, NCLOB) might appear distorted in a federated database

Data code conversion for NCHAR, NVARCHAR2, and NCLOB data types in a Unicode federated environment might not be performed correctly if you are using a version of Oracle earlier than Oracle 9i. Data can appear distorted because the national character set has changed from non-Unicode, to Unicode in Oracle9i. DB2 Information Integrator provides NCHAR, NVARCHAR2, and NCLOB data type support that is consistent with Oracle9i specifications.

LIKE predicate applied to a CHAR column

For Informix, Microsoft SQL Server, Oracle, and Sybase data sources, the LIKE predicate that is applied to a CHAR column is not processed at the data source. These data sources use different blank padding rules than DB2 Universal Database. For example, if a CHAR(10) column contains 'a', then the predicate char_col LIKE 'a' is false in DB2 Universal Database, 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 processed at the data source.

You can improve performance of LIKE predicates by enabling them to be processed at the data source. To do so, use the following temporary CREATE FUNCTION MAPPING statement syntax to create a function mapping for LIKE(CHAR, *column_type* (*length*)) so that the LIKE predicate is processed at the data source. For example:

```
CREATE FUNCTION MAPPING my_mapping FOR SYSIBM.LIKE(SYSIBM.CHAR(), SYSIBM.VARCHAR())
   SERVER TYPE INFORMIX OPTIONS(REMOTE_NAME ':1P LIKE :2P');
```

However, if you use this function mapping, LIKE predicates might return different results than DB2 Universal Database returns.

Unexpected results when using a Unicode federated database and a remote database with DBCS codepage

When the federated database uses Unicode and the remote database uses a DBCS codepage, you can get unexpected results because of the different ways spaces are represented:

- In DBCS databases, CHAR and GRAPHIC columns are usually padded with the DBCS space.
- In federated Unicode databases, CHAR and GRAPHIC columns are padded with the Unicode space (U+0020).
- When character data is retrieved from a DBCS database to a federated Unicode database, DBCS spaces are usually converted to Unicode ideographic spaces (U+3000).
- When character data is sent from a federated Unicode database to a remote DBCS database, Unicode spaces (U+0020) are usually converted to the substitute character (because most DBCS codepages have no equivalent to a Unicode space).

As a result of these differences, you might get unexpected results when specifying any of the following:

- Predicates using CHAR columns, GRAPHIC columns, or character data containing Unicode spaces.

- Inserts into CHAR or GRAPHIC columns, or inserts of character data that contain Unicode spaces into a Unicode nickname.
- Updates of CHAR or GRAPHIC columns, or updates using character data that contains Unicode spaces on a Unicode nickname.

To avoid these problems:
- Use only ideographic spaces (U+3000) in the federated Unicode database.
- Avoid padding by casting CHAR and GRAPHIC columns to VARCHAR and VARGRAPHIC. For example: CAST(my_col AS VARCHAR(x)).

For additional related information see Japanese or Traditional Chinese EUC, and UCS-2 Considerations in C and C++ in the DB2 information center.

Cursor with hold semantics for certain relational wrappers

The Microsoft SQL Server, ODBC, OLE DB, Oracle, Sybase, and Teradata wrappers do not currently support DECLARE CURSOR WITH HOLD semantics.

Limitation on file code page specifications for Unicode support of table-structured file and XML data sources (Windows only)

On Windows operating systems, you can create file names and directory names using any Unicode character. However, the FILE_PATH and DIRECTORY_PATH options of the CREATE NICKNAME and ALTER NICKNAME statements must use characters that the operating system code page supplies. If you specify Unicode characters in either of the PATH nickname options using either the table-structured file wrapper or the XML wrapper, and those Unicode characters are not part of the Windows code page, the wrapper issues an error that indicates that the resource cannot be located.

For example, on an English language Windows-based computer, you cannot specify a FILE_PATH option value that uses Japanese Unicode characters, because the create or alter nickname operation will fail. In this situation, you must specify a FILE_PATH option value that uses only English characters. You can also use ASCII characters for FILE_PATH and DIRECTORY_PATH option values in a Japanese Windows environment. The CREATE or ALTER NICKNAME operation will fail if you use non-ASCII and non-Japanese characters (like Chinese).

Using the CREATE TYPE MAPPING statement with ODBC, Microsoft SQL Server, and Teradata data sources

You cannot omit a remote data type length with a data type that uses parameters. Specify a length or empty parentheses. For example, for Teradata, the default length for the CHAR data type is 1. If you want to map CHAR(1) of Teradata to VARCHAR(1) of DB2 Information Integrator, issue the following CREATE TYPE MAPPING statement:

```
CREATE TYPE MAPPING tm1 from local type varchar(1)
 to server type teradata remote type char(1);
```

You cannot use the synonym of the remote data type. You must use the short data type name. For example, in case of Teradata, the following CREATE TYPE MAPPING works correctly:

```
CREATE TYPE MAPPING tm1 from local type varchar()
 to server type teradata remote type char();
```

The following CREATE TYPE MAPPING does not work because the synonym, character, is used for the character data type. The short data type name, char, must be used.

```
CREATE TYPE MAPPING tm1 from local type varchar()
 to server type teradata remote type character();
```

DRDA wrapper support

Insert, update, and delete operations not allowed on graphic data types for UTF-8 databases

If the federated database uses the UTF-8 code page (Unicode), you cannot perform insert, update, and delete operations on GRAPHIC, VARGRAPHIC, and DBCLOB data types for DB2 Server for VM and VSE data source objects. These operations are not allowed for graphic data types because DB2 Server for VM and VSE does not contain a conversion table for UTF-16 (code page 1200). UTF-8 (Unicode) federated databases require a UTF-16 conversion table for graphic encoding.

Using the DRDA wrapper to connect to partitioned databases

When you use the DRDA wrapper to connect to a DB2 Universal Database system with remote partitions, and that remote system uses different partition numbers than the federated server uses, the connection might fail. The connection failure results in an SQL1822N error message that refers to error SQL1469N. Error SQL1469N is a remote data source error that indicates that node numbers were not specified for a particular database instance.

Before you use the DRDA wrapper to connect to a DB2 Universal Database system with remote partitions, you must ensure that the federated server is aware of the partition number that is associated with port 0 on that DB2 Universal Database system.

To determine the partition number that is associated with port 0 on a remote DB2 Universal Database system:

1. Open the sqllib/db2nodes.cfg file on the remote system.
2. Look for the partition number that is associated with port 0 on that remote system, and record that partition number.

   For example, if the host name of the remote system is nori, then look for the partition number that is associated with port 0 for host nori.
3. Close the db2nodes.cfg file.

Use one of the following three methods to notify the federated server of the partition number that is associated with port 0 on a remote DB2 Universal Database system:

- Edit the db2cli.ini file on the federated server. Set the CONNECTNODE parameter to the partition number that is associated with port 0 on the remote DB2 Universal Database system.

  This method changes the federated server environment. Using this method can potentially affect your applications.
- Set environment variable DB2NODE on the federated server to the partition number that is associated with port 0 on the remote DB2 Universal Database system. You must stop and then restart DB2 Universal Database on the federated server after you set the environment variable.

  This method changes the federated server environment. Using this method can potentially affect your applications.

- If you use SQL statements from the DB2 Command Line Processor, issue the following command on the federated server before you connect to the remote DB2 Universal Database system:

```
SET CLIENT CONNECT_DBPARTITIONNUM partition_number
```

partition_number is the partition number that is associated with port 0 on the remote DB2 Universal Database system.

This method sets the partition number only for the duration of the Command Line Processor session.

Example:

The host name of the DB2 Universal Database system to which you want to connect is palmetto. In the db2nodes.cfg file on the palmetto system, partition number 10 is associated with port 0. To notify the federated server that partition number 10 is associated with port 0 on palmetto, you can perform any one of the following actions:

- Specify the value `CONNECTNONDE=10` in the db2cli.ini file on the federated server.
- Set the DB2NODE environment variable to 10 on the federated server, stop DB2 Universal Database, and start DB2 Universal Database again. For example, on a UNIX system, use the following commands:

    db2stop

    export DB2NODE=10

    db2start

- Before you connect to palmetto, issue this command:

```
SET CLIENT CONNECT_DBPARTITIIONNUM 10
```

Nicknames on DB2 Universal Database tables that contain LONG VARCHAR FOR BIT DATA columns

When you create a nickname on a DB2 Universal Database table that contains a LONG VARCHAR FOR BIT DATA data type column, the corresponding local nickname column is automatically set to the BLOB data type. If you assign a column of data type BLOB from a nickname to a LONG VARCHAR FOR BIT DATA data type column, an SQL0408N error occurs.

To avoid this data type incompatibility, alter the local data type of the BLOB nickname column from BLOB to CLOB.

Example:

You create a nickname, MY_NICK, on a table with column LONG_COL1 of data type LONG VARCHAR FOR BIT DATA. You create a local table, MY_TABLE, with column LONG_COL2 of data type LONG VARCHAR FOR BIT DATA. Then you try to perform the following insert operation:

```
INSERT INTO my_table(long_col2) SELECT long_col1 FROM my_nick;
```

The operation fails with error SQL0408N. The local column data type of LONG_COL1 was automatically set to BLOB, which is not compatible with the data type of LONG_COL2 (LONG VARCHAR FOR BIT DATA).

To avoid the data type incompatibility, alter the local data type of the LONG_COL1 column to CLOB using the ALTER NICKNAME statement:

```
ALTER NICKNAME my_nick ALTER COLUMN long_col1 LOCAL TYPE CLOB(32770);
```

Potential error when using the DRDA wrapper for multiple local database connections on the same computer (AIX only)

If a 32-bit DB2 Universal Database server is run on an AIX system, and an application that is running on the same system has more than one federated database connection through the DRDA wrapper, then the application might get following error:

```
SQL1822N  Unexpected error code "-1224" received from data source
"W3_SERVER2". Associated text and tokens are "func="DriverConnect" msg="
SQL1224N A database agent".  SQLSTATE=560BD
```

There are two potential workarounds for this error:

- Add the entry EXTSHM=ON to the federated configuration file, instance_directory/cfg/db2dj.ini. Setting the EXTSHM variable to ON increases the amount of shared memory segments that can be used.

  When you make changes to the federated configuration file, you must stop and then restart DB2 Universal Database to activate the changes.

- Catalog the federated database on a TCP/IP node. For example:

```
CATALOG TCPIP NODE my_node REMOTE my_host SERVER 123;
CATALOG DB mydb AT NODE my_node;
CREATE WRAPPER drda;
CREATE SERVER my_server TYPE DB2/UDB VERSION 8 WRAPPER drda
    AUTHORIZATION "my_id" PASSWORD "my_pw"
    OPTIONS(ADD DBNAME 'MYDB');
```

If the error persists after you try both of the workarounds, contact IBM Support.

Documentum wrapper support

Custom function query limitation

You cannot use Documentum custom functions that do not include a column reference as an argument in queries with joins between multiple Documentum nicknames. Some examples of functions of this type are CABINET and FOLDER.

The following example shows a query that will fail:

```
SELECT D.object_name, L.sv_char1
FROM dmdoc_sr_1 D,  lsdoc_sr_1 L
WHERE DCTM.CABINET('/Cabinet1')=1 AND D.object_name = L.object_name;
```

Use Documentum custom functions that do not have a column reference argument (such as CABINET and FOLDER) only in queries that involve one Documentum nickname.

Microsoft SQL Server wrapper support

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

When you issue an insert, update, or delete statement on a Microsoft SQL Server table through a nickname, and that table has an insert, update, or delete trigger, Microsoft SQL Server will not process any further statements until the current statement has been closed.

To close the insert, update, or delete statement, issue a commit statement immediately after issuing the original insert, update, or delete statement. Issuing a commit statement after issuing the original the insert, update, or delete statement closes the original statement, and enables other statements to process.

ODBC wrapper support

CREATE TABLE restrictions

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

```
CREATE TABLE my_table(c1 INT) OPTIONS(remote_server 'MY_ODBC_SERVER',
 remote_schema 'J15USER1', invalid_option 'non option');
```

The following options are valid CREATE TABLE options:

- REMOTE_SERVER
- REMOTE_SCHEMA
- REMOTE_TABNAME
- SQL_SUFFIX

Problems after migrating wrapper user-defined reverse data type mappings

ODBC wrapper user-defined reverse data type mappings for data types that do not allow parameters (for example, the INTEGER data type) might not work correctly after migration.

If the user-defined data type mappings for data types that do not allow parameters do not work after you migrate, you must drop the problem data type mappings, and then redefine the mappings in DB2 Information Integrator Version 8.2.

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:

```
lslpp -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.

Using the Oracle wrapper with different code pages (Windows only)

This information applies to both the SQLNET and NET8 versions of the Oracle wrapper.

If you use more than one Oracle wrapper in different federated databases in the same instance at the same time, and each federated database uses a different code page, the Oracle client performs the correct code page conversion for only the first federated database that connects to an Oracle server.

When a wrapper connects to the Oracle server, the wrapper uses the code page of the federated database to determine the setting of the NLS_LANG environment variable. However, Oracle does not allow the code page portion of the NLS_LANG

variable setting to be changed after the variable is set in a process. So when wrappers in the other federated databases that have different code pages connect to the Oracle server, those wrappers use the value of NLS_LANG that the first wrapper invocation set.

Place the federated databases with different code pages in different instances so that all of the federated databases in one instance use the same code page.

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.

Recommendation: Use Sybase Adaptive Server Enterprise version 12.5.0.3 or later. Other versions of Sybase Adaptive Server Enterprise have a known array overrun problem.

Calculations on SMALLINT columns in remote tables

An SQL statement that includes a calculation on a SMALLINT column in a remote table might cause an arithmetic overflow error. To avoid this error, explicitly define the column as an INTEGER data type.

This problem occurs with both the DBLIB and CTLIB wrappers.

Run the CTLIB wrapper in trusted mode to optimize performance (UNIX systems only)

On UNIX systems, the Sybase CTLIB wrapper is not threadsafe, and should be run in trusted mode for best performance.

Update and delete operation limitation (CTLIB wrapper only)

If you create a nickname for a Sybase view, and that Sybase view was created using a Sybase subquery, you might receive error SQL1822N, "unexpected error code 7732", when you use that nickname to perform update or delete operations from the federated server.

To avoid errors when performing update or delete operations, do not use nicknames that are based upon a Sybase view that uses a Sybase subquery. Create a new nickname that directly references the table in the required Sybase view, then create a federated view that uses the new nickname. Run your update or delete operations on that federated view.

For more information about creating and using federated views, see the "Creating and using federated views" topic in the DB2 Information Center.

Using host variables in pass-through sessions (DBLIB wrapper only)

When you use statements that reference a host variable in a pass-through session of the DBLIB wrapper, you must add a trailing space to the host variable. For example, to insert data into remote table, you might issue a statement like this one:

```
INSERT INTO remote_table_name VALUES (:H0 , :H1 );
```

In the previous example, `:H0` , `:H1` are host variables to which trailing spaces have been added.

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.

Teradata wrapper support

Run the Teradata wrapper in trusted mode to optimize performance (UNIX systems only)

On UNIX systems, the Teradata wrapper is not threadsafe, and should be run in trusted mode for best performance.

UPDATE or DELETE operation errors on nicknames

By default, rows are not uniquely identified on Teradata data source tables. You might receive an SQL30090N, RC="21" error when you try to update or delete a nickname that is associated with a Teradata table or a Teradata view.

If the SQL30090N, RC="21" error occurs, create at least one unique index on the Teradata table that is being updated or deleted, and then try the operation again.

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

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 for 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 SQL3324N.

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.

WebSphere Business Integration wrapper support

New nickname option for detecting business object not found conditions

The WebSphere Business Integration wrapper is including a nickname option for detecting business object not found conditions. Because SAP and PeopleSoft APIs do not follow a standardized error reporting model, it is possible for each API to return a different message for the same error condition. Specifically, SAP and PeopleSoft business application APIs do not issue consistent error messages for a record not found condition.

For example, when a record for a given identifier cannot be located in the SAP application, the getdetail2 BAPI for the CUSTOMER business object might return error code 502, while the getdetail BAPI for the COMPANY business object might return a different error code.

Because of the inconsistent error messages, the WebSphere Business Integration adapters for mySAP.com and PeopleSoft cannot map these errors to the specific BO_NOT_FOUND status flag, which is in the response business object that is returned to the wrapper. The lack of information in the status flag can cause certain types of DB2 Information Integrator application queries to fail, even if a single record is missing or cannot be located in the application data source. The application failure might not be the desired behavior for the DB2 Information Integrator application.

The following example is a join of a WebSphere Business Integration SAP nickname that represents the Customer.get_detail2 BAPI with a local table that contains the customer identifiers:

```
Select name from sap_bapi_customer_getdetail2_NN a,
 local_table b where b.customerid=a.customerno;
```

If the local_table contains a single customer ID that does not exist in the SAP application, the query fails.

As a workaround for this problem, the WebSphere Business Integration wrapper provides an additional nickname option that allows users to specify specific error codes that should be processed as a record not found rather than an error for the nickname. The nickname option is: BO_NOT_FOUND_CODES. The value is a comma delimited list of business application error codes. The following example shows the new nickname option:

```
CREATE NICKNAME sap_bapi_customer_getdetail2_NN
   (
    CUSTOMER VARCHAR(10) OPTIONS
        (XPATH './ns3:sap_customeraddress/
            ns1:sap_customeraddress/ns1:CUSTOMER/text()'),
    ...
    FOR SERVER
        sap_server
    OPTIONS(XPATH '//ns3:sap_bapi_customer_getdetail2',
            ...
            BO_NOT_FOUND_CODES '502,503,501'
            ....
            );
```

DB2 Control Center

db2updv8 command: updating a federated database to use the DB2 Control Center or utility stored procedures like nickname statistics (SYSPROC.NNSTAT)

If you are using a federated database with DB2 Information Integrator, Version 8.2 that was created using DB2 Universal Database, Version 8.1 or later, you must upgrade your federated database to use the DB2 Control Center or utility stored procedures, such as nickname statistics (SYSPROC.NNSTAT), with DB2 Information Integrator, Version 8.2. To update your federated database, use the **db2updv8** command.

An example of this command is:

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

For more information about using this command, see the "db2updv8 - Update Database to Version 8 Current Level Command" topic in the DB2 Information Center.

Federated client and server compatibility

If you use DB2 Universal Database, Version 7 or Version 8.1 clients with a DB2 Information Integrator, Version 8.2 server, you must upgrade your clients to DB2 Information Integrator, Version 8.2 to use the new federated function in the DB2 Control Center.

If you use DB2 Information Integrator, Version 8.2 clients and a DB2 Universal Database Version 7 or Version 8.1 server, you must upgrade your server to DB2 Information Integrator, Version 8.2 to use the new federated function in the DB2 Control Center.

Blast data sources

If the BlastSeq column for a BLAST nickname is defined as a CLOB data type, you must create sequence match function templates before you can query BLAST data sources. The SQL statements that you issue on the federated database to create the function templates are documented in "Fix Pack 8 Documentation updates" on page 49.

HMMER data sources

If the HMMQSEQ column for a HMMER nickname is defined as a CLOB data type, you must create sequence match function templates before you can query HMMER data sources. The SQL statements that you issue on the federated database to create the function templates are documented in "Fix Pack 8 Documentation updates" on page 49.

Web services

The sample deployment descriptor file (dds-example.xml) is not available in the apache-services.war file

Deployment descriptor file
The sample deployment descriptor file (dds-example.xml) is not available in the apache-services.war file. If you are using the Apache SOAP engine, use the default SOAP engine configuration manager. Do not create the soap.xml file, or if the file already exists, delete the soap.xml file in the root directory of your Web application. You can also omit the deployment descriptor generation (Dadx2Dd) steps.

Dadx2Dd parameters -c and -u do not work
The parameters that were added in fix pack 8, in a topic named "Generating deployment descriptors", included -c and -u in the Dadx2Dd deployment descriptor class. These parameters do not work as described. If you are using the Apache SOAP engine, use the default SOAP engine configuraiton manager. Do not create the soap.xml file, or if the file already exists, delete the soap.xml file inthe root directory of your Web application. You can also omit the deployment descriptor generation (Dadx2Dd) steps.

The Web services wrapper does not support HTTPS/SSL

The Web services wrapper supports HTTPS/SSL encryption on Linux 2.6 and later. The Web services wrapper does not support HTTPS/SSL encryption on Linux 2.4. If a URL that starts with https:// is used for a nickname, the wrapper will return an error from the SELECT statement on that nickname when you are using Linux 2.4.

WebSphere MQ

WebSphere MQ 5.3 configuration Issues for Solaris

After you install WebSphere MQ Version 5.3 with corrective service diskette 05 (CSD05) on a Solaris Operating Environment system, open a command prompt, and type:ulimit -n 1024.

This command sets the limit for the number of files that are opened and will allow you to create a queue manager in the Solaris Operating Environment.

Wrapper development

Workaround for ClassNotFoundException error in Java™ wrappers that use J2EE Application Programming Interfaces

A Java™ wrapper might fail to load a valid class that is specified in the CLASSPATH. The error is thrown when the wrapper tries to load the initial class context.

For example, the following exception information was logged from a wrapper that attempted to create a new javax.naming.InitialContext instance:

```
Exception :javax.naming.NoInitialContextException: Cannot instantiate class:
org.jnp.interfaces.NamingContextFactory
[Root exception is java.lang.ClassNotFoundException:
org/jnp/interfaces/NamingContextFactory]
javax.naming.NoInitialContextException:
Cannot instantiate class: org.jnp.interfaces.NamingContextFactory.
Root exception is java.lang.ClassNotFoundException:
 org/jnp/interfaces/NamingContextFactory
   at java.lang.Class.forName0(Native Method)
   at java.lang.Class.forName(Class.java:256)
   at com.sun.naming.internal.VersionHelper12.loadClass(VersionHelper12.java:59)
   at javax.naming.spi.NamingManager.getInitialContext(NamingManager.java:661)
   at javax.naming.InitialContext.getDefaultInitCtx(InitialContext.java:256)
   at javax.naming.InitialContext.init(InitialContext.java:232)
   at javax.naming.InitialContext. <init> (InitialContext.java:208)
```

In the example, the org.jnp.interfaces.NamingContextFactory class is included in the CLASSPATH, but the Java Virtual Machine fails to load the class. This error occurs because the com.sun.naming.internal.VersionHelper12 class loads the application-specific InitialContext subclass by using the class loader of the current thread, as returned by the Thread.currentThread().getContextClassLoader() method. However, the value that the method returned is null.

A potential workaround for class load failures that have a similar structure to the example error is to explicitly set the class loader for the current thread as a valid non-null ClassLoader instance, and to add the following code in the wrapper. Insert the following code immediately before the call that caused the ClassNotfoundException to be thrown:

```
final ClassLoader myClassLoader = this.getClass().getClassLoader();
ClassLoader threadClassLoader = (ClassLoader) AccessController.doPrivileged(new PrivilegedAction()
    {
      public Object run()
      {
        return Thread.currentThread().getContextClassLoader();
      }
    }
    );
```

```
if( myClassLoader != null && threadClassLoader == null)
{
  AccessController.doPrivileged(new PrivilegedAction()
     {
       public Object run()
       {
         Thread.currentThread().setContextClassLoader(myClassLoader);
         return null;
       }
     }
     );
}
```

The class to which this code is added might need to import the package java.security.* to successfully compile.

XML Metadata Registry

Required TCP/IP ports to enable Web browser access to the XML Metadata Registry

To allow Web browsers to access the XML Metadata Registry, you must enable unrestricted access to the following TCP/IP ports on the computer where the XML Metadata Registry is deployed:

- 20000 (for http)
- 20010 (for htpps)
- 20020 (for the XSLTLoader utility)

Reset Password utility

The XML Metadata Registry administrator account password on the application server for DB2 is used when you deploy the XML Metadata Registry. If you change the XML Metadata Registry administrator account password on the application server for DB2, you must also immediately register that password change in the XML Metadata Registry program. If you do not register a password change, the XML Metadata Registry program will not function correctly.

To register changes to the XML Metadata Registry administrator account password on the application server for DB2 in the XML Metadata Registry, use the Reset Password utility. The Reset Password utility is in the uil directory.

You must stop the Enterprise Application Server before you run the Reset Password utility. You can usually stop the Enterprise Application Server by using the old password. If you cannot stop the Enterprise Application Server with the old password, you can safely terminate the Java process that is associated with the application server for DB2.

To run the Reset Password utility, issue the following command:

resetPassword -u *<userid>* -o *<oldpassword>* -n *<newpassword>*

- *<userid>* is the XML Metadata Registry administrator account name.
- *<oldpassword>* is the old XML Metadata Registry administrator account password on the application server for DB2.
- *<newpassword>* is the new XML Metadata Registry administrator account password on the application server for DB2.

To help recover in case of failure, the Reset Password utility creates a copy of the security.xml configuration file in the directory where the application server for DB2 is installed.

Scripts to start and stop the XML Metadata Registry

You can now stop and start the XML Metadata Registry after it is deployed without stopping the application server for DB2, and any other application server for DB2 programs. The scripts (for UNIX and Linux™ systems) or batch files (for Windows systems) that enable you to perform these stop and start operations are in the *<eas_install_dir>*/xmr/bin directory. *<eas_install_dir>* is the directory where the application server for DB2 is installed.

The application server for DB2 must be running to use the scripts or batch files.

To stop the XML Metadata Registry, enter the following command:

```
stopXMR -user <userID> -password <password>
```

To start the XML Metadata Registry, enter the following command:

```
startXMR -user <userID> -password <password>
```

<userID> and *<password>* is the same user ID and password combination that is used for the **deployXMR** command (the XML Metadata Registry administrator ID and password combination).

Large document storage

To prevent errors when large documents are stored in the XML Metadata Registry, you can use the Jython wsadmin script nlc.py. You can also use this script to tune the performance of the XML Metadata Registry program in general.

The nlc.py script changes the noLocalCopies (pass by reference) setting of the Object Request Broker (ORB) that the XML Metadata Registry uses. The script changes the noLocalCopies current setting (true or false) to its opposite setting. The default setting for noLocalCopies when the XML Metadata Registry is installed is false. Changing noLocalCopies to true can eliminate errors, and can improve performance when large documents are stored in the XML Metadata Registry.

The nlc.py script is in the util directory. To run the nlc.py script on UNIX systems, issue the following command:

```
<application server for DB2 installation directory>/bin/wsadmin
-user <xmradmin>
-password <password>
-lang jython
-f <application server for DB2 installation directory>/xmr/util/nlc.py
```

To run the nlc.py script on Windows systems, issue the following command:

```
<application server for DB2 installation directory>\bin\wsadmin
-user <xmradmin>
-password <password>
-lang jython
-f <application server for DB2 installation directory>\xmr\util\nlc.py
```

<xmradmin> is the XML Metadata Registry administrator account name.

<password> is the XML Metadata Registry administrator account password.

After you run the nlc.py script, you must restart the application server for DB2 to activate the changes.

Patch required when using Microsoft Internet Explorer browser

Anyone who accesses the XML Metadata Registry by using the Internet Explorer browser must download and apply Microsoft critical update patch Q831167.

The easiest way to notify XML Metadata Registry users about this requirement is to modify the "About This Registry" message that displays on the XML Metadata Registry Log In page.

The XML Metadata Registry administrator can change the "About This Registry" message text from the Welcome page. The Welcome page appears after you successfully log in to the XML Metadata Registry. To change the "About This Registry" message text, go to the appropriate area of the Welcome page, enter the text changes in HTML, then click **Save**. The changed text appears on the XML Metadata Registry Log In page for all users.

The following example shows sample HTML that you can use to document the Microsoft patch requirement in the "About This Registry" message.

```
If you encounter the following problems when you attempt to register
or log in using Microsoft Internet Explorer, you need to download
and install a patch from Microsoft.<br>
1. You receive an invalid user ID or password error when you use a
valid user ID and password to log in.<br>
2. The Registration page indicates that required fields are empty
when they are not.<br>
See the Microsoft Knowledge Base Article 831167 at
<a href "http://support.microsoft.com/default.aspx?kbid=831167">
http://support.microsoft.com/default.aspx?kbid=831167</a> for
instructions on downloading and installing the patch.
```

Changing the Administrator ID name value

The XML Metadata Registry administrator user account is given the name "XMR Administrator" by default when the XML Metadata Registry is deployed. The administrator user account name, not the administrator user ID value, appears in the XML Metadata Registry user interface.

To change the name of the XML Metadata Registry administrator account:
1. In the XML Metadata Registry program, browse through Business Objects.
2. Select the XML Metadata Registry administrator user name.
3. Click the user name to display the user information.
4. On the Basics page, change the value in the Name field.
5. Click Save to store the changes.

Do not use the back function in a Web browser

Do not use the back function (including the Back button) of your Web browser with the XML Metadata Registry interface. Using the Back function can cause problems in the XML Metadata Registry interface.

Double byte character strings do not work in Mozilla Web browsers

Web browsers that use Mozilla file upload controls do not properly handle file names that contain double-byte character set (DBCS) and multi-byte character set (MBCS) characters. Any attempt to upload files with names that contain DBCS or MBCS characters in Mozilla Web browsers will not work.

To avoid this problem, you can either rename the file that you are trying to upload, or use a Web browser that does not use the Mozilla file upload controls (for example, Microsoft Internet Explorer).

HTML limitation affects navigation tree display

XML Metadata Registry objects can contain consecutive white space characters (for example, spaces and tabs). However, because of HTML constraints, the object names as displayed in the XML Metadata Registry navigation tree show any consecutive white space in the name as a single space character.

Invalid HTML markup in About This Registry message causes display problems

The "About This Registry" message text is shown on the XML Metadata Registry Log In page. You can customize the "About This Registry" message by adding HTML markup as needed.

However, if invalid HTML is used (for example, a tag is missing an end tag), the Log In window might be displayed incorrectly. For example, it might be difficult to identify the Username and Password text boxes and controls.

Fix Pack 9 Documentation updates

Corrections to documentation

The following information describes corrections to documentation for fix pack 9.

All of the navigational references to the DB2 Information Center in this section assume that you are using the DB2 Information Center navigation tree as the starting point.

Product overviews

Topic title: Supported data sources

Location in the DB2 Information Center:
Product overviews —> Information Integration —> DB2 Information Integrator overview —> Federated systems overview

Addition 1 of 7
For DB2 Universal Database for iSeries, 5.3 is a supported version.

Addition 2 of 7
For Microsoft SQL Server: If you use the Microsoft SQL Server wrapper with a UTF-8 database on UNIX you must use DataDirect Connect for ODBC 4.2 Service Pack 2 or later.

Addition 3 of 7
For Informix data sources:
- On Solaris, the Informix client version 2.81.xC2 is not supported. If you are using the Informix client version 2.81.xC2, update the client to version 2.81.xC2R1 or later.
- On Windows, the Informix SDK client version 2.81.TC2 or later is required.
- If you use Informix Dynamic Server 9.3 as a data source, you must use Informix Dynamic Server 9.30.xC4 or later.
- On the 64-bit mode zLinux operating system, you must use the Informix client version 2.81.FC3.

Addition 4 of 7
For Teradata data sources: On Windows, you must use the Teradata client TTU 7.0 or later and the Teradata API library CLIv2 4.7.0 or later on the federated server.

Addition 5 of 7
For Oracle data sources: Oracle 8.0.6 and 8.1.6 are not supported versions.

Addition 6 of 7
For BioRS and Entrez, the access method should be HTTP.

Addition 7 of 7
For BioRS the supported versions are: 5.0.14, 5.2.x.x.

Installing

Topic title: Web services consumer functions

Location in the DB2 Information Center:

Installing —> Information Integration —> Installing DB2
Information Integrator on Windows and UNIX —> Installing Web
services components —> Installing the Web services consumer

Addition

To use HTTP proxies, you must set two environment variables for
configuring to DB2 Universal Database:

- Set the DB2SOAP_PROXY variable to include the host name of
  the computer with the HTTP proxy.
- Set the DB2SOAP_PORT to the port of the HTTP proxy, such as
  8080.

The SOAP traffic goes through the system that tunnels the firewall.

You can set the DB2SOAP_TIMEOUT environment variable to
specify a maximum number of minutes to wait for both a network
transfer and the computation of a result. When you use a proxy
server, make sure that the TIMEOUT setting of the proxy server is
longer than the TIMEOUT value of the Web services wrapper
SERVER OPTION or the DB2SOAP_TIMEOUT value. Otherwise,
the proxy might report a timeout to Information Integrator even
though Information Integrator would have waited longer for a
result.

Configuring

Topic title: Registering nicknames for SQL Server tables and views

Location in the DB2 Information Center

Configuring –> Federated systems and data sources –>
Configuring data sources –> Configuring Microsoft SQL Server
data sources –>Adding SQL Server data sources to a federated
system

Correction

The following information should be included in the introduction
to registering nicknames.

To ensure that the federated database has current and complete
statistics, run the SQL Server sp_createstats stored procedure and
the SQL Server CREATE STATISTICS command from the SQL
Server database before creating the nickname.

The sp_createstats stored procedure gathers statistics on all default
columns in a table in an SQL Server data source, but does not
gather statistics for columns that appear first within an index. To
ensure that the federated database has complete statistics on the
SQL Server table, you also must use the SQL Server CREATE
STATISTICS command to gather statistics for each column that
appears first in an index.

When you use the CREATE STATISTICS command from the SQL
Server database, you must give the statistic the same name as the
column on which the statistics are being collected. By giving the
statistic the same name as the column, you ensure that when you
register the nickname with the CREATE NICKNAME statement,
the federated database reads the statistics that are collected by the
SQL Server CREATE STATISTICS command.

Topic title: Configuring the WebSphere Business Integration Adapters

Location in the DB2 Information Center:
Configuring –> Federated systems and data sources –>
Configuring data sources –> Configuring business application data
sources

Correction
In step 2 in the subsection named "Standard configuration
properties", the third list item incorrectly defines the type of
delivery transport as WMQI-MQ. You should instead specify the
type of delivery transport as JMS.

Topic title: Registering the server definition for business application data sources

Location in the DB2 Information Center:
Configuring –> Federated systems and data sources –>
Configuring data sources –> Configuring business application data
sources –> Adding business applications to a federated system

Correction
The example in the CREATE SERVER statement should contain
TYPE wbi:

```
CREATE SERVER siebel_server
  TYPE wbi
  VERSION 2.4
  WRAPPER wbi_wrapper
  OPTIONS ( App_Type 'siebel',
     Request_Queue 'myqueue3',
     Response_Queue 'myqueue4',
     Fault_Queue 'myqueue5',
     MQ_Manager 'mymq'
     MQ_REPONSE_TIMEOUT '55000',
     MQ_CONN_NAME '9.30.76.151(1420)',
     MQ_SVRCONN_CHANNELNAME 'SYSTEM.DEF.SVRCONN'
     )
```

Topic title: Fixed columns for BLAST nicknames

Location in the DB2 Information Center:
Configuring –> Federated systems and data sources –>
Configuring data sources –> Configuring BLAST data sources –>
Adding BLAST to a federated system –> Registering nicknames for
BLAST data sources

Correction 1 of 3
There are six new input columns for BLAST nicknames. You must
drop any existing nicknames and create the nicknames again before
you can use the new input columns.

Table 3. New fixed input columns for BLAST nicknames

Name	Data type	Operators	Description
Mask_Lower_Case	CHAR(1)	=	Use lowercase filtering with a FASTA sequence.
Query_Genetic_Code	INTEGER	=	Query genetic code uses default = 1.
DB_Genetic_Code	INTEGER	=	DB genetic code for tblastn and tblastx queries uses default = 1.

Table 3. New fixed input columns for BLAST nicknames (continued)

Name	Data type	Operators	Description
XDropoff_Ungapped	DOUBLE	=	The X dropoff value for ungapped extension, measured in bits. The value 0.0 invokes the default behavior. For blastn queries, the default is 20 bits. For megablast queries, the default is 10 bits. For all other query types, the default is 7 bits (REAL data types).
XDropoff_Gapped	INTEGER	=	The X dropoff value for gapped alignment, measured in bits. The value 0.0 invokes the default behavior. For blastn queries, the default is 30 bits. For megablast queries, the default is 20 bits. For all other query types, the default is 15 bits (INTEGER data types).
Final_XDropoff	INTEGER	=	The X dropoff value for the final gapped alignment, measured in bits. The value 0.0 invokes the default behavior. For blastn and megablast queries, the default is 50 bits. For tblastx queries, the default is 0 bits. For all other query types the default is 25 bits (INTEGER data types).

Correction 2 of 3

There are new BLAST switches for the new input columns.

Table 4. BLAST search types and switches supported by the new input fixed columns

Name	BLAST search types	BLAST switch	Required	Default
Mask_Lower_Case	n, p, x, tn, tx	-U	No	F
Query_Genetic_Code	n, p, x, tn, tx	-Q	No	1
DB_Genetic_Code	tn, tx	-D	No	1
XDropoff_Ungapped	n, p, x, tn, tx	-y	No	0.0
XDropoff_Gapped	n, p, x, tn, tx	-X	No	0
Final_XDropoff	n, p, x, tn, tx	-Z	No	0

Correction 3 of 3

There is one new fixed output column for BLAST nicknames.

Table 5. New fixed output columns for BLAST nicknames

Name	Data type	Description
Hsp_Score	DOUBLE	Use lowercase filtering with a FASTA sequence.

Topic title: Writing BLAST SQL queries

Location in the DB2 Information Center:

Configuring –> Federated systems and data sources –> Configuring data sources –> Configuring BLAST data sources –> Adding BLAST to a federated system

Correction 1 of 3

Add the following text to the Restrictions section in this topic:

If the BlastSeq column for a BLAST nickname is defined as a CLOB data type, you must create sequence match function templates before you can query BLAST data sources. The SQL statements that you issue on the federated database to create the function templates are:

```
CREATE FUNCTION lblast.sequence_match (CLOB, CLOB)
    RETURNS INTEGER AS TEMPLATE DETERMINISTIC NO EXTERNAL ACTION;
CREATE FUNCTION lblast.sequence_match (CLOB, VARCHAR())
    RETURNS INTEGER AS TEMPLATE DETERMINISTIC NO EXTERNAL ACTION;
CREATE FUNCTION lblast.sequence_match (VARCHAR(), CLOB)
    RETURNS INTEGER AS TEMPLATE DETERMINISTIC NO EXTERNAL ACTION;
CREATE FUNCTION lblast.sequence_match (VARCHAR(), VARCHAR())
    RETURNS INTEGER AS TEMPLATE DETERMINISTIC NO EXTERNAL ACTION;
```

Correction 2 of 3

You can use the SUBSTR scalar function in the query predicate. The SUBSTR function returns part of a string.

For example:

```
SELECT BlastSeq FROM t1, t2
    WHERE BlastSeq = SUBSTR (t2.sequence, 15, 300)
```

Correction 3 of 3

You can use the SUBSTR function in a SEQUENCE_MATCH function.

For example:

```
SELECT BlastSeq FROM t1, t2
    WHERE LSBlast.SEQUENCE_MATCH
    (BlastSeq,SUBSTR(t2.sequence, 15, 300)) = 1
```

Topic title: CREATE SERVER statement options - BioRS wrapper

Location in the DB2 Information Center

Configuring –> Federated systems and data sources –> Configuring data sources –>

Correction

The description for the VERSION parameter has changed to reflect the support for BioRS server version 5.2. The updated description is:

VERSION The version of the BioRS server that you want to access. The supported BioRS versions are 5.0.14 and 5.2. If you are accessing a BioRS server that is version 5.2, you must specify 5.2 as the value for the VERSION parameter. You do not need to specify this option if you are using version 5.0.14. The default value of 1.0, which equates to version 5.0.14, is used for this parameter if you do not specify the value.

Topic title: Registering nicknames for XML data sources

Location in the DB2 Information Center

Configuring –> Federated systems and data sources –> Configuring data sources –> Configuring XML data sources –> Adding XML to a federated system

Correction: Restriction with Windows 2003 federated servers

If you try to access XML data sources that are on a shared drive from a federated server that runs Windows 2003, your query might fail with the following error message:

```
SQL1822N  Unexpected error code "ERRNO = 2"
received from data source "XML_SERVER".
Associated text and tokens are
"Unable to read file".
SQLSTATE=560BD
```

This is a limitation of Windows 2003. You can avoid this problem by specifying the absolute path in the FILE_PATH or DIRECTORY_PATH option in the CREATE NICKNAME statement.

The following example shows a CREATE NICKNAME statement with an abbreviated path that is specified in the FILE_PATH option (X:\ is the drive that maps to the remote computer).

```
CREATE NICKNAME customers
(
id VARCHAR(5)  OPTIONS(XPATH '.\@id'),
name VARCHAR(16) OPTIONS(XPATH '.\name'),
address VARCHAR(30) OPTIONS
(XPATH '.\address\@street'),

 FOR SERVER xml_server
   OPTIONS(FILE_PATH 'X:\customers.xml',
     XPATH '\doc\customer');
```

Queries that use this nickname might fail because you specified the abbreviated path.

For a federated server that runs Windows 2003, specify the absolute path in the FILE_PATH or DIRECTORY_PATH option in the CREATE NICKNAME statement.

For example:

```
CREATE NICKNAME customers
(
id      VARCHAR(5)  OPTIONS(XPATH '.\@id'),
name    VARCHAR(16) OPTIONS(XPATH '.\name'),
address VARCHAR(30) OPTIONS(XPATH '.\address\@street'),
 FOR SERVER xml_server
  OPTIONS(FILE_PATH '\\host.svl.ibm.com\D$\customers.xml',
    XPATH '\doc\customer');
```

Topic title: Registering nicknames for table-structured files

Location in the DB2 Information Center

Configuring –> Federated systems and data sources –> Configuring data sources –> Configuring table-structured file data sources

Correction: Restriction with Windows 2003 federated servers

If you attempt to access a table-structured file data sources that are on a shared drive from federated server that runs Windows 2003, your query might fail with the following error message:

```
SQL1822N  Unexpected error code "ERRNO = 2"
received from data source "SERVERNAME1".
Associated text and tokens are
 "Unable to read file".
SQLSTATE=560BD
```

This is a limitation of Windows 2003. You can avoid this problem by specifying the absolute path in the FILE_PATH option in the CREATE NICKNAME statement.

The following example shows a CREATE NICKNAME statement with an abbreviated path specified in the FILE_PATH option:

```
CREATE NICKNAME nickname
(COL1 CHARACTER (10) NOT NULL)
    FOR SERVER servername1
    OPTIONS (FILE_PATH '\textfile1.txt');
```

Queries that use this nickname might fail because you specified the abbreviated path.

For a federated server that runs Windows 2003, specify the absolute path in the FILE_PATH option in the CREATE NICKNAME statement.

For example:

```
CREATE NICKNAME nickname
(COL1 CHARACTER (10) NOT NULL)
    FOR SERVER servername1
    OPTIONS (FILE_PATH
'\\host.svl.ibm.com\D$\textfile1.txt');
```

Topic title: Manually linking the wrapper libraries to the data source client software

Location in the DB2 Information Center:
Configuring —> Federated systems and data sources —>Configuring a federated system —> Checking the setup of the federated server —>Confirming the link-edit to the data source client libraries (UNIX)

Correction
The first bullet in step 2, which specifies the Typical installation option, is incorrect. The correct text is:

- For Informix data sources, run the DB2 server installation again and specify the Custom installation option.

Topic title: Registering nicknames for Teradata tables and views

Location in the DB2 Information Center:
Configuring –> Federated systems and data sources –> Configuring data sources –> Configuring Teradata data sources –>Adding Teradata data sources to a federated system

Correction

The following information should be included in the introduction to registering nicknames.

To ensure that the federated database has complete statistics on Teradata tables, use the COLLECT STATISTICS Teradata command.

From the Teradata server, use the COLLECT STATISTICS Teradata command to collect statistics on all columns and indexes in a table before creating the nickname.

When you register the nickname with the CREATE NICKNAME statement, the federated database reads the statistics from the Teradata system catalog and updates the local statistics for the nickname.

Topic title: Registering nicknames for Teradata tables and views

Location in the DB2 Information Center:
Configuring –> Federated systems and data sources –> Configuring data sources –> Configuring Teradata data sources –>Adding Teradata data sources to a federated system

Correction

The following information should be included in the introduction to registering nicknames:

When you register a nickname on a Teradata view, the federated database recognizes all columns of the view as nullable, even if the columns in the Teradata view do not allow null values. There is no workaround for this limitation.

Topic title: Adding Teradata data sources to a federated system

Location in the DB2 Information Center:
Configuring –> Federated systems and data sources –> Configuring data sources –> Configuring Teradata data sources

Correction: Restriction with Windows 2003 federated servers.

Teradata data sources cannot be added to a federated server that runs on Windows 2003 because the supported Teradata client does not support Windows 2003. If you try to install the Teradata client (CLIv2 4.7) on Windows 2003, the installation might fail with the following error message:

```
Teradata CLIv2 does not support your operating system.
Installation will be aborted.
```

There is no workaround for this restriction.

Administering

Topic title: Default forward data type mappings

Location in the DB2 Information Center:
Administering —> Federated systems —> Mapping data types —> Data type mappings in a federated system —> Forward and reverse data type mappings

Correction 1 of 1

Table: Oracle NET8 forward default data type mappings
Add a table row for the BIGINT data type mapping:

Table 6. Oracle NET8 forward default data type mappings

NUMBER	10	18	0	0	-	\0	BIGINT	0	0	N

Topic title: Default reverse data type mappings

Location in the DB2 Information Center:
Administering —> Federated systems —> Mapping data types —>
Data type mappings in a federated system —> Forward and
reverse data type mappings

Correction 1 of 1

Table: Oracle NET8 reverse default data type mappings
Make the following corrections and additions to this table:

- In the table row for the SMALLINT data type, change
  the REMOTE_LENGTH value from 4 to 5.
- In the table row for the INTEGER data type, change the
  REMOTE_LENGTH value from 9 to 10.
- Add a table row for the BIGINT data type mapping:

Table 7. Oracle NET8 reverse default data type mappings

BIGINT	0	8	0	0	N	\0	NUMBER	19	0	N

Developing

**Topic title: Installing or migrating WORF to work with WebSphere Application
Server Version 5 or later for Windows and UNIX**

Location in the DB2 Information Center:
Developing —> Web services and federated Web applications —>
Installing the Web services provider —> Installing or migrating
WORF —> For Windows and UNIX

Addition

If you create a Web services application in WebSphere Studio
Application Developer Version 5, the Web applications that are
generated are not compatible with the DB2 Information Integrator
Version 8.2 Web services engine. You will not be able to deploy
that Web application into an application server that uses DB2
Information Integrator Version 8.2 Web services provider. Also, if
you manually create a Web application that uses the DB2
Information Integrator Version 8.2 Web services engine, and then
import that application into WebSphere Studio Application
Developer Version 5, you cannot run that Web application.

You must manually update a Web application that is generated
with WebSphere Studio Version 5 to deploy the application in an
application server that has DB2 Information Integrator Web
Services Provider Version 8.2.

To update the Web application:

1. Optional: Copy the Web application into a file system to
   modify the files in the application. Use the following command
   from the command line:

   ```
   jar xf myWebapp.war
   ```

2. Replace the worf-servlets.jar file in the \WEB-INF\lib directory
   of the Web application with the worf-servlets.jar file that is in
   the dxxworf.zip file.

3. If you are using the custom configuration manager that is used
   by WebSphere Application Server 5.1, perform one of the
   following actions:

- Generate the deployment descriptors by using the default Apache SOAP configuration manager. The Web services provider automatically redeploys all of the Web services when you access them.
- Manually regenerate the Apache SOAP deployment descriptors. Modify the dds.xml file by regenerating the deployment descriptors:

    a. Issue the following command:

    ```
    com.ibm.etools.webservice.rt.dadx2dd.Dadx2Dd
       -r list.dadx
       -p /db2sample
       -n /dxx_sample
       -i "classes\groups\dxx_sample\list.dadx"
       -o "classes\dds\dxx_sample\list.isd"
    ```

    You must precede the group name and the directory name with a slash.

    b. Add or replace the generated deployment descriptor (list.isd) to the dds.xml file.

    The dds.xml file contains all of the Apache SOAP deployment descriptors for the current Web application.

    Note: If a soap.xml file exists in the root of the Web application, and it contains the com.ibm.soap.server.XMLDrivenConfigManager specification (this driver is the custom configuration manager that is used by WebSphere Application Server 5.1 when using Apache SOAP), then you are using the custom configuration manager.

4. For each application that you deployed, replace the JavaServer Pages in the worf directory of that application, with the files in the worf directory of the apache-services.war or the axis-services.war.

5. Create the WAR file again with the following command:

    ```
    jar cf myWebapp.war *
    ```

6. Deploy the application again.

Topic title: Apache SOAP configurations

Location in the DB2 Information Center:
Developing —> Web services and federated Web applications —> Deploying and testing Web applications

axis–services.war file
The axis-services.war file does not contain the dds-example.xml file. Only the apache-services.war file contains the dds-example.xml file.

Topic title: Starting and stopping the application server for DB2 Information Integrator

Location in the DB2 Information Center:
Developing —> Web services and federated Web applications —> Deploying and testing Web applications

Addition 1 of 2

To start the application server for DB2, you must use server1 as the server name.

```
startServer server1
stopServer  server1
```

Addition 2 of 2

You must define port number 20000 for the Web services that are running under the application server for DB2. Invoke the WORF samples with the following URL:

```
http://localhost:20000/services
```

In this example, `services` is the context root that you specify when you install the Web services.

Reference

Topic title: Server options for federated systems

Location in the DB2 Information Center:

Reference —> Configuration parameters —> Federated systems and data sources

Correction

Table: Server options for nonrelational data sources.
In the Web services row, X has been added to the TIMEOUT column.

Table: Server options and their settings
In the TIMEOUT row, the following column has been added:

Table 8.

TIMEOUT	Web services: Specifies the time, in minutes, that DB2 should wait for a network transfer and the computation of a result.

New documentation

This section provides new information for this release.

Updates to the KEGG user-defined function information

This section summarizes the key KEGG user-defined function information that was updated in the DB2 Information Center for this release.

New KEGG user-defined functions

The new KEGG user-defined functions are:

BtitS The BtitS function is a scalar function that returns a semicolon-delimited list of definitions. The list is a character string with a data type of VARCHAR and an actual length that is not greater than 32767 bytes.

The maximum number of gene identifiers that you can specify for the entrylist argument is 100.

For example, to retrieve the definitions for the M. loti, P. abyssi, and X. fastidiosa genes, the clause that you use is:

```
VALUES CAST(DB2LS.BtitS ('mlo:mll1558','pab:PAB1288','xfa:XF2434' ))
    AS VARCHAR(1000));
```

MotifsByGeneS

The MotifsByGeneS function is a scalar function that returns a space-delimited list of motif identifiers. The list is a character string with a data type of VARCHAR and an actual length that is not greater than 32767 bytes.

You can search for the motifs in a single database or in all of the valid databases. To search all of the valid databases, you specify `all` in the database argument of the function.

The databases that you can specify are:

Database name	Abbreviation
Pfam	pfam
TIGRFAM	tfam
PROSITE pattern	pspt
PROSITE profile	pspf

To search for all of the motifs for the E. coli gene in the PROSITE profile database, the clause that you use is:

```
VALUES CAST(DB2LS.MotifsByGeneS ('eco:b0002','pspf')
    AS VARCHAR(1000));
```

The KEGG code for the E. coli gene is *eco* and the gene name is *b0002*. The KEGG name for the PROSITE profile database is *pspf*.

To search for all of the motifs for the E. coli gene in all four of the valid databases, the clause that you use is:

```
VALUES CAST(DB2LS.MotifsByGeneS ('eco:b0002','all')
    AS VARCHAR(1000));
```

The KEGG code for the E. coli gene is *eco* and the gene name is *b0002*.

MotifsByGeneT

The MotifsByGenesT function is a table function that returns the motif row. You can search for the motifs in a single database or in all of the valid databases. To search all of the valid databases, you specify `all` in the database argument of the function.

The databases that you can specify are:

Database name	Abbreviation
Pfam	pfam
TIGRFAM	tfam
PROSITE pattern	pspt
PROSITE profile	pspf

To search for all of the motifs for the E. coli gene in the PROSITE profile database, the SELECT statement that you use is:

```
SELECT * FROM TABLE(DB2LS.MotifsByGeneT
    ('eco:b0002','pfam')
    AS t;
```

The KEGG code for the E. coli gene is *eco* and the gene name is *b0002*. The KEGG name for the Pfam profile database is *pfam*.

To search for all of the motifs for the E. coli gene in all of the valid
databases, the SELECT statement that you use is:

```
SELECT * FROM TABLE(DB2LS.MotifsByGeneT
    ('eco:b0002','all')
    AS t;
```

The KEGG code for the E. coli gene is *eco* and the gene name is *b0002*.

Removed KEGG user-defined functions

The following KEGG user-defined functions were removed from DB2 Information
Integrator. Use the alternative functions that are shown in the table.

Table 9. KEGG user-defined functions that are no longer supported

Function name	Alternative function
BstBstHmlgByGenes	Use the BstBstNbrsByGene function.
BestHmlgsByGenes	Use the BstBstNbrsByGene function.
DefinitionByGeneS	Use the new BtitS function.

Changes to the function arguments

Most of the function arguments have new names. There are some arguments that
are no longer used by the KEGG user-defined functions and several new
arguments have been added.

The following table describes the changes to the arguments that are used with the
KEGG user-defined functions.

Table 10. KEGG UDF argument changes

Former argument name	New argument name	New description
cpdlist	compound_id	The unique identifier for a compound. The format that you use for the compound_id argument is `cpd:compound`. The cpd is the abbreviation for the compound database. The *compound* is the compound number that is in the KEGG LIGAND composite database. For example, the compound identifier for Methylparabanic acid is: `cpd:C11116`.
	compound_id_list	This argument is new in Fixpack 9. A space-delimited list of compound identifiers.
enzymelist	enzyme_id	The unique identifier for an enzyme. The format that you use for the enzyme_id argument is `ec:enzyme`. The ec is the abbreviation for the enzyme database. The *enzyme* is the enzyme code that is in the enzyme database. For example, the compound identifier for Sulfite oxidase is: `ec:1.8.3.1`.
	enzyme_id_list	This argument is new in Fixpack 9. A space-delimited list of enzyme identifiers.

Table 10. KEGG UDF argument changes (continued)

Former argument name	New argument name	New description
keggid	genes_id	A unique identifier for a gene. The format that you use for the genes_id argument is org:*gene_name*. The org is the 3-letter KEGG code for an organism. The *gene_name* is the name of the gene.
		For example, the genes identifier for the E. coli gene is: eco:*b0001*.
	gene_id_list	This argument is new in Fixpack 9.
		A space-delimited list of genes identifiers.
	max_result	This argument is new in Fixpack 9.
		An integer that is used to specify the number of results that are returned. The max_results argument is used with the start argument to control the results that are returned. For example, if 5 is specified for start and 12 is specified for max_results, the results that are returned start with result 5 and end with result 16.
midlist	motif_id	A unique identifier for a motif. The format that you use for the motif_id argument is motif_database_identifier:*motif_name*. The motif_database_identifier is the 2-letter code for a motif database. The *motif_name* is the name of the motif. The valid motif database identifiers are: • The bl Identifier for the Blocks Motif database • The pd Identifier for the Prodom Motif database • The pf Identifier for the Pfam Motif database • The pr Identifier for the Prints Motif database • The ps Identifier for the Prosite Motif database
		For example, the motif identifier for the DnaJ entry in the pfam database is: pf:*DnaJ* .
orglist	org	A 3-letter KEGG code for an organism. If this argument is not specified, the entire organism list is searched.
		Each organism in the KEGG databases is assigned a code. The list of organisms changes frequently. Check the current list of genomes at: http://www.genome.ad.jp/kegg/kegg2.html#genes for the correct codes.

Table 10. KEGG UDF argument changes (continued)

Former argument name	New argument name	New description
pathwayid	pathway_id	A unique identifier for a pathway. The format that you use for the pathway_id argument is `path:`*`pathway_number`*. The `path` is the abbreviation for the pathway database. The *pathway_number* consists of a prefix and a number. The valid prefixes are `map` and `org`. The `map` prefix indicates that you want to search for the reference pathway. The `org` prefix is a 3-letter KEGG code for an organism.
		For example, the pathway identifier for the reference pathway for the cytrate cycles is: `path:`*`map00020`*
		For example, the pathway identifier for the specific pathways for the E.coli genes is: `path:`*`eco00020`*.
	start	This argument is new in Fixpack 9.
		An integer that is used to specify the first result that is returned. The start argument is used with the max_results argument to control the results that are returned. For example, if 5 is specified for start and 12 is specified for max_results, the results that are returned start with result 5 and end with result 16.
threshold		This argument is no longer supported.

Changes to the Pathway function names and accepted arguments

The following table lists the changes to the KEGG user-defined functions for the Pathway database. The function arguments and descriptions have changed.

Table 11. Changes to the Pathway functions and arguments

Function name	New function arguments	Description
CompoundsByPathwy	pathway_id	Use the CompoundsByPathwyS and CompoundsByPathwyT functions to search for all of the compounds on a pathway. These functions now return compound identifiers. In version 2.3, these functions returned the compound.
EnzymesByPathwy	pathway_id	Use the EnzymesByPathwyS and EnzymesByPathwyT functions to search for all of the enzymes on a pathway. These functions now return enzyme identifiers. In version 2.3, these functions returned the enzyme.
GenesByPathwy	pathway_id	Use the GenesByPathwyS and GenesByPathwyT functions to search for all of the genes on a pathway. These functions now return gene identifiers. In version 2.3, these functions returned the gene.
PathwysByCompnds	compound_id_list	Use the PathwysByCompndsS and PathwysByCompndsT functions to search for all of the pathways that contain the compounds that you specify. These functions now return pathway identifiers. In version 2.3, these functions returned the pathway.

| *Table 11. Changes to the Pathway functions and arguments (continued)*

Function name	New function arguments	Description
PathwysByEnzymes	enzyme_id_list	Use the PathwysByEnzymesS and PathwysByEnzymesT functions to search for all of the pathways that contain the enzymes that you specify. These functions now return pathway identifiers. In version 2.3, these functions returned the pathway.
PathwysByGenes	gene_id_list	Use the PathwysByGenesS and PathwysByGenesT functions to search for all of the pathways that contains the genes that you specify. These functions now return pathway identifiers. In version 2.3, these functions returned the pathway.

Changes to the SSDB function names and accepted arguments

The following table lists the changes to the KEGG user-defined functions for the SSDB database. Some of the function names and most the function arguments have changed.

Table 12. Changes to the SSDB functions and arguments

Former function name	New function name	New function arguments	New function description
AllNbrsByGene	GetNbrsByGene	genes_id, org, start, max_results	Use the GetNbrsByGeneS and GetNbrsByGeneT functions to search for organisms that are homologous neighbors to the gene that you specify.
BestNbrsByGene	no name change	genes_id, start, max_results	Use the BestNbrsByGeneS and BestNbrsByGeneT functions to search for the best neighbors or the best homologous neighbors of a gene in all organisms.
BstBstNbrsByGene	no name change	genes_id, start, max_results	Use the BstBstNbrsByGeneS and BstBstNbrsByGeneT functions to search for the best-best neighbors or the best-best homologous neighbors of a gene in all organisms.
GenesByMotifsT	no name change	motif_id_list, start, max_results	Use the GenesByMotifsT function to search for all of the genes that contain all of the motifs in a list of motifs that you specify. This function now returns a Definition row. In version 2.3, this function returned the keggid and definition.
ParalogsByGene	no name change	genes_id, start, max_results	Use the ParalogsByGeneS and ParalogsByGeneT functions to search for the genes that are paralogous to the gene that you specify.
RevBstNbrsByGene	no name change	genes_id, start, max_results	Use the RevBestNbrsByGeneS and RevBestNbrsByGeneT functions to search for the reverse best neighbors of a gene in all organisms.
ScoreBtwGenes	SimilarBtwGenesT	genes_id, genes_id2	Use the SimilarBtwGenesT function to determine the Smith-Waterman score and the alignment positions between two genes. This function now returns an SSDB row. In version 2.3, this function returned a DOUBLE data type.

Changes to the columns that are returned from SSDB database queries (table functions)

The columns that are returned depend on the function that you use and the type of row that is returned. The table functions in the version 2.3 KEGG user-defined functions returned SSDB rows.

In the version 3.2 KEGG user-defined functions, the following types of rows are returned:

- SSDB rows
- Motif rows
- Definition rows

Changes to the SSDB rows: The columns that are retuned with the SSDB rows have changed. The following table shows the changes to the columns.

Table 13. SSDB rows from the SSDB database

Version 2.3 column name	Version 3.2 column name	Column data type	Description
keggid1	genes_id1	VARCHAR (100)	The identifier for the gene that is specified in the query.
keggid2	genes_id2	VARCHAR (100)	The identifier for the gene that is returned from the query.
sw_score	sw_score	INTEGER	The Smith-Waterman score between genes_id1 and genes_id2.
N/A	bit_score	DOUBLE	The bit score between genes_id1 and genes_id2
identity	identity	DOUBLE	The identity percent between genes_id1 and genes_id2.
overlap	overlap	INTEGER	The overlap length between genes_id1 and genes_id2.
s1_start	start_position1	INTEGER	The start position of the alignment in genes_id1.
s1_end	end_position1	INTEGER	The end position of the alignment in genes_id1.
s2_start	start_position2	INTEGER	The start position of the alignment in genes_id2.
s2_end	end_position2	INTEGER	The end position of the alignment in genes_id2.
best1	best_flag_1to2	INTEGER	The flag that indicates the best hit from genes_id1 to genes_id2.
best2	best_flag_2to1	INTEGER	The flag that indicates the best hit from genes_id2 to genes_id1.
def1	definition1	VARCHAR (1000)	The definition for the genes_id1 column.
def2	definition2	VARCHAR (1000)	The definition for the genes_id2 column.
length1	length1	INTEGER	The length of the amino acid in the genes_id1 column.
length2	length2	INTEGER	The length of the amino acid in the genes_id2 column.

Motif rows: Motif rows are new with the version 3.2 KEGG user-defined functions. The following table lists the columns that are returned with Motif rows:

Table 14. Motif rows from the SSDB database

Column name	Column data type	Description
motif_id	VARCHAR (100)	The identifier for the motif.
definition	VARCHAR (1000)	The definition for the motif.
genes_id	VARCHAR (100)	The identifier for the gene that contains the motif.
start_position	INTEGER	The start position of the motif_id that matches the query.
end_position	INTEGER	The end position of the motif_id that matches the query.
score	DOUBLE	The Smith-Waterman score of the motif in the TIGRFAM and PROSITE databases that matches the query. The data type of the score in the PROSITE database is FLOAT.
evalue	DOUBLE	The e-value of the motif in the Pfam database that matches the query.

Definition rows: Definition rows are new with the version 3.2 KEGG user-defined functions. The following table lists the columns that are returned with Definition rows:

Table 15. Definition rows from the SSDB database

Column name	Column data type	Description
entry_id	VARCHAR (100)	The identifier for the database entry.
definition	VARCHAR (1000)	The definition for the entry.

New CREATE SERVER statement options

The following information is new in the DB2 Information Center for fix pack 9:

Topic title: CREATE SERVER statement options - Web services wrapper

Location in the DB2 Information Center
Configuring —> Federated systems and data sources

Optional parameters
Options for the CREATE SERVER statement for Web services are:

TIMEOUT
Optional: Specifies the time, in minutes, that the DB2 Universal Database server should wait for a network transfer and the computation of a result. The following example sets the maximum time out value as 60 minutes:

```
CREATE SERVER my_server WRAPPER my_wrapper
OPTIONS (TIMEOUT '60')
```

PROXY_TYPE
Optional: Type of proxy. The only valid value is HTTP.

PROXY_SERVER_NAME
Optional: The name of the proxy server.

PROXY_SERVER_PORT
Optional: The port number of the proxy server.

Important: If you specify any proxy server setting, you must specify all three.

For example:

```
CREATE SERVER my_server WRAPPER my_wrapper
OPTIONS (PROXY_TYPE 'HTTP', PROXY_SERVER_NAME
'proxy.mycompany.com', PROXY_SERVER_PORT '8080')
```

Fix Pack 8 Documentation updates

Corrections to documentation

The following information describes corrections to documentation for fix pack 8.

All of the navigational references to the DB2 Information Center in this section assume that you are using the DB2 Information Center navigation tree as the starting point.

Product overviews

Topic title: Supported data sources

> **Location in the DB2 Information Center:**
>> Product overviews —> Information Integration —> DB2 Information Integrator overview —> Federated systems overview

> **Addition 1 of 6**
>> For Microsoft SQL Server: If you use the Microsoft SQL Server wrapper with a UTF-8 database on UNIX you must use DataDirect Connect for ODBC 4.2 Service Pack 2 or later.

> **Addition 2 of 6**
>> For Informix data sources:
>> - On Solaris, the Informix client version 2.81.xC2 is not supported. If you are using the Informix client version 2.81.xC2, update the client to version 2.81.xC2R1 or later.
>> - On Windows, the Informix SDK client version 2.81.TC2 or later is required.
>> - If you use Informix Dynamic Server 9.3 as a data source, you must use Informix Dynamic Server 9.30.xC4 or later.
>> - On the 64-bit mode zLinux operating system, you must use the Informix client version 2.81.FC3.

> **Addition 3 of 6**
>> For Teradata data sources: On Windows, you must use the Teradata client TTU 7.0 or later and the Teradata API library CLIv2 4.7.0 or later on the federated server.

> **Addition 4 of 6**
>> For Oracle data sources: Oracle 8.0.6 and 8.1.6 are not supported versions.

> **Addition 5 of 6**
>> For BioRS and Entrez, the access method should be HTTP.

> **Addition 6 of 6**
>> For BioRS the supported versions are: 5.0.14, 5.2.x.x.

Migrating

Topic title: Verifying that your databases are ready for migration

> **Location in the DB2 Information Center:**
>> Migrating —> Database Systems —> DB2 Universal Database

Version 8.2 —> Migrating from previous versions of DB2 Universal
Database —> Servers —> Migrating DB2 Servers (Linux and
UNIX) —> Premigration tasks

Correction 1 of 2

The second sentence in the topic should read as follows:

Use the db2ckmig command to verify that databases that are
owned by the current instance are ready to be migrated.

Correction 2 of 2

The following paragraph should be added to step 3:

On UNIX (AIX), you can run the db2ckmig command from the
install_directory/instance where install_directory is
/usr/opt/db2_08_01/ on AIX and /opt/IBM/db2_V* on all other
UNIX systems.

Installing

**Topic title: Editing the Oracle genclntsh script and creating the libclntsh file
after you install DB2 Information Integrator (HP-UX, Linux, Solaris)**

Location in the DB2 Information Center:

Installing –> Installing DB2 Information Integrator on Windows
and UNIX

Correction 1 of 2

The topic, "Editing the Oracle genclntsh script and creating the
libclntsh file after you install DB2 Information Integrator (HP-UX,
Linux, Solaris)" has been updated with the following information
for the Oracle 10g client on Linux.

Back up the following files for Linux:
- $ORACLE_HOME/bin/genclntsh
- $ORACLE_HOME/lib/libclntsh.so.10.1
- $ORACLE_HOME/lib32/libclntsh.so.10.1

Back up the following files for HP-UX:
- $ORACLE_HOME/bin/genclntsh
- $ORACLE_HOME/lib/libclntsh.sl.10.1
- $ORACLE_HOME/lib32/libclntsh.sl.10.1

Correction 2 of 2

To edit the genclntsh script and create the libclntsh file for Oracle
10g on Linux and HP-UX operating systems after you install DB2
Information Integrator:

1. Rename the following files:
   - On Linux, rename the $ORACLE_HOME/lib/libnnz10.so file
     to $ORACLE_HOME/lib/libnnz10.so.orig
   - On HP-UX, rename the $ORACLE_HOME/lib/libnnz10.sl
     file to $ORACLE_HOME/lib/libnnz10.sl.orig and the
     $ORACLE_HOME/lib32/libnnz10.sl file to
     $ORACLE_HOME/lib32/libnnz10.sl.orig
2. In a text editor, open the $ORACLE_HOME/bin/genclntsh
   script.

a. On Linux, add "-Wl,-Bsymbolic" to the link line, and ensure that there are no spaces between Wl, and -Bsymbolic. For example:

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

On HP-UX, add -Bsymbolic to the link line:

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

b. On Linux, change the name of the $ORACLE_HOME/lib/libnnz10.so library to $ORACLE_HOME/lib/libnnz10.a.

3. Run the genclntsh script from a command prompt to create the libclntsh file.

On HP-UX, you can ignore the following error:

```
ld: I/O error, file "/opt/oemclient/oracle/
10client/rdbms/lib/sllfls.o":
No such file or directory Fatal error.
```

4. Rename the following files:

On Linux, rename the $ORACLE_HOME/lib/libnnz10.so.orig file to $ORACLE_HOME/lib/libnnz10.so.

On HP-UX, rename the $ORACLE_HOME/lib/libnnz10.sl.orig file to $ORACLE_HOME/lib/libnnz10.sl, and the $ORACLE_HOME/lib32/libnnz10.sl.orig file to $ORACLE_HOME/lib32/libnnz10.sl

5. Run the djxlinkOracle script from a command prompt to update the Oracle wrapper library.

6. Run the db2iupdt command on each DB2 instance to enable federated access to the data sources.

7. Restart all of the instances that use the NET8 wrapper.

Topic title: Supported operating systems for DB2 Information Integrator (32-bit)

Location in the DB2 Information Center

Installing —> Information Integration —> Planning to install DB2 Information Integrator —> Hardware and software requirements

Correction

The Web services wrapper is supported on Solaris.

Topic title: Installing DB2 Information Integrator (Windows) and Installing DB2 Information Integrator (UNIX)

Location in the DB2 Information Center:

Installing —> Information Integration —> Installing DB2 Information Integrator on Windows and UNIX —> Installing DB2 Information Integrator on UNIX and Windows

Correction

To use Q replication or event publishing, you must install WebSphere MQ version 5.3 Fix Pack 05 (CSD05).

Topic title: deployXMR syntax

Location in the DB2 Information Center:

Installing —> Information integration —> Installing DB2 Information Integrator on Windows and UNIX —> Installing the DB2 XML Metadata Registry - overview

Correction

The following correction shows the updated explanation of the -doc parameter of the deployXMR command:

-doc helpSystemHomePageURL:Port

The URL of the help system home page and its port number. If you do not specify a port number, the default port number for the protocol is used. The documentation for the XML Metadata Registry is in the DB2 Information Center. The information center can be installed locally or on an intranet server from the CD that is included with the DB2 Universal Database product CDs.

For example, if you installed the information center on myhost.example.org on port 8081, specify **-doc** `http://myhost.example.org:8081/db2help/index.jsp`. If you choose not to install the DB2 Information Center, you can specify the URL of the online information center:

`-doc http://publib.boulder.ibm.com/infocenter/db2help/index.jsp`

Topic title: Changing to a different edition of DB2 Information Integrator

Location in the DB2 Information Center:

Installing —> Installing DB2 Information Integrator on Windows and UNIX

Addition 1

The following information is added to the existing information in this topic.

If DB2 Information Integrator Advanced Edition or DB2 Information Integrator Advanced Edition Unlimited is installed and you change to another edition of DB2 Information Integrator, you must remove the IBM Tivoli® License Manager (ITLM) file from your system before you install a different edition.

The following table shows the name of the ITLM file depending on your operating system and edition.

Table 16. Table of file names for the DB2 Information Integrator ITLM file

Operating system	Advanced Edition ITLM file name	Advanced Edition Unlimited ITLM file name
AIX	db2iiaex080200.sys	db2iiuex080200.sys
Linux™	db2iiael080200.sys	db2iiuel080200.sys
HP-UX	db2iiaeh080200.sys	db2iiueh080200.sys
Solaris	db2iiaes080200.sys	db2iiues080200.sys
Windows	db2iiaew080200.sys	db2iiuew080200.sys

Addition 2

Administrator authority is also required for removing the ITLM file from your system.

Addition 3

The following step is added between steps 4 and 5 of the original topic.

5. If Advanced Edition or Advanced Edition Unlimited is installed, remove the ITLM file from one of the following directories depending on your operating system.

Topic title: Removing the DB2 Information Integrator and DB2 product license keys

Location in the DB2 Information Center:
Installing —> Information integration —> Removing DB2 Information Integrator from your system

Addition 1
If you are removing Advanced Edition or Advanced Edition Unlimited, you must remove the IBM Tivoli License Manager file in addition to removing the product license keys.

The following table shows the ITLM files for each edition and operating system.

Table 17. Table of file names for the DB2 Information Integrator ITLM file

Operating system	Advanced Edition ITLM file name	Advanced Edition Unlimited ITLM file name
AIX	db2iiaex080200.sys	db2iiuex080200.sys
Linux	db2iiael080200.sys	db2iiuel080200.sys
HP-UX	db2iiaeh080200.sys	db2iiueh080200.sys
Solaris	db2iiaes080200.sys	db2iiues080200.sys
Windows	db2iiaew080200.sys	db2iiuew080200.sys

Addition 2
To remove the DB2 Information Integrator license key and the ITLM file, you must have administrator authority.

Addition 3
The following step is added to the original topic:

9. If you are removing Advanced Edition or Advanced Edition Unlimited, remove the ITLM file from one of the following directories depending on your operating system:

Windows

x:\Program Files\IBM\DB2InformationIntegrator\V8\

x is the drive where DB2 Information Integrator is installed.

UNIX /opt/IBM/DB2InformationIntegrator/V8/

Configuring

Topic title: Required settings for WebSphere MQ objects

Location in the DB2 Information Center:
Configuring –> Replication and event publishing –> Configuring for Q replication and event publishing –> Setting up WebSphere MQ

Correction
The DISCINT parm is listed as a valid parameter for sender channels, but it is also valid for receiver channels.

Topic title: Memory used by the Q Apply program

Location in the DB2 Information Center:
Configuring –> Replication and event publishing –> Configuring for Q replication and event publishing –> Planning memory and storage requirements –> Planning memory requirements

Correction
If you are running Q replication or event publishing in a 64-bit environment on the HP-UX or Solaris platforms, catalog the database (the Q Apply server) as a loop back database and create an entry for this database in the password file. If you do not catalog the Q Apply server as a loop back database, you will encounter a Semaphore Wait problem, sqlcode 1224, when the Q Apply program reaches the shared memory limit.

Topic title: WebSphere Business Integration wrapper library files

Location in the DB2 Information Center
Configuring —> Federated systems and data sources —> Configuring data sources —> Configuring business application data sources

Correction
The WebSphere Business Integration wrapper is supported on Solaris and Linux systems. The name of the library file for Solaris and Linux is libdb2ws.so. The name of the path is /opt/IBM/db2/V8.1/lib

Topic title: Transparent DDL

Location in the DB2 Information Center:
Configuring —> Federated systems and data sources —> Configuring a federated system —> Creating and modifying remote tables by using transparent DDL

Addition
In a multiple partitioned environment, you must run transparent DDL statements at the catalog partition.

Topic title: Adding data sources to a federated server using the DB2 UDB Control Center

Location in the DB2 Information Center
Configuring —> Federated systems and data sources —> Configuring data sources —> Data source configuration - overview

Correction
The following prerequisite should be added to this topic:

Prerequisite: To use the discover feature to create server definitions for Microsoft SQL Server data sources, you must edit the data source configuration file before you use the discover feature.

To edit the data source configuration file:

1. Determine the name of the data source configuration file. The data source configuration file is specified in the db2dj.ini file in the ODBCINI entry. Typically, the name of the data source configuration file is odbc.ini or .odbc.ini.
2. Edit the data source configuration file.
3. In the data source configuration file, locate the Microsoft SQL Server data sources in the [ODBC Data Sources] section.

4. Add entries for the Microsoft SQL Server data sources that you want to access, if entries for the data sources do not already exist. The entry must include the text SQL Server.

   For example, the entry for a DSN with the name mssqldsn is:

   ```
   mssqldsn=Microsoft SQL Server 2003
   ```

5. Save the changes to the data source configuration file.

Topic title: Registering server definitions for a data source

Location in the DB2 Information Center

Configuring —> Federated systems and data sources —> Configuring data sources —> Data source configuration - overview

Correction

The federated server does not verify whether the version of the server that you specify in the CREATE SERVER statement matches the version of the data source server. If you specify an incorrect version in the CREATE SERVER statement, you might receive SQL errors. Errors can occur when you issue SQL statements that specify a nickname that relies on the inaccurate server definition.

Topic title: CREATE SERVER statement options - BioRS wrapper

Location in the DB2 Information Center

Configuring —> Federated systems and data sources —> Configuring data sources —> Configuring BioRS data sources —>Adding BioRS data sources to a federated server

Correction

The text that describes the VERSION parameter is not correct. The correct text is:

VERSION

Specifies the version of the BioRS server. Examples of valid values are 5.0 or 5.2.1.10. If you to do not specify the VERSION parameter, the default value 1.0 is used. This default value is used for backward compatibility, and is equivalent to specifying the version 5.0.

Topic title: Constructing BLAST SQL queries

Location in the DB2 Information Center

Configuring —> Federated systems and data sources —> Configuring data sources —> Configuring BLAST data sources —> Adding BLAST to a federated system

Correction

Wrappers that require predicates, such as the BlastSeq fixed input column, cannot process queries that result in a left outer join on the required predicate.

For example, the following query returns an SQL0901N error:

```
SELECT n1.Score FROM blastNN1 n1
    LEFT OUTER JOIN myseqs n2 ON N1.BlastSeq = n2.seq
```

Addition

If the BlastSeq column for a BLAST nickname is defined as a CLOB data type, you must create sequence match function templates before you can

query BLAST data sources. The SQL statements that you issue on the federated database to create the function templates are:

```
CREATE FUNCTION lblast.sequence_match (CLOB, CLOB)
    RETURNS INTEGER AS TEMPLATE DETERMINISTIC NO EXTERNAL ACTION;
CREATE FUNCTION lblast.sequence_match (CLOB, VARCHAR(1))
    RETURNS INTEGER AS TEMPLATE DETERMINISTIC NO EXTERNAL ACTION;
CREATE FUNCTION lblast.sequence_match (VARCHAR(1), CLOB)
    RETURNS INTEGER AS TEMPLATE DETERMINISTIC NO EXTERNAL ACTION;
CREATE FUNCTION lblast.sequence_match (VARCHAR(1), VARCHAR(1))
    RETURNS INTEGER AS TEMPLATE DETERMINISTIC NO EXTERNAL ACTION;
```

Topic title: Registering nicknames for business application data sources

Location in the DB2 Information Center:
Configuring —> Federated systems and data sources —> Configuring data sources —> Configuring business application data sources —> Adding business applications to a federated system

Correction
The following sentence should be included in the introduction to registering nicknames:

The DB2 Control Center generates unique nicknames by grouping part names or element names with the column name from the XML schema document.

Topic title: Custom functions and Entrez queries

Location in the DB2 Information Center
Configuring —> Federated systems and data sources —> Configuring data sources —> Configuring Entrez data sources —> Adding Entrez to a federated system —> Queries and custom functions for Entrez data sources

Correction
When you use the CONTAINS function on a fetch key, the results of the query might not be accurate. The fetch key for PubMed nicknames is pmid. The fetch key for Genbank nicknames is gi.

For example, the following SELECT statement uses the CONTAINS function. When you issue this statement, no rows are returned:

```
SELECT g.gi,g.seqlength FROM gbseq g
   WHERE entrez.contains(g.gi,'23273757')=1;
```

If you issue the SELECT statement with an equal predicate, the rows are returned:

```
SELECT g.gi,g.seqlength FROM gbseq g
   WHERE g.gi = '23273757';
```

Rows that are not returned when you use the CONTAINS function, but are returned when you use an equal predicate, might be marked as dead in the Entrez database and not available for search.

Topic title: What is Extended Search?

Location in the DB2 Information Center:
Configuring —> Federated systems and data sources —> Configuring data sources —> Configuring Extended Search data sources

Addition

>To use Extended Search with WebSphere, you must have Extended Search 4.0.2 and WebSphere 5.1.1.

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

Correction

>The last table in this topic lists the fixed columns that the wrapper creates when the VERTICAL_TABLE option is used. In the list of valid data types that the column FIELD_DATATYPE can represent, the DOUBLE data type is also a valid data type.

Topic title: Extended Search wrapper - query guidelines

Location in the DB2 Information Center

>Configuring —> Federated systems and data sources —> Configuring data sources —> Configuring Extended Search data sources

Correction 1 of 5

>When an SQL query contains predicates in the WHERE clause that the Extended Search wrapper cannot process, the following SQL0901N error is returned:

```
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
```

The following examples show queries that return the SQL0901N error.

Example of a query without a predicate:

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

Example of query that uses the ANY, ALL, SOME, or EXIST predicates:

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

Example of a query that uses the NOT IN predicate and is not supported with a FULL SELECT clause:

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

Example of a query that uses the IN predicate and is not supported with a FULL SELECT clause:

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

Example of a query that uses a fixed column predicate:

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

If the SQL query contains at least one predicate that the Extended Search wrapper can process, the query is accepted and run.

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

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

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

Correction 2 of 5

The Extended Search wrapper can process queries that specify the SQL LIKE predicate if the General Query Language (GQL) that Extended Search uses can process the LIKE predicate. The data sources that support the LIKE predicate are listed in the GQL support matrix in the *IBM Lotus Extended Search Programming*.

Examples of valid SQL statements with the LIKE predicate:

```
SELECT * FROM nickname WHERE title LIKE "%defense%"

SELECT * FROM nickname WHERE title LIKE "%defense"

SELECT * FROM nickname WHERE title LIKE "defense%"
```

Example of an invalid SQL statement with the LIKE predicate:

```
SELECT * FROM nickname WHERE title LIKE "defen_e"
```

In this example, the Extended Search wrapper cannot convert the single character specification in the SQL query into a valid GQL statement.

Correction 3 of 5

Limitations of the Extended Search data sources affect the queries that you can run. Before you run queries on Extended Search data sources, check the data types, the operators in the predicates, and the target Extended Search data sources. For example, you cannot specify combinations of AND, OR, or SUBTRACT predicates in the es_search() table function queries for File System data sources on the Extended Search server. For a complete list of supported predicates, see the GQL support matrix in *IBM Lotus Extended Search Programming*.

Correction 4 of 5

When you use an SQL date expression with the Extended Search custom function ESWRAPPER.ES_SEARCH, you must use the YYYYMMDD date format. For example, for December 31, 2004, the date that you specify is 20041231.

Correction 5 of 5

Some of the Extended Search wrapper data sources do not support comparison operators in string fields. Examples of data sources that do not support comparison operators are File System, Lotus Notes Formula, Lotus Notes Full Text, and Web. The types of operators that are not supported include the greater than, less than, less than or equal to, and greater than or equal to operators.

For example, the following query returns an SQL1822N error. This query selects data from a VARCHAR column with the column name LA.

```
SELECT YEAR(JournalDate) FROM notes.alzheimer_13 WHERE LA > 'A';
```

The following error is returned from this query:

```
SQL1822N  Unexpected error code "ES Non-critical Error" received from data
source "ES Server: iisvtnt2.sv". Associated text and tokens are "DBCQAE0011
Cannot transla".  SQLSTATE=560BD
```

The GQL support matrix that is in the *IBM Lotus Extended Search Programming*document lists the comparison operator limitations.

Topic title: Registering the Extended Search wrapper

Location in the DB2 Information Center

Configuring —> Federated systems and data sources —> Configuring data sources —> Configuring Extended Search data sources —> Adding Extended Search data sources to a federated system

Correction

The following prerequisite applies to this topic:

Prerequisite: With IBM Lotus Extended Search, Version 4.0, values that are returned from the Extended Search server might not be converted from the Extended Search server operating system code page to UTF-8. Before you register the Extended Search wrapper, you must do either of these tasks:

- Install fix pack 1 for IBM Lotus Extended Search, Version 4.0 on the Extended Search server.
- Change the Extended Search server operating system code page. To change the Extended Search server operating system code page, you must add a line to the esNLS.ini file on the Extended Search server. The esNLS.ini file is in the *Extended_Search_installation_path*/config directory. Add the following line to the esNLS.ini file:

  ` A.ALL.File\ System.Title=1`

  Then stop and restart the Extended Search Server and the Extended Search Remote Method Invocation Server.

Topic title: Construct new HMMER queries with samples

Location in the DB2 Information Center

Configuring —> Federated systems and data sources —> Configuring data sources —> Configuring HMMER data sources —> Adding HMMER to a federated server

Correction

Wrappers that require predicates, such as the HmmerQSeq fixed input column, cannot process queries that result in a left outer join on the required predicates. For example, the following query returns an SQL0901N error:

```
SELECT n1.Score FROM hmmerNN1 n1
   LEFT OUTER JOIN myseqs n2 ON N1.HmmerQSeq = n2.seq
```

Addition

If the HMMQSEQ column for a HMMER nickname is defined as a CLOB data type, you must create sequence match function templates before you can query HMMER data sources. The SQL statements that you issue on the federated database to create the function templates are:

```
CREATE FUNCTION lshmmer.sequence_match (CLOB, CLOB)
   RETURNS INTEGER AS TEMPLATE DETERMINISTIC NO EXTERNAL ACTION;
CREATE FUNCTION lshmmer.sequence_match (CLOB, VARCHAR(1))
   RETURNS INTEGER AS TEMPLATE DETERMINISTIC NO EXTERNAL ACTION;
```

```
CREATE FUNCTION lshmmer.sequence_match (VARCHAR(1), CLOB)
    RETURNS INTEGER AS TEMPLATE DETERMINISTIC NO EXTERNAL ACTION;
CREATE FUNCTION lshmmer.sequence_match (VARCHAR(1), VARCHAR(1))
    RETURNS INTEGER AS TEMPLATE DETERMINISTIC NO EXTERNAL ACTION;
```

Topic title: KEGG user-defined functions - overview

Location in the DB2 Information Center

Configuring —> Federated systems and data sources —> Configuring data sources —> Configuring KEGG data sources using KEGG user-defined functions

Correction 1 of 2

The KEGG user-defined functions use the KEGG API, Release 2.3.

Correction 2 of 2

On some networks, access to the Internet must go through a firewall. To use the KEGG user-defined functions with a firewall, you can set up proxy support using the steps described in the topic "The Web service consumer functions" in the *IBM DB2 Information Integrator Application Developer's Guide*.

Topic title: Registering the KEGG user-defined functions

Location in the DB2 Information Center

Configuring —> Federated systems and data sources —> Configuring data sources —> Configuring KEGG data sources using KEGG user-defined functions

Correction 1 of 2

To use the KEGG user-defined functions, you must install XML Extender and then enable XML Extender, the SOAP user-defined functions, and the KEGG user-defined functions.

- To install XML Extender:
  1. After you install DB2 Universal Database and DB2 Information Integrator, go to the DB2 ESE installation directory and run ./db2setup again (as root) to start the installation wizard.
  2. On the Install Products page select **DB2 Application Development Client**.
  3. On the Welcome page, select **Install Additional Features** and click **Next**.
  4. On the Installation Type page, select **Custom** and click **Next**.
  5. On the Installation Action page, click **Next**.
  6. On the Select Features to Install page, select **+** in front of Client Support to expose all of the software items.
  7. Check **XML Extender** and click **Next**.
  8. On the Languages page, click **Finish**.
- To enable XML Extender, the SOAP user-defined functions, and the KEGG user-defined functions:
  1. Run the `dxxadm enable_db database_name` command.
  2. Run the `db2enable_soap_udf -n database_name` command.
  3. Run the `enable_KEGGFunctions` command.

Correction 2 of 2

There is a new flag for the command to enable the KEGG user-defined functions. The correct syntax is:

```
enable_KEGGFunctions -n dbName
 -u userID -p password [-force] [-url endpointURL]
```

The definition for endpointURL flag is: The endpoint URL for the KEGG Web services API, version 2.3. The default URL is set to http://soap.genome.ad.jp/keggapi/request_v2.3.cgi. The endpointURL flag is necessary only if the KEGG Web services API, version 2.3 changes the name or location of the endpoint URL.

Topic title: Disabling the KEGG user-defined functions

Location in the DB2 Information Center
>Configuring —> Federated systems and data sources —> Configuring data sources —> Configuring KEGG data sources using KEGG user-defined functions

Correction
>After you disable the KEGG user-defined functions, you can disable the XML Extender and the SOAP user-defined functions:
>
>1. Run the `disable_KEGGFunctions` command.
>2. Run the `db2disable_soap_udf -n database_name` command.
>3. Run the `dxxadm disable_db database_name` command.

Topic title: Adding ODBC data sources to a federated server

Location in the DB2 Information Center
>Configuring —> Federated systems and data sources —> Configuring data sources —> Configuring ODBC data sources

Correction
>The following text has been added to the Restrictions section in this topic:
>
>• The ODBC wrapper does not support operations on tables that contain columns with data types that use driver-specific SQL data type indicators. The type of operations that are not supported include the CREATE NICKNAME and SELECT statements in the pass-thrugh mode. The ODBC wrapper supports only the SQL data type indicators that are defined by the ODBC standard in the *Microsoft ODBC Programmer's Reference*.

Topic title: Tuning and troubleshooting the configuration to ODBC data sources

Location in the DB2 Information Center
>Configuring —> Federated systems and data sources —> Configuring data sources —> Configuring ODBC data sources —> Adding ODBC data sources to federated servers

Correction
>If queries fail because of data source syntax errors, set the PUSHDOWN server option value to 'N', or remove the PUSHDOWN server option.

Topic title: Testing the connection to the Teradata server

Location in the DB2 Information Center
>Configuring —> Federated systems and data sources —> Configuring data sources —> Configuring Teradata data sources —> Adding Teradata data sources to a federated system

Correction

> To verify that the character set that you want to use is installed on the Teradata server:
>
> 1. Log on to the Teradata server by using the BTEQ utility or any other valid logon utility.
> 2. Issue the following statement to display the dbc.chartranslations table:
>
>    ```
>    select * from dbc.chartranslations;
>    ```
>
> 3. Check the value in the third column, InstallFlag, of the table that is returned. The value 'Y' in the third column indicates that the character set is installed and in use on the Teradata server.
>
>   Use the following table to determine if you have the correct character set installed:

Table 18. Character sets for Teradata

Double- byte character set	Single- byte character set	Teradata character set	Language	IBM DB2 code set
941	897	"KanjiSJIS_0S"	Japanese	IBM-943
1362	1126	"HANGULKSC5601_2R4"	Korean	1363
1385	1114	"SCHGB2312_1T0"	Simplified Chinese	GBk
380	1115	"SCHGB2312_1T0"	Simplified Chinese	IBM-1381
947	1114	"TCHBIG5_1R0"	Traditional Chinese	big5
1200	1208	"UTF8"	Unicode	UTF-8
0	819	"Latin1_0A"	English (Latin 1)	ISO8859-1
0	1252	"Latin1252_0A"	English (Win Latin)	ISO8859-1/15

> 4. If you do not have the required character set installed, install the character set to use the Teradata wrapper.
>   - If the character set that you want to use is listed in the dbc.chartranslations table, but the InstallFlag value is set to 'N', issue the following statement to change the InstallFlag to 'Y':
>
>      ```
>      update dbc.chartranslations
>          set installflag='Y' where CharSetName= 'character_set_name';
>      ```
>
>   - If the character set that you want to use is not listed in the dbc.chartranslations table, contact Teradata customer support.
> 5. Restart the Teradata server to update the list of character sets. In a Teradata command window, enter:
>
>    ```
>    tpareset -f reason_for_restart
>    ```

Topic title: Registering nicknames for Web services data sources

Location in the DB2 Information Center:

> Configuring —> Federated systems and data sources —> Configuring data sources —> Configuring Web services data sources —> Adding Web services to a federated system

Correction 1 of 2

> The following sentence should be included in the introduction to registering nicknames:

If a WSDL operation has an input variable that contains optional grouping elements for the value, then the DB2 Control Center generates input-only child nicknames. You can use these nicknames to specify optional input by joining the input-only nickname with the root of the nickname hierarchy for that operation.

Correction 2 of 2

The following sentence should be included in the introduction to registering nicknames:

When the DB2 Control Center generates a child nickname that is used only for input, the XPATH nickname option contains a period, as in the following example:

XPATH '.'

Topic title: Web services wrapper library files

Location in the DB2 Information Center

Configuring —> Federated systems and data sources —> Configuring data sources —> Configuring Web services data sources

Correction

The Web services wrapper is supported on Sun and Linux. The name of the library file for Sun and Linux is libdb2ws.so. The name of the path is /opt/IBM/db2/V8.1/lib

Topic title: What is XML?

Location in the DB2 Information Center

Configuring —> Federated systems and data sources —> Configuring data sources —> Configuring XML data sources

Correction

Vendor-defined characters are retrieved as alternative characters. The XML parser extracts the data from the XML document and converts it into UTF-16, the native code page of the XML parser. The XML wrapper then extracts the data from the XML parser. The wrapper receives the data in UTF-16 and converts it into the federated database code page.

The XML wrapper does not support vendor-defined characters in the XML instance documents.

Topic title: Planning for SQL replication

Location in the DB2 Information Center:

Configuring –> Replication and event publishing –> Configuring for SQL replication

Correction

In the section entitled "Planning for conflict detection," the referential integrity rules are not restricted for conflict detection.

Topic title: Setting up WebSphere MQ for Q replication and event publishing—Overview

Location in the DB2 Information Center:

Configuring —> Replication and event publishing —> Configuring for Q replication and event publishing

Correction

Q replication and event publishing require WebSphere MQ Express Version 5.3 Fix Pack 05 (CSD05).

Topic title: WebSphere MQ objects required for bidirectional or peer-to-peer replication (two remote servers)

Location in the DB2 Information Center:
Configuring —> Replication and event publishing —> Configuring Q replication and event publishing —> Setting up WebSphere MQ —> WebSphere MQ objects required for Q replication and event publishing

Correction
For bidirectional or peer-to-peer replication, only one transmission queue and channel are needed at each server for each remote queue manager to which you need to send messages. The diagram in this topic should also show only one transmission queue and channel in each direction between the two servers.

Topic title: Configuring servers for Q replication and event publishing—Overview

Location in the DB2 Information Center:
Configuring –> Replication and event publishing –> Configuring for Q replication and event publishing –> Configuring servers for Q replication and event publishing

Correction
New functionality: Q replication and event publishing now support the WebSphere MQ Client. Q replication and event publishing programs previously ran as 32-bit but now they run as 64-bit on 64-bit platforms for all platforms except HP-UX, z/OS, and Windows. You can run the programs as 64-bit by using the MQ Client. If you want to continue running the programs as 32-bit on 64-bit platforms, see the technote entitled "Running Q replication and event publishing as 32-bit on 64-bit platforms with the MQ Client" at the DB2 Information Integrator support Web site: www.ibm.com/software/data/integration/db2ii/support.html.

Topic title: Configuring the source database to work with the Q Capture program (Linux, UNIX, Windows)

Location in the DB2 Information Center:
Configuring —> Replication and event publishing —> Configuring Q replication and event publishing —> Configuring servers —> Configuring databases for Q replication and event publishing (Linux, UNIX, Windows)

Correction
The command in step 1 of the procedure is incorrect. To check your current setting to see what the LOGRETAIN parameter is set to, enter the following command:

```
db2 get database configuration for database
```

Where *database* is the database that contains your source data.

Administering

Topic title: Starting bidirectional or peer-to-peer replication with two servers

Location in the DB2 Information Center:
Administering –> Replication and event publishing –> Administering Q replication and event publishing –> Setting up Q

replication –> Setting up multidirectional replication from sources to targets –> Starting multidirectional replication

Correction

When you are starting bidirectional replication for the first time, you must cold start the Q Capture programs at both servers before you start the Q Apply programs at either server. If you did not cold start both Q Capture programs and the second Q Capture program did not receive the signal from the Q Apply program, update the STATE column in the IBMQREP_SUBS table to 'I '(which sets the Q subscription to inactive state) and do a manual CAPSTART signal.

Topic title: Registering tables and views as SQL replication sources

Location in the DB2 Information Center:

Administering –> Replication and event publishing –> Administering SQL replication –> Setting up SQL replication

Correction

Restrictions: For non-DB2 UDB sources for SQL replication:
- COMMIT_COUNT must be null.
- Referential integrity is not supported.

Topic title: Utilities used for automatic load option for Q

Location in the DB2 Information Center:

Administering –> Replication and event publishing –> Administering Q replication and event publishing –> Setting up Q replication –> Options for loading target tables for Q replication –> Automatic load option

Correction

Table 16 is incorrect. Q replication and event publishing does not support DB2 for OS/390 Version 6. For DB2 UDB z/OS for Version 8 and DB2 for OS/390 Version 7, manual load or no load are the only load options that are available.

Topic title: Creating Q subscriptions for peer-to-peer replication with two servers

Location in the DB2 Information Center:

Administering –> Replication and event publishing –> Administering Q replication and event publishing –> Setting up Q replication –> Setting up multidirectional replication from sources to targets –> Creating Q subscriptions for peer-to-peer replication

Correction

Restrictions:
- If the z/OS source table includes a LONG VARCHAR column type, that table cannot participate in a peer-to-peer replication. For peer-to-peer replication, the Replication Center or ASNCLP must add two columns to the source table. DB2 UDB for z/OS does not allow the replication administration tools to add columns to a table that includes a LONG VARCHAR column.
- For peer-to-peer and bidirectional configurations, do not use the import utility. The import utility logs the inserts and, therefore, the inserts will be recaptured.
- In peer-to-peer and bidirectional replication, you must use the same constraints in both the source and target.

Topic title: Nickname statistics update facility - overview

Location in the DB2 Information Center:
Administering —> Federated systems —> Updating nickname statistics

Correction 1 of 2
The introductory sentence for the first list of statistics should read: You can update the following statistics for relational nicknames if the wrapper retrieved them when the nickname was created.

Correction 2 of 2
The introductory sentence for the second list of statistics should read: You can update the following statistics for nonrelational nicknames and relational nicknames that the wrapper cannot retrieve statistics for when the nickname was created.

Addition
Excel is added to the list of data sources that support nickname statistics.

Topic title: Event Publishing

Location in the DB2 Information Center:
Administering —> Replication and event publishing —> Administering Q replication and event publishing

Correction
Tip: The asnqwxml sample program provides an example of a Web-based application that consumes the XML messages that the Q Capture program publishes. The sample demonstrates how to use XML publications in a business scenario.

Topic title: Q subscriptions

Location in the DB2 Information Center:
Administering —> Replication and event publishing —> Administering Q replication and event publishing —> Objects for Q replication and event publishing

Correction
Recommendation: In a multiple partitioned database environment, if there are tables with referential integrity relationships, ensure that both the parent and child are on the same partition. If the parent and child are in a referential integrity relationship and are not on the same partition, you might see referential integrity issues at the target that might result in SQLSTATES 23504, 23001, or 23503 (which correspond to SQLCODES 530 and 532).

Topic title: Starting bidirectional or peer-to-peer replication with two servers and Starting a peer-to-peer group with three or more servers

Location in the DB2 Information Center:
Administering —> Replication and event publishing —> Administering Q replication and event publishing —> Setting up Q replication —> Setting up multidirectional replication from sources to targets –> Starting multidirectional replication

Correction

You must start the Q Capture programs before you start the Q Apply programs for the Q subscriptions in bidirectional or peer-to-peer replication.

Topic title: Changing attributes of Q subscriptions

Location in the DB2 Information Center:
Administering —> Replication and event publishing —>
Administering Q replication and event publishing —> Changing a
Q replication environment

Correction
You can change attributes for Q subscriptions that participate in
unidirectional, bidirectional, or peer-to-peer replication. In addition
to the attributes that are listed, you can change the following
attributes for Q subscriptions:

- source_colname
- target_colname
- target_colno
- description
- error_action
- has_loadphase
- load_type
- src_nickname
- src_nickname_owner
- sendq
- recvq
- target_name
- target_owner
- target_type
- conflict_action
- before_values
- conflict_rule
- changed_cols_only
- is_key

For instructions on how to change these additional attributes for Q
subscriptions, see the technote entitled "Changing attributes for Q
replication and event publishing objects" at the DB2 Information
Integrator support Web site:
www.ibm.com/software/data/integration/db2ii/support.html.

Topic title: Adding columns to existing Q subscriptions

Location in the DB2 Information Center:
Administering —> Replication and event publishing —>
Administering Q replication and event publishing —> Changing a
Q replication environment

Correction

Perform an ALTER TABLE ADD COLUMN operation for the
source table in the same transaction as the insert of the ADDCOL
signal. You can perform the two SQL statements in any order.

Topic title: Changing attributes of replication queue maps

Location in the DB2 Information Center:
> Administering —> Replication and event publishing —>
> Administering Q replication and event publishing —> Changing a
> Q replication environment

Correction
> In addition to the attributes that are listed, you can change the
> following attributes for replication queue maps:
>
> - sendq
> - recvq
> - description
>
> For detailed instructions on how to change these additional
> attributes for replication queue maps, see the technote entitled
> "Changing attributes for Q replication and event publishing
> objects" at the DB2 Information Integrator support Web site:
> www.ibm.com/software/data/integration/db2ii/support.html.

Topic title: Changing attributes of XML publications

Location in the DB2 Information Center:
> Administering —> Replication and event publishing —>
> Administering Q replication and event publishing —> Changing
> an event publishing environment

Correction
> In addition to the attributes that are listed, you can change the
> following attributes for XML publications:
>
> - before_values
> - changed_cols_only
> - description
> - sendq
> - src_colname
> - is_key
> - topic
>
> For detailed instructions on how to change these additional
> attributes for XML publications, see the technote entitled
> "Changing attributes for Q replication and event publishing
> objects" at the DB2 Information Integrator support Web site:
> www.ibm.com/software/data/integration/db2ii/support.html.

Topic title: Adding columns to existing XML publications

Location in the DB2 Information Center:
> Administering —> Replication and event publishing —>
> Administering Q replication and event publishing —> Changing
> an event publishing environment

Correction

> Perform an ALTER TABLE ADD COLUMN operation for the
> source table in the same transaction as the insert of the ADDCOL
> signal. You can perform the two SQL statements in any order.

Topic title: Changing attributes of publishing queue maps

Location in the DB2 Information Center:
>Administering —> Replication and event publishing —>
>Administering Q replication and event publishing —> Changing
>an event publishing environment

Correction
>In addition to the changeable attributes that are listed, you can
>change the following attributes for publishing queue maps:
>
>- sendq
>
>- message_format
>
>For detailed instructions on how to change these additional
>attributes for publishing queue maps, see the technote entitled
>"Changing attributes for Q replication and event publishing
>objects" at the DB2 Information Integrator support Web site:
>www.ibm.com/software/data/integration/db2ii/support.html.

Topic title: Running and saving SQL scripts from the Replication Center

Location in the DB2 Information Center:
>Administering —> Replication and event publishing —>
>Administering Q replication and event publishing —> Running
>SQL scripts and commands from the Replication Center

Correction

>**Prerequisites**
>>Before you run an SQL script from a command line, you
>>must connect to the server and specify a user ID and
>>password for the server. If you will connect to the database
>>from the command line, specify the CONNECT TO
>>statement as follows:
>>
>>```
>>CONNECT TO database USER user_ID USING password
>>```
>>
>>Where *database* is the name of the server that you are
>>connecting to, *user_ID* is a valid user ID, and *password* is
>>the password for that user ID. If you specify the
>>CONNECT TO statement in a file, use a semicolon (;) at
>>the end of the statement. Here is an example:
>>
>>```
>>CONNECT TO database USER user_ID USING password;
>>```
>>
>>*database* is the name of the server that you are connecting
>>to, *user_ID* is a valid user ID, and *password* is the password
>>for that user ID.

Topic title: Considerations for general data types for Q replication and event publishing

Location in the DB2 Information Center:
>Administering —> Replication and event publishing —>
>Administering Q replication and event publishing —> Data type
>considerations

Correction
>Columns with the GRAPHIC data type at the source and target
>might not match when you use the tdiff utility to check that the
>source and target tables are the same. DB2 Universal Database
>columns with the GRAPHIC data type have blank padding after
>the graphic data. This padding might be single-byte or double-byte
>spaces, depending on the code page that the database was created

in. This padding potentially can cause data to not match between the source and the target tables, especially if the source and target tables are in different code pages. This padding applies only to GRAPHIC data types and not other graphic data types such as VARGRAPHIC or LONG VARGRAPHIC.

To compare columns with GRAPHIC data types, you must remove the blank padding in the data before you compare the source and target tables by using the DB2 scalar function:

```
rtrim(<column>)
```

This function eliminates the code page differences for single-byte or double-byte spaces and ensures that the tdiff utility compares the GRAPHIC data in a consistent manner.

Topic title: Starting Q Capture program

Location in the DB2 Information Center:
Administering —> Replication and event publishing —> Q replication and event publishing —> Using system services to operate the replication programs—> Operating the Q replication and event publishing programs by using system services (z/OS) —> Using JCL

Correction

Specifying the directory for the transaction log on z/OS

You must specify a path where the data set that contains the transaction log will be stored. The CAPTURE_PATH parameter contains the path where the replication and publishing programs send their transaction logs. If you do not specify the CAPTURE_PATH parameter, the Q Capture program writes its log files to the home directory of the user who submits the JCL to start the program.

You can specify the path for the log file in JCL or in the IBMQREP_CAPPARMS table.

Procedure:

To specify the path in JCL:

Specify the CAPTURE_PATH parameter in the PARM field of the JCL statement that will start the Q Capture program.

- Here is an example that you can use to specify the CAPTURE_PATH parameter:

```
// PARM='/CAPTURE_SERVER=DSN7  CAPTURE_PATH=//JAYQC
//          LOGSTDOUT capture_schema=JAY'
```

In this example, the Q Capture program writes its log files to the USER1.JAYQC.D7DP.JAY.QCAP.LOG file. USER1 is the user who submits the JCL.

- If you want the data set for the log to have a specific high level qualifier, use this example:

```
// PARM='/capture_server=DSN7 capture_schema=JAY
//          CAPTURE_PATH=//''OEUSR01'
```

Now the Q Capture program writes its log files to the OEUSR01.DSN7.JAY.QCAP.LOG file.

- If you want to specify the path to SYSADM.XYZ, use one of the following examples:

```
//  PARM='/CAPTURE_server=DSN7  Capture_path=//''SYSADM.XYZ
//              capture_schema=JAY'

//  PARM='/CAPTURE_server=DSN7 capture_schema=JAY
//               capture_PATH=//''SYSADM.XYZ'
```

Ensure that the path name does not exceed the 44 character limit for MVS™ data sets. The user ID that runs this JCL must be authorized to write to the above data set.

To specify the path in the IBMQREP_CAPPARMS table:

- Use the following example:

```
INSERT INTO JAY.IBMQREP_CAPPARMS
 (qmgr, restartq, adminq, startmode, memory_limit, commit_interval,
 autostop,monitor_interval,monitor_limit, trace_limit, signal_limit,
 prune_interval, sleep_interval, logreuse, logstdout, term,
 capture_path, arch_level )
  VALUES
  ( 'CSQ1', 'IBMQREP.ASN.RESTARTQ',
            'IBMQREP.ASN.ADMINQ',
          , 'WARMSI', 32, 500,
'N',300, 10080, 10080, 10080, 300, 5000, 'N', 'N', 'Y',
 '//JAYQC', '802' ) ;
```

- If you want the data set for the log to have a specific high level qualifier, use this example:

```
INSERT INTO JAY.IBMQREP_CAPPARMS
(qmgr, restartq, adminq, startmode, memory_limit, commit_interval,
autostop,monitor_interval,monitor_limit, trace_limit, signal_limit,
prune_interval, sleep_interval, logreuse, logstdout, term,
capture_path, arch_level )
 VALUES
( 'CSQ1', 'IBMQREP.ASN.RESTARTQ',
          'IBMQREP.ASN.ADMINQ',
        , 'WARMSI', 32, 500,
'N',300, 10080, 10080, 10080, 300, 5000, 'N', 'N', 'Y',
 '//''OEUSR01','802');
```

- To specify the path to SYSADM.XYZ, use this example:

```
INSERT INTO JAY.IBMQREP_CAPPARMS
 (qmgr, restartq, adminq, startmode, memory_limit, commit_interval,
 autostop,monitor_interval,monitor_limit, trace_limit, signal_limit,
 prune_interval, sleep_interval, logreuse, logstdout, term,
 capture_path, arch_level )
  VALUES
  ( 'CSQ1', 'IBMQREP.ASN.RESTARTQ',
            'IBMQREP.ASN.ADMINQ',
          , 'WARMSI', 32, 500,
 'N',300, 10080, 10080, 10080, 300, 5000, 'N', 'N', 'Y',
  '//''SYSADM.XYZ','802');
```

Topic title: Tdiff: Table difference utility

Location in the DB2 Information Center:

Administering —> Replication and event publishing —> Administering Q replication and event publishing —> Detecting and repairing differences between source and target tables

Correction

When you run the **asntdiff** command, you specify an SQL WHERE clause that uniquely identifies the Q subscription or subscription set member:

Q replication

The WHERE clause identifies a row in the IBMQREP_SUBS control table at the Q Capture server, based on the value of the SUBNAME column, for example:

```
                        where="subname = 'my_qsub'"
```

SQL replication

The WHERE clause identifies a row in the IBMSNAP_SUBS_MEMBR table at the Apply control server, based on the value of the SET_NAME column. For example:

```
where="set_name = 'my_set' and source_table='EMPLOYEE'"
```

You might need to use more predicates in the WHERE clause to uniquely identify the subscription set member. For example, you might need to add the APPLY_QUAL, the SOURCE_OWNER, the TARGET_OWNER, or the TARGET_TABLE column from the IBMSNAP_SUBS_MEMBR table to the clause.

Difference table:

The difference table uses three identifiers that indicate the operation that is needed to change the target table so that it matches the source table:

- D (delete): Indicates that a row with the key value exists at the target and not at the source.
- U (update): Indicates that rows with the same key value exist at both the source and target, but at least one non-key column is different at the target.
- I (insert): Indicates that a row with the key value exists at the source and not at the target.

A value of ? 1 indicates that there is an invalid character in one or more source columns.

A value of ? 2 indicates that there is an invalid character in one or more target columns.

Example:

The following list of values is returned by comparing an EMPLOYEE table at the source with a target copy of the same table. The key column for replication is the employee number, EMPNO:

```
DIFF      EMPNO
U 2         000010
I 2         000020
I 2         000040
D 2         000045
I 2         000050
D 2         000055
```

The second row in the example shows that a row with the key value 000010 exists at both the source and target tables, but at least one non-key column at the target has a different value. The next two rows show that rows with the key values 000020 and 000040 exist only at the source. The fourth row shows that a row with the key value 000045 exists only at the target.

In some cases, differences between the source and target tables are intentional, for example, if you use a search condition in Q replication to filter which rows are replicated. The **asntdiff** command does not show differences between source and target tables that are a result of the following options for Q replication or SQL replication:

Q replication
- A search condition that limits which rows are replicated
- The option to suppress replication of delete operations from the source

SQL replication
- A subscription set member predicate
- An SQL expression used to create the target column contents

Topic title: Considerations for rebinding packages and plans for Q replication and event publishing

Location in the DB2 Information Center:
Administering —> Replication and event publishing —> Q replication and event publishing —> Maintaining a Q replication and event publishing environment —> Maintaining control tables

Correction

For z/OS

The Q Capture program, Q Apply program, and Common packages are bound automatically. You can use the z/OS sample ASNQBNDL to bind ASNCOMMON, ASNQCAPTURE, ASNQAPPLY, and ASNMON packages at a DB2 Universal Database subsystem.

DB2 Information Integrator Replication for z/OS specifies VERSION AUTO when it precompiles. So, DB2 Universal Database for z/OS automatically frees any packages that are older than the two most current versions.

Topic title: Customizing the XML Metadata Registry home page

Location in the DB2 Information Center:
Administering —> Administering XML metadata with the XML Metadata Registry

Correction
The sentence in step 1 about where the registry home page is displayed has changed to include the login page:

The information in the About section of the home page is also displayed on the login page.

Topic title: Enabling version control for objects in the XML Metadata Registry

Location in the DB2 Information Center:
Administering —> Administering XML metadata with the XML Metadata Registry —> Controlling versions

Correction
You must have full access to an object to enable version control.

Topic title: Access groups in the XML Metadata Registry

> **Location in the DB2 Information Center:**
>> Administering —> Administering XML metadata with the XML Metadata Registry —> Managing user registration, object access, and security

> **Correction**
>> You must have Write access to an access group to add or remove members from it.

Topic title: Disabling access to the XML Metadata Registry for users or access groups

> **Location in the DB2 Information Center:**
>> Administering —> Administering XML metadata with the XML Metadata Registry —> Managing user registration, object access, and security

> **Correction**
>> When you click **Enable** or **Disable** to grant or revoke access to the XML Metadata Registry, the change occurs immediately. If you are logged in as the user that you disable access for, the Login page is displayed when you click **Disable**. Otherwise, the notebook for the user or access group that you are editing does not change.

Topic title: Enabling automatic registration for the XML Metadata Registry

> **Location in the DB2 Information Center:**
>> Administering —> Administering XML metadata with the XML Metadata Registry —> Managing user registration, object access, and security

> **Correction**
>> This topic incorrectly states that the **Enable the Register a user feature** link is in the About this Registry section of the home page. The link is in the Using the Registry section of the home page if automatic registration is not enabled. You might need to scroll down on the section to see the link.

Topic title: Disabling automatic registration for the XML Metadata Registry

> **Location in the DB2 Information Center:**
>> Administering —> Administering XML metadata with the XML Metadata Registry —> Managing user registration, object access, and security

> **Correction**
>> This topic incorrectly states that the **Disable the Register a user feature** link is in the About this Registry section of the home page. The link is in the Using the Registry section of the home page if automatic registration is enabled. You might need to scroll down on the section to see the link.

Monitoring

Topic title: The Replication Alert Monitor

> **Location in the DB2 Information Center:**
>> Monitoring —> Replication and event publishing

> **Correction**

>> **Restrictions on z/OS servers**

- The Replication Alert Monitor can send e-mail notifications by using an SMTP server, but cannot use the ASNMAIL exit routine to handle notification.
- A monitor that runs from a z/OS server can monitor status only for Capture, Q Capture, Apply and Q Apply programs that run on z/OS servers. The Replication Alert Monitor cannot check Linux, UNIX, and Windows environments for the following alert conditions when it runs from a z/OS server:
  - CAPTURE_STATUS
  - QCAPTURE_STATUS
  - APPLY_STATUS
  - QAPPLY_STATUS

  However, a monitor that runs on a Linux, UNIX, or Windows server can monitor a Capture, Q Capture, Apply or a Q Apply program that runs on z/OS. Install the DB2 Administration Server on the system where the replication programs will run. For example, MONITOR1, which runs on SERVER_LINUX1 can monitor a Q Apply program that runs on SERVER_ZOS1 if the DB2 Administration Server is installed on SERVER_ZOS1.

Topic title: Alert Conditions for the Replication Alert Monitor

Location in the DB2 Information Center:
Monitoring –> Replication and event publishing –> Alert Conditions and notifications for the Replication Alert Monitor

Correction
To enable the QAPPLY_SPILLQDEPTH and the QAPPLY_QDEPTH alert conditions for z/OS platforms, you must apply the z/OS fix for APAR PQ94364.

Topic title: Setting up the Replication Alert Monitor

Location in the DB2 Information Center:
Monitoring —> Replication and event publishing

Correction
Linux, UNIX, and Windows: To use the QAPPLY_QDEPTH and QAPPLY_SPILLQDEPTH alert conditions, copy the following library from the sqllib/bin/routine directory to the sqllib/function/unfenced directory:

Linux and UNIX
asnqmon

Windows
asnqmonasnqmon.dll

Developing

Topic title: Why develop a wrapper?

Location in the DB2 Information Center:
Developing —> Connections (wrappers) to customer data sources —> Developing wrappers – overview

Correction

In the figure entitled "Query to request IDs of molecules similar to molecules with a result > 0.8 in a stomach experiment", the last line E.MOLE_ID should be E.MOLECULE_ID.

Topic title: Wrapper development kit

Location in the DB2 Information Center:

Developing —> Connections (wrappers) to customer data sources —> Developing wrappers – overview

Correction

The Windows environment variable %DB2PATH% should be %DB2TEMPDIR%, which represents the environment variable that is used to specify a temporary directory path on Windows.

Topic title: Control flow for query planning

Location in the DB2 Information Center:

Developing —> Connections (wrappers) to customer data sources —> Developing and documenting wrappers —> Overview of data flows —> Control flows for processes

Correction

Some of the step numbers that are referenced from the process are incorrect:

- Step 10 should read: "The wrapper determines whether a data source can compute the value of the head expression that is represented by Request_Exp. The wrapper repeats step 9 [not step 12] recursively descending the Request_Exp tree until a the wrapper makes a decision."

- Step 12 should read: "Repeat steps 8-11 [not steps 9-12] for each additional head expression in the request."

- Step 17 should read: "Repeat steps 14-16 [not steps 15-17] for each additional predicate in the request."

Topic title: Dynamic database queries that use the Web services provider

Location in the DB2 Information Center:

Developing —> Web services and federated Web applications —> Developing applications that use the Web services provider —> Using dynamic database queries

Correction

The supported functions for dynamic query services should include the delete function:

With dynamic query services, you can dynamically build and submit queries at run time that select, insert, update, and delete application data, and call stored procedures rather than run queries that are predefined at deployment time.

Topic title: Starting and stopping the application server for DB2 in Information Integrator

Location in the DB2 Information Center:

Developing —> Web services and federated Web applications —> Deploying and testing Web applications

Correction

You can install and configure the Web services provider, the Web services provider applications (including the Web services provider

samples), and JDBC providers and enable and disable a trace on an application server for DB2 by using a script that is included in the \bin directory of the dxxworf.zip file. The application server for DB2 must be running.

Prerequisite

To install and configure the Web services provider runtime on an application server for a DB2 Universal Database system, follow installation instructions steps 1– 4 in the topic "Installing or migrating WORF to work with WebSphere Application Server Version 5 or later for Windows and UNIX."

You can use the worf_eas_admin.jacl script in these ways:

* **Install the Web services provider samples on an application server for DB2**

  The following examples assume that you are familiar with deploying the Web services provider examples, as described in topic "Deploying WORF examples on WebSphere Application Server Version 5.1 or later for Windows and UNIX."

  – Optional: Use the following command to install a JDBC provider that is used by the Web services provider samples Web application:

  ```
  $appserv_install_dir\bin\wsadmin.bat -f worf_eas_admin.jacl
    configureJDBC
      -name db2jdbc
      -classpath "C:\\SQLLIB\\java\\db2java.zip"
      -iClassName COM.ibm.db2.jdbc.app.DB2Driver
  ```

  From the command line in Windows, use either \\ or /as the directory delimiter.

  – Optional: Use the following command to install the Web services provider samples:

  ```
  $appserv_install_dir\bin\wsadmin.bat -f worf_eas_admin.jacl
    installApp
      -warPath "C:\\worf\\lib\\axis-services.war"
      -contextRoot services
      -appName WorfAxis
  ```

  From the command line in Windows, use either \\ or /as the directory delimiter.

* **Install user-created Web services provider Web applications on an application server for DB2**

  You can issue the following command to install and configure Web Services provider Web applications:

  – worf_eas_admin.jacl (use one of the following six optional keywords)

  > installApp <parameters>
  > configureJDBC <parameters>
  > uninstallApp <parameters>
  > removeJDBC <parameters>
  > enableTrace
  > disableTrace

installApp

Installs a Web application that is based on the following parameters:

warPath

The path to the Web archive (WAR) package.

contextRoot

The context root of the application in the application server for DB2.

appName

The name of the application in the application server for DB2. The name must not contain any blank characters.

The following example shows the installApp mode:

```
$appserv_install_dir\bin\wsadmin.bat -f worf_eas_admin.jacl
   installApp
   -warPath "C:\\My Files\\WORF\\axis-services.war"
   -contextRoot services
   -appName WorfAxis
```

From the command line in Windows, use either \\ or /as the directory delimiter.

configureJDBC

Configures a JDBC provider. The following parameters are required:

name   The name of the JDBC provider in the application server for DB2.

classPath

The CLASSPATH to the Java™ archive (JAR) packages of the JDBC driver.

iClassName

The name of the implementation class.

The following example shows the configureJDBC mode:

```
$appserv_install_dir\bin\wsadmin.bat -f worf_eas_admin.jacl
  configureJDBC
   -name db2jcc
   -classPath "C:\\SQLLIB\\java\\db2jcc.jar;
       C:\\SQLLIB\\java\\db2jcc_license_cu.jar;
       C:\\SQLLIB\\java\\db2jcc_license_cisuz.jar"
   -iClassName com.ibm.db2.jcc.DB2Driver
```

From the command line in Windows, use either \\ or /as the directory delimiter.

uninstallApp mode

Removes an installed application. You must provide the name of the application in the application server for DB2, as in the following example:

```
$appserv_install_dir\bin\wsadmin.bat -f worf_eas_admin.jacl
  uninstallApp -appName WorfAxis
```

removeJDBC

Removes a JDBC provider. You must provide the name of the JDBC provider in the application server for DB2, as in the following example:

```
$appserv_install_dir\bin\wsadmin.bat -f worf_eas_admin.jacl
   removeJDBC -name db2jcc
```

enableTrace

Enables the trace for the application server for DB2. The application server for DB2 must be running. The following example enables the trace:

```
$appserv_install_dir\bin\wsadmin -f worf_eas_admin.jacl
 enableTrace
```

disableTrace

Disables the trace for the application server for DB2. The application server for DB2 must be running. The following example disables the trace:

```
$appserv_install_dir\bin\wsadmin -f worf_eas_admin.jacl
 disableTrace
```

Topic title: Generating deployment descriptors

Location in the DB2 Information Center:

Developing —> Web services and federated Web applications —> Deploying and testing Web applications

Correction 1 of 3

The following parameters are new for the Dadx2Dd deployment descriptor class:

-t   Optional: Specifies the SOAP encoding style. The valid values are document for document style and RPC for remote procedure call style. The default value is RPC. The style that you select must match the style that is specified in the group.properties file of the group that contains the DADX file.

-c   Specifies the context path of the Web service. The valid values are any string values that are accepted by the Web server. The default is services. The context path must match the path that is used for the installation of the Web services provider Web application on the application server.

-u   Specifies the target namespace that is used in the generated WSDL document and the XML schema files (XSD) for Web services in the group. The default is http://tempuri.org/<context_path>. The target namespace must match the namespace that you specified in the group.properties file of the group containing the DADX file.

The default for the following parameter is changed:

-s   If the soap-engine parameter is not specified in the web.xml file, the default soap engine is Apache SOAP.

Correction 2 of 3

All references to the Java package com.ibm.etools.webservice.rt.dadx should be com.ibm.etools.webservice.rt.dadx2dd.

The correction to the example is:

```
java com.ibm.etools.webservice.rt.dadx2dd.Dadx2Dd -r ZipCity.dadx
  -p /travel
   -n /dxx_travel -i classes\groups\dxx_travel\ZipCity.dadx
    -o classes\dds\dxx_travel\ZipCity.isd -c services -t rpc
```

Correction 3 of 3

The dds.xml example is updated to specify an Apache service provider, and an Apache DxxMappingRegistry. The changes in the example are highlighted in the following example:

```
<isd:service xmlns:isd='http://xml.apache.org/xml-soap/deployment'
  id='http://tempuri.org/travel/ZipCity.dadx'>
   <isd:provider
    type='com.ibm.etools.webservice.rt.framework.apache.ApacheServiceProvider'
    scope='Request'
    methods='findCityByZipCode insertZipCodeAndCity
        updateCityForZipCode deleteZipCode'>
    <isd:java class='com.ibm.etools.webservice.rt.dxx.DxxService'/>
    <isd:option key='group.name' value='/dxx_travel'/>
    <isd:option key='group.path' value='/travel'/>
    <isd:option key='group.class.name'
            value='com.ibm.etools.webservice.rt.dxx.DxxGroup'/>
   </isd:provider>
 <isd:faultListener>org.apache.soap.server.DOMFaultListener
 </id:faultListener>
 <isd:mappings
    defaultRegistryClass=
    'com.ibm.etools.webservice.rt.dxx.apache.DxxMappingRegistry'/>
</isd:service>
...
...
</dds>
```

Topic title: Enabling tracing for the DB2 Web services provider—Apache Tomcat Version 4.0 or later Web application server

Location in the DB2 Information Center:
Developing —> Web services and federated Web applications —> Deploying and Testing Web applications —> Enabling Web services provider tracing

Correction
Specify the default level of trace information as INFO, instead of DEBUG, to maintain better performance. The following example is correct:

```
log4j.rootCategory=INFO, console, rollingFile
```

Topic title: The Web service consumer functions

Location in the DB2 Information Center:
Developing —> Web services and federated Web applications —> Installing the Web services consumer

Correction
You cannot secure the Web services consumer by encrypting messages through HTTPS.

Tuning

Topic title: Creating a federated materialized query table

Location in the DB2 Information Center:
Tuning —> Federated systems —> Materialized query tables

Addition
When you populate or update a user-maintained materialized query table by using the INSERT statement in a subselect statement, issue the following commands before you issue the INSERT statement so that the same materialized query table is not used as the source of the insert:

```
SET CURRENT REFRESH AGE 0
SET CURRENT MAINTAINED TABLE TYPE FOR OPTIMIZATION SYSTEM
```

Topic title: Data source specific restrictions for materialized query tables

Location in the DB2 Information Center:
Tuning —> Federated systems —> Materialized query tables

Addition 1 of 3
For federated queries, you must specify the -u option for the DB2 Design Advisor command, db2advis. If this option is not specified, the DB2 Design Advisor recommends only refresh immediate materialized query tables. Because the refresh immediate function is not supported for materialized query tables that reference nicknames, no recommendations are provided by the DB2 Design Advisor.

Addition 2 of 3
If the materialized query table recommendation that is provided by the DB2 Design Advisor does not comply with the materialized query table restrictions of the data source, the materialized query table can be created, but it cannot be refreshed. To prevent this situation from occurring, follow the recommendations provided by the DB2 Design Advisor for materialized query tables and ensure that the recommended materialized query table definition is acceptable to the data source.

Addition 3 of 3
In a multi-partitioned environment, the DB2 Design Advisor does not provide materialized query table recommendations that involve nicknames.

Reference

Topic title: Wrapper class

Location in the DB2 Information Center:
Reference –> APIs –> Wrapper development Java APIs –> Wrapper classes

Correction

destroy method

Purpose
Destroy this wrapper object and free any resources that are associated with it. The default implementation is no action. This method is called by the federated server before the wrapper object is removed so that the wrapper can free any allocated resources. To free wrapper-specific resources, you can implement this method in the wrapper-specific subclasses UnfencedGenericWrapper and FencedGenericWrapper.

Syntax
```
protected void destroy()
                throws java.lang.Exception
```

Parameters
None.

Return value
None.

Throws

An Exception object if the processing fails.

Topic title: UnfencedWrapper class

Location in the DB2 Information Center:

Reference –> APIs –> Wrapper development Java APIs –> Wrapper classes

Correction

getFencedWrapperClass method

Purpose

Retrieve the name of the class that must be loaded for the fenced part of the wrapper. The class name is specified as the value of the FENCED_WRAPPER_CLASS wrapper option.

Syntax

```
public final java.lang.String
getFencedWrapperClass(WrapperInfo wrapperInfo)throws
WrapperException
```

Parameters

wrapperInfo

The WrapperInfo object that stores the wrapper catalog information.

Return value

The name of the FencedWrapper subclass or null if no class name was specified.

Throws

A WrapperException object if the processing fails.

Topic title: Nickname class

Location in the DB2 Information Center:

Reference –> APIs –> Wrapper development Java APIs –> Nickname classes –> Nickname class

Correction

destroy method

Purpose

Destroy this nickname object and free any resources associated with it. The default implementation does nothing. This method is called by the federated server before the nickname object is removed and gives the wrapper a chance to free any resources allocated. If you need to free wrapper-specific resources, you can implement this method in the wrapper-specific subclass **UnfencedGenericNickname** and **FencedGenericNickname**.

Syntax

```
protected void destroy()
                    throws java.lang.Exception
```

Parameters

None.

Return value

None.

Throws
>> An Exception object if the processing fails.

Topic title: RemoteUser class

>**Location in the DB2 Information Center:**
>> Reference –> APIs –> Wrapper development Java APIs –> User
>> classes –> RemoteUser class

>**Correction**

>>**destroy method**

>>**Purpose**
>>> Destroy this user mapping object and free any resources
>>> associated with it. The default implementation does
>>> nothing. This method is called by the federated server
>>> before the user mapping object is removed and gives the
>>> wrapper a chance to free any resources allocated. If you
>>> need to free wrapper-specific resources, you can implement
>>> this method in the wrapper-specific subclass
>>> **UnfencedGenericNickname** and **FencedGenericNickname**.

>>**Syntax**
>>>
>>> ```
>>> protected void destroy()
>>>                 throws java.lang.Exception
>>> ```

>>**Parameters**
>>> None.

>>**Return value**
>>> None.

>>**Throws**
>>> An Exception object if the processing fails.

Topic title: Server class

>**Location in the DB2 Information Center:**
>> Reference –> APIs –> Wrapper development Java APIs –> Server
>> classes –> Server class

>**Correction**

>>**destroy method**

>>**Purpose**
>>> Destroy this server object and free any resources associated
>>> with it. The default implementation does nothing. This
>>> method is called by the federated server before the server
>>> object is removed and gives the wrapper a chance to free
>>> any resources allocated. If you need to free
>>> wrapper-specific resources, you can implement this method
>>> in the wrapper-specific subclass
>>> **UnfencedGenericNickname** and **FencedGenericNickname**.

>>**Syntax**
>>>
>>> ```
>>> protected void destroy()
>>>                 throws java.lang.Exception
>>> ```

>>**Parameters**
>>> None.

>>**Return value**
>>> None.

Throws
    An Exception object if the processing fails.

Topic title: IBMQREP_APPLYMON table

Location in the DB2 Information Center:
    Reference –> Control tables –> Replication and event publishing table structures –> Q replication and event publishing table structures –> Tables at the Q Apply server

Correction
    In the IBMQREP_APPLYMON table, the value in the OLDEST_TRANS column tells which transaction the Q Apply program has applied at the target based on the transaction commit time at the source server. After each monitor interval, the value in the OLDEST_TRANS column has the following meaning:

- The commit time at the source server of the oldest transaction that the Q Apply program has applied, if the Q Apply program is processing transactions.
- The latest heartbeat time, if no transactions are being processed and the heartbeat message arrived after the Q Apply program applied the oldest transaction.
- The value '1900-01-01-00.00.00.000000', if the Q Apply program has not received any messages (transaction or heartbeat messages).

Topic title: asnqcap: Starting a Q Capture program

Location in the DB2 Information Center:
    Reference —> Commands —> Q replication and event publishing commands

Correction

sleep_interval=*n*
    Specifies the number of milliseconds that a Q Capture program is idle after processing the active log and any transactions that remain in memory. The default is 5000 milliseconds (5 seconds).

Topic title: asntdiff: Comparing data in source and target tables

Location in the DB2 Information Center:
    Reference —> Commands —> Q replication and event publishing commands

Correction

The following examples show how to use the **asntdiff** command.

Example 1

In Q replication, to find the differences between a source and target table that are specified in a Q subscription named my_qsub, on a Q Capture server named source_db with a Q Capture schema of myschema:

```
asntdiff DB=source_db SCHEMA=myschema WHERE="subname = 'my_qsub'"
```

Example 2

In SQL replication, to find the differences between a source and target table that are specified in a subscription set named my_set,

with a target table named trg_table, on an Apply control server named apply_db, with an Apply schema of asn, and to name the difference table my_diff_table:

```
asntdiff DB=apply_db SCHEMA=asn WHERE="set_name = 'my_set'
 and target_table = 'trg_table'" DIFF=my_diff_table
```

Topic title: asntrep: Repairing differences between source and target tables

Location in the DB2 Information Center:
Reference —> Commands —> Q replication and event publishing commands

Correction

The following examples show how to use the **asntrep** command.

Example 1

In Q replication, to synchronize a source and target table that are specified in a Q subscription named my_qsub, on a Q Capture server named source_db, with a Q Capture schema of asn, and whose differences are stored in a table called q_diff_table:

```
asntrep DB=source_db SCHEMA=asn
 WHERE="subname = 'my_qsub'" DIFF=q_diff_table
```

Example 2

In SQL replication, to synchronize a source and target table that are specified in a subscription set called my_set, with a target table named trg_table, on an Apply control server named apply_db, with an Apply schema of asn, and whose differences are stored in a table called sql_diff_table:

```
asntrep DB=apply_db SCHEMA=asn WHERE="set_name = 'my_set'
 and target_table = 'trg_table'" DIFF=sql_diff_table
```

Topic title: Server options for federated systems

Location in the DB2 Information Center:
Reference —> Configuration parameters —> Federated systems and data sources

Addition 1 of 2
The following text should be added to the INFORMIX_LOCK_MODE server option description column:

If a deadlock or timeout error occurs when a federated server attempts to connect to an Informix data source, changing the lock mode setting on the federated server can often resolve the error. Use the ALTER SERVER statement to change the lock mode setting on the federated server.

For example:
```
ALTER SERVER TYPE informix VERSION 9 WRAPPER informix
          OPTIONS (ADD informix_lock_mode '60')
```

For more information about using the ALTER SERVER statement, see the topic "Altering server definitions and server options" in the DB2 Information Center. For more information about the INFORMIX_LOCK_MODE setting, see the topic "Server options for federated systems" in the DB2 Information Center.

If the deadlock or timeout error persists, contact IBM Software Support.

Addition 2 of 2

The following text should be added to the USE_CLOB_SEQUENCE server option description and default value columns:

If you specify a value of N, the data type is VARCHAR(32000). If you specify a value of Y, the data type is CLOB(5M). The default value is N, not Y.

Samples

Topic title: Samples to set up Q replication and event publishing (Linux, UNIX, Windows)

Location in the DB2 Information Center:

Samples —> Samples by functional area —> Q replication and event publishing samples —> Linux, UNIX, Windows

Correction

asnqdefq

This sample script shows which WebSphere MQ objects you need and how to create them. It includes commands to create queue managers and queues for two servers in unidirectional, remote replication. Use this sample script to quickly create and set up the WebSphere MQ objects for this type of Q replication environment.

Topic title: Samples to set up Q replication and event publishing (z/OS)

Location in the DB2 Information Center:

Samples —> Samples by functional area —> Q replication and event publishing samples —> z/OS

Correction

asnqdefq

This sample script shows which WebSphere MQ objects you need and how to create them). It includes commands to create queue managers and queues for two servers in unidirectional, remote replication. Use this sample script to quickly create and set up the WebSphere MQ objects for this type of Q replication environment.

New documentation

The following topic will be added to the *DB2 Information Integrator Installation Guide*.

Installing the IBM Tivoli License Manager (ITLM) file on your system

When you install DB2 Information Integrator Advanced Edition or DB2 Information Integrator Advanced Edition Unlimited, the IBM Tivoli License Manager (ITLM) file is added to your system and renamed after the product license key is registered. If the ITLM file is not installed on your system successfully, you must install the file manually.

When it is installed, the ITLM file is renamed from db2ii080200.sys to one of the following file names depending on your edition and operating system.

Table 19. Table of file names for the DB2 Information Integrator ITLM file

Operating system	Advanced Edition ITLM file name	Advanced Edition Unlimited ITLM file name
AIX	db2iiaex080200.sys	db2iiuex080200.sys
Linux	db2iiael080200.sys	db2iiuel080200.sys
HP-UX	db2iiaeh080200.sys	db2iiueh080200.sys
Solaris	db2iiaes080200.sys	db2iiues080200.sys
Windows	db2iiaew080200.sys	db2iiuew080200.sys

Procedure

To install the db2ii080200.sys ITLM file on your system, copy the <prod_cd>\license\db2ii080200.sys to one of the following directories, and rename it to the name that corresponds to your edition and operating system in Table 19. <prod_cd> is the root level of the DB2 Information Integrator product CD.

Windows
> x:\Program Files\IBM\DB2InformationIntegrator\V8\

> x is the drive where DB2 Information Integrator is installed.

UNIX   /opt/IBM/DB2InformationIntegrator/V8/

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 106-0032, Japan

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

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

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

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

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

IBM 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 contains 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 contains 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. 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 IBM's application programming interfaces.

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
DataJoiner
DB2
DB2 Connect
DB2 Universal Database
DRDA
Informix
iSeries
Lotus
MVS
OS/390
WebSphere
z/OS

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

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

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo 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.

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

Contacting IBM

To contact IBM customer service in the United States or Canada, call
1-800-IBM-SERV (1-800-426-7378).

To learn about available service options, call one of the following numbers:
- In the United States: 1-888-426-4343
- In Canada: 1-800-465-9600

To locate an IBM office in your country or region, see the IBM Directory of
Worldwide Contacts on the Web at www.ibm.com/planetwide.

Product information

Information about DB2 Information Integrator is available by telephone or on the
Web.

If you live in the United States, you can call one of the following numbers:
- To order products or to obtain general information: 1-800-IBM-CALL
  (1-800-426-2255)
- To order publications: 1-800-879-2755

On the Web, go to www.ibm.com/software/data/integration/db2ii/support.html.
This site contains the latest information about:
- The technical library
- Ordering books
- Client downloads
- Newsgroups
- Fix packs
- News
- Links to Web resources

Comments on the documentation

Your feedback helps IBM to provide quality information. Please send any
comments that you have about this book or other DB2 Information Integrator
documentation. You can use any of the following methods to provide comments:
- Send your comments using the online readers' comment form at
  www.ibm.com/software/data/rcf.
- Send your comments by e-mail to comments@us.ibm.com. Include the name of
  the product, the version number of the product, and the name and part number
  of the book (if applicable). If you are commenting on specific text, please include
  the location of the text (for example, a title, a table number, or a page number).

IBM ®

Printed in USA