IBM DB2 Information Integrator

Release Notes

Version 8.1, Fix Pack 5

IBM DB2 Information Integrator

Release Notes

Version 8.1, Fix Pack 5

Before using this information and the product it supports, be sure to read the general information under "Notices" on page 41.

Contents

About the release notes

The release notes contain the latest information about DB2® Information Integrator, Version 8.1, Fix Pack 5.

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

DB2 Information Integrator support on the Web

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

Known problems, limitations, and workarounds

The following information describes the currently known limitations, problems, and workarounds for DB2® Information Integrator Version 8.1, Fix Pack 5. Any limitations and restrictions might or might not apply to other releases of the product.

DB2 Information Integrator installation issues

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, you take one of the following actions:

- Install DB2 Information Integrator installer using the provided CDs.
- 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.

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

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

This message is incorrect, because the required product is already installed on your system. However, the DB2 Information Integrator wrappers are not installed. You must install the DB2 Information Integrator wrappers separately. You can use the db2_install script or the RedHat Package Manager to install the DB2 Information Integrator wrappers.

After you install the wrappers, you must reinstall the DB2 Information Integrator fix pack that was originally used to update your DB2 Information Integrator, Version 8.1 product. This reinstallation of the fix pack is necessary to ensure that all of your components are at the same fix pack level.

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

The rest of this topic describes how to install the DB2 Information Integrator wrappers. For information about installing DB2 Information Integrator fix packs, see the DB2 Information Integrator Support site at www.ibm.com/software/data/integration/db2ii/support.html.

Installing wrappers by using the db2_install script

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

Installing wrappers by using RedHat Package Manager

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

   Each DB2 Information Integrator product CD provides a file that lists the components that are available for installation. The component list is in a file named ComponentList.htm, and identifies which components are required, typical, and optional. Record the names of the components for use in step 5 of this procedure.

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

   ```
   rpm -ivh component_name
   ```

   component_name is the name of the component that you want to install. You can specify multiple components with the **rpm** command.

   For example, you can install the Signature for DB2 Information Integrator Nonrelatational Wrappers component (file name IBM_db2lssg81-

8.1.0.16.i386.rpm) and the Structured Files Data Sources component (file name IBM_db2lsstf81-8.1.0.16.i386.rpm) at the same time by using the following command:

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

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

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

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

Changing library permissions

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

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

Installing life sciences user-defined functions

The life sciences user-defined functions are basic and frequently used life sciences functions that simplify operations when you are querying life sciences data sources. The life sciences user-defined functions are part of DB2 Information Integrator and are installed with DB2 Information Integrator, Version 8.1., Fix Pack 5.

The LSGeneWise user-defined function is available for AIX®, HP–UX, Linux, and Solaris Operating Environment 32-bit operating systems. The other life sciences user-defined functions are available for Windows®, AIX, HP–UX, Linux, and Solaris Operating Environment 32-bit operating systems.

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

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

Windows
> sqllib\function\db2lsudfs.dll

AIX

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

HP-UX

- sqllib/function/libdb2lsudfs.sl
- sqllib/function/libdb2lsgenewise.sl

Solaris Operating Environment

- sqllib/function/libdb2lsudfs.so
- sqllib/function/libdb2lsgenewise.so

Linux

- sqllib/function/libdb2lsudfs.so
- sqllib/function/libdb2lsgenewise.so

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

The following life sciences user-defined functions are installed with DB2 Information Integrator, Version 8.1, Fix Pack 5:

- LSPep2AmbNuc
- LSPep2ProbNuc
- LSDeflineParse
- LSPatternMatch
- LSPrositePattern
- LSGeneWise (AIX, HP-UX, Linux, and Solaris Operating Environment only)
- LSMultiMatch
- LSMultiMatch3
- LSBarCode
- LSTransAllFrames
- LSRevNuc
- LSRevPep
- LSRevComp
- LSNuc2Pep

Prerequisites: You must install DB2 Information Integrator, Version 8.1 and the Life Sciences User-Defined Functions component of the nonrelational wrappers before you install DB2 Information Integrator, Version 8.1, Fix Pack 5. See the DB2 Information Integrator Support site at www.ibm.com/software/data/integration/db2ii/support.html for information about installing DB2 Information Integrator fix packs.

To install the life sciences user-defined functions:

1. Start the DB2 Information Integrator, Version 8.1 installation wizard.
2. Select the **Nonrelational wrappers** check box on the Product Selection window and click **Next**. Follow the instructions in the wizard to complete the installation. When you install the nonrelational wrappers, select the Life Sciences User-defined Functions component. You do not need to reinstall DB2 Universal Database or any wrappers that are already installed.
3. Install DB2 Information Integrator, Version 8.1, Fix Pack 5.
4. If you plan to use the LSGeneWise user-defined function, run the djxlinkLSGeneWise script to create the sqllib/function/libdb2lsgenewise.a library.

You can now register the life sciences user-defined functions. See the "Registering life sciences user-defined functions" topic in the DB2 Information Center for instructions on registering the life sciences user-defined functions.

Updated information for DB2 Universal Database documentation

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

Table 1. Updated file names for DB2 Universal Database documentation

Title and form number	PDF file name
CLI Guide and Reference, Volume 1, SC09-4849-01	db2l1x80
CLI Guide and Reference, Volume 2, SC09-4850-01	db2l2x80
Installation and Configuration Supplement, GC09-4837-00	db2iyx80

Migration issues

Accessing DB2 family remote data sources after migration

If you need to access a DB2 family remote data source after you migrate to DB2 Information Integrator from either DB2 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 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 issues

Data type restrictions

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

Unsupported data types

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

Table 2. Unsupported data source data types

Data source	Unsupported data types
Extended Search	DECIMAL

Table 2. Unsupported data source data types (continued)

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

For the data types that DB2 Information Integrator does not support, you can create a view at the data source that is based on the data source object, and create a nickname for the view. The view must not contain any columns that use the unsupported data types. 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 3. Write operation restrictions on data types

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

Table 3. Write operation restrictions on data types (continued)

Data source	Unsupported data types
Informix®	BLOB CLOB BYTE TEXT Exceptions: You can perform select, insert, update, and delete operations on TEXT data types if you alter the nickname column to the VARCHAR data type. You can perform select, 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
ODBC	SQL_LONGBINARY (length > 255) SQL_LONGVARCHAR (length > 255) SQL_WLONGVARCHAR (length > 255)

Oracle (NET8 wrapper only)	INTERVAL DAY (day_precision) TO SECOND
	    (fractional_seconds_precision)
	INTERVAL YEAR (year_precision) TO MONTH
	LONG
	LONG RAW
	NCHAR
	NVARCHAR2
	TIMESTAMP (fractional_seconds_precision) WITH TIMEZONE
	TIMESTAMP (fractional_seconds_precision)
	    WITH LOCAL TIME ZONE
	
	Exceptions: You can perform select, insert, update, and delete operations on LONG data types if you alter the nickname column to the VARCHAR data type. You can perform select, insert, update, and delete operations on LONG RAW data types if you alter the nickname column to the VARCHAR FOR BIT DATA data type.

Oracle (SQLNET wrapper only)	BLOB CLOB INTERVAL DAY (day_precision) TO SECOND (fractional_seconds_precision) INTERVAL YEAR (year_precision) TO MONTH NCHAR NVARCHAR2
Sybase (CTLIB wrapper only)	image text unichar univarchar

Table 3. Write operation restrictions on data types (continued)

Data source	Unsupported data types
Sybase (DBLIB wrapper only)	All data types. Write operations are not supported by the DBLIB wrapper.
Teradata	char (32673–64000) varchar (32673–64000) byte (32673-64000) varbyte (32673-64000) graphic (16337-32000) vargraphic (16337-32000)

Mapping to the federated GRAPHIC and VARGRAPHIC data types

Only the DB2 family and Teradata data sources can override the default data type mappings to map remote data types to the federated GRAPHIC and VARGRAPHIC data types. You use either the CREATE TYPE MAPPING statement or the ALTER NICKNAME statement to override the default data type mappings, depending on the circumstances in which you want the mapping to apply.

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.

Materialized query table issues

These issues apply only to DB2 Information Integrator configurations that use DB2 Universal Database, Version 8, Enterprise Server Edition in a partitioned database environment.

Limitation: Query optimization is supported only for aggregate materialized query tables. However, query optimization is supported for materialized query tables that are defined for a single nickname for queries that involve joins of local tables and nicknames.

Creating materialized query tables that reference nicknames

A materialized query table that references nicknames cannot be created if the CREATE TABLE statement is issued from a database partition without a federated catalog. If you try to create such a materialized query table, error SQL20208N with reason code 3 is issued.

Refreshing materialized query tables that reference nicknames

If you try to refresh a materialized query table that references a nickname, the REFRESH TABLE command returns error SQL0668N with reason code 6. To avoid this problem, define a materialized query table with the MAINTAINED BY USER option. Then, issue an INSERT statement with a subquery to populate the materialized query table with the proper query results.

Workaround example:

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

1. Define the ast2 materialized query table:

   ```
   CREATE TABLE ast2 AS (SELECT c2, COUNT(*) AS CNT FROM loopback_t2
   WHERE c2 > 10  GROUP BY c2) DATA INITIALLY DEFERRED REFRESH DEFERRED
   ENABLE QUERY OPTIMIZATION MAINTAINED BY USER;
   ```

2. Take the ast2 materialized query table out of the check pending state:

   ```
   SET INTEGRITY FOR ast2 ALL IMMEDIATE UNCHECKED;
   ```

3. Populate the ast2 materialized query table with valid query data:

   ```
   INSERT INTO ast2 (SELECT c2, COUNT(*) FROM loopback_t2
   WHERE c2 > 10 GROUP BY c2 );
   ```

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

LIKE predicate 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 by creating a function mapping for the LIKE(CHAR, *column_type* (*length*)) function so that the LIKE predicate is processed at the data source. For example:

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

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

CHAR columns that contain DBCS data

Predicates that compare CHAR columns that contain DBCS data might return incorrect results when the federated server and the data source use different blank padding characters. You might avoid the incorrect results by using the ALTER NICKNAME statement to change the local column type from CHAR to VARCHAR.

Entrez wrapper support

Using the Entrez wrapper with network proxies

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

Changing the EMAIL option of the Entrez wrapper

You cannot alter the Entrez wrapper after it has been created.

If you need to change the value of the EMAIL option for the Entrez wrapper, you must drop the wrapper and all of its subordinate server and nickname objects. Then, you must recreate the wrapper and all of its subordinate server and nickname objects, using the correct value for the EMAIL option. Finally, you must reprocess any other objects that are dependent on the Entrez wrapper, such as bound programs.

Microsoft SQL Server wrapper support

Restrictions on COMMIT statements as CURSOR_CLOSE_ON_COMMIT set YES

The Microsoft SQL Server wrapper does not support DECLARE CURSOR WITH HOLD semantics. If you declare a cursor on Microsoft SQL Server, the cursor will be closed on the server when a COMMIT statement is issued while the cursor is open. The next fetch or close cursor statement on Microsoft SQL Server will also produce errors.

The following sequence of events returns the invalid cursor state:

1. Declare cursor (with hold semantics not allowed in Microsoft SQL Server).
2. Fetch rows (SQLExtendedFetch).
3. User application issues intermittent COMMIT statements, which causes Microsoft SQL Server to close cursors.
4. Fetch continues successfully until the end of the last block fetch.
5. Close Cursor returns the error SQLState: 24000 [Microsoft][ODBC SQL Server Driver]Invalid cursor state.

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

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

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

These errors occur because the wrapper tried to reuse connections or statement handles from pending Microsoft SQL Server operations on tables with triggers.

Unicode restriction

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

Code page conversion requirements

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

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

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

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

The DataDirect Technologies Connect ODBC driver is used to access Microsoft SQL Server data sources. To use a Japanese code page on the federated server, you must add the following settings in the db2dj.ini file:

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

In the reference information provided with the DataDirect Technologies Connect ODBC driver, the correct code page is:

```
11 = Microsoft CP 932 DBCS
```

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

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

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

ODBC wrapper support

Using the PUSHDOWN server option with ODBC data sources
If queries fail because of remote source syntax errors, set the PUSHDOWN
server option value to 'N', or remove the PUSHDOWN server option.

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

To set these server options, use the ALTER SERVER statement:

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

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

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

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

Code page conversion
Code page conversions are performed by the ODBC driver or ODBC Driver
Manager. The ODBC wrapper does not perform any code page conversions.

Using the ODBC wrapper instead of an existing data source wrapper

Use the wrapper that is designed to access a specific data source, not the ODBC wrapper.

DB2 for Linux, UNIX, and Windows data sources

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

Informix data sources

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

ODBC data sources with indexes

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

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

ODBC wrapper restriction

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

Oracle wrapper support

djxlinkOracle script errors

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

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

```
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 section 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

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

This problem occurs with both the DBLIB and CTLIB wrappers.

Using host variables in pass-through sessions

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.

This requirement does not apply to the CTLIB wrapper.

Using SELECT statements on LOB columns (DBLIB wrapper only)

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

This requirement does not apply to the CTLIB wrapper.

Extended Search wrapper support

Incorrect date format in Extended Search queries

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

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

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

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

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

The following examples show queries that return the SQL0901N error.

A query without any predicate:

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

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

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

The NOT IN predicate is not supported with the FULL SELECT statement:

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

The IN predicate is not supported with the FULL SELECT statement:

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

A fixed column predicate:

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

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

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

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

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

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

Accessing a file system data source using the Extended Search wrapper
This problem occurs with IBM Lotus Extended Search, Version 4.0.

The Extended Search wrapper can access a file system data source. When you create a nickname for a file system data source, the title of the file is mapped to the value of the TITLE column that is defined for the nickname. When the value of the TITLE column is returned from the Extended Search server, the value should be converted from the Extended Search server operating system code page to UTF-8. However, IBM Lotus Extended Search, Version 4.0 does not convert the code page, and incorrect results are returned.

To fix this problem, you can either install fix pack 1 for IBM Lotus Extended Search,Version 4.0 on the Extended Search server, or 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 <ES_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 system processes. The system processes are Extended Search Server and Extended Search Remote Method Invocation Server.

Teradata wrapper support

Configuring the Teradata wrapper for Teradata Server Version 2.5
The Teradata wrapper now supports Teradata release V2R5 on AIX Version 5 or later (32-bit) and Windows (32-bit) operating systems.

Prerequisites: To use Teradata release V2R5, you must install and configure Teradata client software that supports the Teradata Call-Level Interface, Version 2 (CLIv2) Release 04.07 or later, on the federated server.

If you are configuring the Teradata wrapper for the first time, specify **2.5** as the version number when you register a server definition for a Teradata data source, as shown in this example:
```
CREATE SERVER TERASERVER1 TYPE TERADATA VERSION 2.5
WRAPPER my_wrapper OPTIONS (NODE 'tera_node')
```

If you already use the Teradata wrapper, you can upgrade to Teradata release V2R5 with either of the following procedures:
- Modify your existing Teradata wrapper configuration.
- Recreate your entire Teradata wrapper configuration of servers, user mappings, type mappings, function mappings, and nicknames.

To modify an existing Teradata wrapper configuration:

1. Change the Teradata server product version to **2.5** on all of the Teradata servers by using the ALTER SERVER statement:

   ```
   ALTER SERVER MYTERASERVER VERSION 2.5
   ```

2. Remove the TUPLE_ID_TYPE nickname option on all of the Teradata nicknames by using the ALTER NICKNAME statement:

   ```
   ALTER NICKNAME TERANICKNAME OPTIONS (DROP TUPLE_ID_TYPE)
   ```

   To get a list of all nicknames that need to be changed, use the following SQL statement to return nickname object names, server object names, and server versions for all nicknames that are associated with a TERADATA server type:

   ```
   SELECT T.tabname, S.servername, S.serverversion FROM syscat.servers
    AS S, syscat.taboptions AS T WHERE T.option = 'SERVER' AND
    S.servername = T.setting AND S.servertype = 'TERADATA' AND
    S.serverversion IN ( '2.3', '2.4', '3', '4' )
   ```

To recreate an entire Teradata wrapper configuration:

1. Use the **db2look** command to get the correct syntax for each data source configuration statement. The **db2look** command generates DDL statements from the metadata in the DB2 catalog.

2. In the **db2look** command output file, change all references to the Teradata product version to **2.5** for all server objects.

3. Drop the Teradata wrapper.

   When you drop a wrapper, all server definitions, user-defined function mappings, and user-defined data type mappings that are dependent on the wrapper are dropped. All user-defined function mappings, nicknames, user-defined data type mappings, and user mappings that are dependent on the dropped server definitions are also dropped. All packages dependent on the dropped objects and inoperative views are invalidated.

4. If necessary, drop type and function mappings that are not explicitly or implicitly associated with the wrapper you dropped.

   The DROP statement drops only the objects that are explicitly or implicitly associated with a wrapper. If other objects are associated with the name of the wrapper that you dropped, you must manually drop those objects.

5. Issue the SQL statements in the db2look output file to reconfigure access to Teradata.

   For example, if the name of the output file is `teradata.sql`, issue the following command:

   ```
   db2 -tvf teradata.sql
   ```

6. Optional: Bind any application packages that were invalidated when you dropped the wrapper.

   DB2 Universal Database automatically rebinds invalidated packages the next time those packages are accessed. However, to control when the

rebinding occurs, you might want to explicitly rebind invalid packages rather than have DB2 Universal Database automatically rebind them.

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.

Specifying version numbers in Teradata server definitions

You cannot use a single-digit number when you specify the Teradata version number in Teradata server definitions. Instead, you must use a two-digit number with a decimal point. For example, to create a server definition that specifies Teradata V2R3, use the following command:

```
CREATE SERVER TERASERVERA TYPE TERADATA VERSION 2.3
 WRAPPER my_wrapper OPTIONS (NODE 'tera_node')
```

You do not need to alter existing Teradata server definitions that use a single-digit number to specify the Teradata version.

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 SQL3324.

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

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

BLAST and HMMER wrapper support

Missing library for BLAST and HMMER wrapper daemons (Linux only)

The BLAST wrapper requires a BLAST daemon to listen for BLAST job requests from the wrapper. The HMMER wrapper requires a HMMER daemon to listen for HMMER job requests from the wrapper. Both the BLAST daemon and the HMMER daemon require a library named libcxa.so.1 on the computer where the daemon is installed. If you did not install the BLAST utility or the HMMER utility and the federated server on the same computer, the libcxa.so.1 library might not be installed. If the libcxa.so.1 library is missing, you will receive an error message from the daemon that indicates no such file or directory.

Copy the libcxa.so.1 library to the computer where you installed the daemon. The library is included with DB2 Information Integrator. The libcxa.so.1 library is in the installation directory /opt/IBM/db2/V8.1/lib.

You can copy the library to the same directory where you installed the BLAST daemon or HMMER daemon, or you can copy the library to another directory on the computer where you installed the BLAST utility or the HMMER utility. If you copy the libcxa.so.1 library to a different directory from where the daemon is installed, you must also change the LD_LIBRARY_PATH environment variable on the computer where you installed the BLAST utility or the HMMER utility to point to that directory. See the topic "UNIX Environment Variable Settings" in the DB2 Information Center for more information about the Linux LD_LIBRARY_PATH setting.

Accessing Microsoft Excel data sources

You can access Microsoft Excel spreadsheets by using either the Excel wrapper or the ODBC wrapper. Additional details on the advantages and disadvantages of each access method, the configuration requirements, and the limitations of each method are described in the white paper "Choosing the best wrapper to access Excel data." You can get this white paper from the DB2 Information Integrator Support site on the Web at www.ibm.com/software/data/integration/db2ii/support.html.

Using the CREATE TYPE MAPPING statement

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

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

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

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

The get_stats and get_stats_nr tools

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

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

get_stats

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

XML   The get_stats tool can collect the statistics of the root nickname only when that nickname was created with the FILE_PATH or DIRECTORY_PATH option.

Table-structured files
The get_stats tool can run only if the nickname was not created with the DOCUMENT option.

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

get_stats_nr

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

The get_stats_nr tool runs only on Windows and AIX.

You can download the get_stats and get_stats_nr tools from the DB2 Information Integrator Web site at www.ibm.com/software/data/integration/db2ii/support.html.

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

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

```
SELECT * FROM db2_table LEFT OUTER JOIN non_relational_nickname
ON db2_table.col = non_relational_nickname.col
WHERE <optional_predicates>
```

db2_table can be a local table or a nickname.

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

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

Informix data source connection errors (Solaris 64-bit operating environments only)

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 "Altering server definitions and server options" topic in the DB2 Information Center. For more information about the INFORMIX_LOCK_MODE setting, see the "Server options for federated systems" topic in the DB2 Information Center.

If the deadlock or timeout error persists, contact IBM Software Support.

New features added to BioRS, Documentum, and Entrez wrapper custom functions

New features have been added to the BioRS, Documentum, and Entrez wrapper custom functions in DB2 Information Integrator, Version 8.1, Fix Pack 5.

If you did not register the custom functions for the BioRS, Documentum, or Entrez wrapper in a previous release of DB2 Information Integrator, then you do not need to perform any additional tasks. The new features are available

after you complete the normal registration process for the custom functions for each wrapper, as described in the topics about registering custom functions in the DB2 Information Center.

If you registered the custom functions for the BioRS, Documentum, or Entrez wrapper in a previous release of DB2 Information Integrator, then you must perform some additional tasks before you can use the new features. You use the create_function_mappings.ddl script file, which is available for each of the three wrappers, to drop and reregister the custom functions.

If you do not drop and reregister the custom functions for these wrappers after you have installed DB2 Information Integrator, Version 8.1, Fix Pack 5, then incorrect results might be received for queries that use the custom functions for these three wrappers.

To drop and reregister the custom functions for the BioRS, Documentum, or Entrez wrapper:

1. Open the create_function_mappings.ddl script file for the wrapper in a text editor.

   BioRS wrapper:
   > The file is in the $DB2_INSTANCE/sqllib/samples/lifesci/biors directory.

   Documentum wrapper:
   > The file is in the $DB2_INSTANCE/sqllib/samples/lifesci/dctm directory.

   Entrez wrapper:
   > The file is in the $DB2_INSTANCE/sqllib/samples/lifesci/entrez directory.

2. Remove the comment characters from each of the DROP statement lines in the create_function_mappings.ddl script file. The comment characters consist of two dashes (--).

3. Save the create_function_mappings.ddl script file.

4. Run the create_function_mappings.ddl script file.

   For example, to run the create_function_mappings.ddl script file for the Entrez wrapper in the DB2 instance db2inst1, issue the following command:

```
$ db2 -tvf  /home/db2inst1/sqllib/samples/lifesci/entrez/create_function_mappings.ddl
```

   After you run the create_function_mappings.ddl file, all of the custom functions for the wrapper are dropped and then reregistered in one task.

Repeat the procedure for each wrapper as necessary.

DB2 Control Center updates

Differences between federated online help and DB2 Control Center object names

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

Table 4. Changed window names

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

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

Table 5. Changed control names

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

db2updv8 command

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

Syntax:

```
►►─db2updv8─-d─database-name─────────────────────────────────►◄
                            └─-u─userid─-p─password─┘
```

An example of this command is:

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

If you do not update your database to DB2 Information Integrator, Version 8.1, you will get error message SQL0444N when you try to create a wrapper from the Federated Database Objects folder in the Control Center.

Federated client and server compatibility

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

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

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

Editing the list of columns before a nickname is created

When you create a nickname for most nonrelational data sources, you can specify a list of columns and column options for the nickname in the Add Nickname window. If you specify a column and then decide that the column does not belong in the list, you can remove the column before you create the nickname. In this release, you can remove more than one column at a time.

You can remove one or more columns in one of the following ways:
- In the Add Nickname window, select the columns and click **Remove**.
- From the Create Nicknames window, click **Properties**. In the Properties window, select the columns and click **Remove**.

If you do not remove the column before you create a nickname, you must drop the nickname and create it again. Omit the column that you do not want from the new list of columns.

Wrapper development limitation

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

Corrections to the documentation

The following information describes corrections to previously released HTML, PDF, and printed documentation.

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

Installing DB2 Information Integrator

Topic title: Fast track to installing DB2 Information Integrator and setting up a federated server and database

Location in the DB2 Information Center
Installing DB2 Information Integrator –> Installing DB2 Information Integrator and setting up a federated server and database

Correction
If you install a DB2 Information Integrator fix pack on a UNIX system, you must run the djxlink script for each data source wrapper that is installed on your computer. The djxlink script links the data source client with the updated wrapper library.

Before you run the script, you must set the following environment variables:

Script	Environment variable
djxlinkInformix	INFORMIXDIR
djxlinkMssql	DJX_ODBC_LIBRARY_PATH
djxlinkOracle	ORACLE_HOME
djxlinkSybase	SYBASE
	Optional: SYBASE_OCS
djxlinkTeradata	COPLIB

Run the djxlink script from the following directories:
- For 32-bit libraries, run the djxlink script from the *db2_install_directory*/bin directory.

- For 64-bit libraries (AIX, HP-UX, and Solaris), run the djxlink script from the *db2_install_directory*/bin64 directory.

After you run the djxlink script, run the db2iupdt script for each DB2 Universal Database instance.

For example, you might run the following command to update the links for an Oracle 32-bit library on AIX:

```
cd /usr/opt/db2_08_01/bin
export ORACLE_HOME=/usr/oracle_client
djxlinkOracle
```

For example, you might run the following command to update the links for an Oracle 64-bit library on AIX:

```
cd /usr/opt/db2_08_01/bin64
export ORACLE_HOME=/usr/oracle_client
djxlinkOracle
```

Topic title: Confirming the link between DB2 and the data source client libraries (UNIX)

Location in the DB2 Information Center
Installing DB2 Information Integrator –> Installing DB2 Information Integrator and setting up a federated server and database –> Post-installation tasks

Correction
If you use the DB2_DJ_INI registry variable, you must set it to an absolute path. If the FEDERATED parameter is set to YES and the DB2_DJ_INI registry variable is set to a relative path, the DB2 Universal Database engine will not start.

Product overview

Topic title: Supported data sources

Location in the DB2 Information Center
Federated systems - overview –> Data sources

Correction
DB2 Information Integrator, Version 8.1, supports the following data source versions:
- DB2 Universal Database for iSeries, Version 5.1 or later
- DB2 for UNIX and Windows, Version 7.2 and Version 8.1
- Microsoft Excel 97, Excel 2000, and Excel 2002

Migrating to DB2 Information Integrator

Topic title: Supported features and functions

Location in the DB2 Information Center
Migrating to DB2 Information Integrator –> Planning to migrate to version 8

Correction: Update to supported data sources
You do not need to install DB2 Information Integrator to access DB2 family data sources, Informix data sources, or OLE DB data sources. Access to these data sources is provided when you install DB2 Universal Database Enterprise Server Edition, Version 8.1.

Topic title: Migration restrictions

Location in the DB2 Information Center
Migrating to DB2 Information Integrator –> Planning to migrate to version 8

Correction: Update to DataJoiner migration restrictions
In DataJoiner Version 2.1.1 you can create nicknames on DB2 aliases; however, you cannot create nicknames on DB2 aliases using DB2 Information Integrator, Version 8.

Topic title: Space considerations for DB2 Migration

Location in the DB2 Information Center
Migrating to DB2 Information Integrator –> Planning to migrate to version 8

Correction: DataJoiner
When you migrate from DataJoiner, the log file space that you need depends on the size of the database that you are migrating. You can avoid the log full error by increasing the logfilsiz database configuration parameter or creating additional log files before you migrate. The total size for all log files should be twice the size of the backup files for the database.

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

```
( 4096 X logfilsiz X ( logprimary + logsecond) ) >
2 X (DataJoiner database backup file size)
```

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

Example:

The size of the backup file for the DataJoiner database is 117.5 MB. To successfully complete the migration, the total size for all log files must be 2 X 117.5 = 235 MB. Currently, the logfilsiz database configuration parameter value is 5000. The DataJoiner database has 6 logprimary files and 4 logsecond files. The total size for all log files is 30 (4096 X 5000 X [6 + 4]) = 204.8 MB. Because the total file size for all log files is less than twice the size of the backup file of the DataJoiner database, the migration process will fail. To complete the migration, increase the logfilsiz database configuration parameter to 235 MB.

Topic title: After you migrate to DB2 Information Integrator

Location in the DB2 Information Center

Migrating to DB2 Information Integrator –> Accessing data sources after migration

Correction: Migrating type 1 indexes to type 2 indexes

After you migrate to DB2 Information Integrator, you can convert type 1 indexes to type 2 indexes for a table.

To determine the index type, issue the following command:

```
db2 load query table <table_name>
```

where *<table_name>* is the name of the table that you want to determine the index type for.

If the index is type 1, the output for the table is:

```
Tablestate:
       Normal
       Type-1 Indexes
```

If the index is type 2, the output for the table is:

```
Tablestate:
       Normal
```

To convert an index from type 1 to type 2, issue the following command:

```
db2 reorg indexes all for table <table_name> allow
[write | read | no] access convert
```

Alternatively, you can drop all indexes on the table and create them again.

Topic title: Accessing Sybase data sources after migrating to DB2 Information Integrator

Location in the DB2 Information Center

Migrating to DB2 Information Integrator –> Accessing data sources after migration

Correction 1 of 2: Updates to step 1b.

Set the following environment variables in the db2dj.ini file as necessary:

SYBASE

The directory path where the Sybase Open Client software is installed. Sybase supports 32-bit and 64-bit environments. You must set the variable paths accordingly. The following examples show the Sybase variable path set for each environment.

For 32-bit environments:

```
SYBASE=/sybase/sdk125_32
```

For 64-bit environments:

```
SYBASE=/sybase/sdk125_64
```

SYBASE_OCS

If you are using Sybase Open Client Version 12 or later, set the SYBASE_OCS environment variable to the name of the OCS directory. For example:

```
SYBASE_OCS=OCS-12_5
```

Correction 2 of 2: Update to step 2

Before you perform step 2, update the .profile file of the DB2 instance with the SYBASE environment variable by issuing the following commands:

```
export SYBASE=sybase_home_directory
export PATH=$SYBASE/bin:$PATH
```

sybase_home_directory is the directory where the Sybase Open Client software is installed. Next, you must perform the following additional step:

Update the .profile file of the DB2 instance to include or set the LD_LIBRARY_PATH and LIBPATH environment variables by issuing the following commands:

```
export LIBPATH=$SYBASE/lib
export LD_LIBRARY_PATH=$SYBASE/lib
```

Configuring federated systems and data sources

Topic title: Configuring the BLAST daemon

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 –> Configuring the
BLAST daemon

Correction 1 of 6: Limitation of the BLAST daemon for UNIX and Windows systems
If the path name for the BLAST daemon, its executable file, or
the database contains spaces, the BLAST daemon might not
run properly. For example, do not install the BLAST
executable file in `C:\Program Files`.

Correction 2 of 6: Update to BLAST daemon and executable files available on Windows systems
The BLAST daemon and executable files are available on
federated servers that run the Windows operating system. The
following sections describe how to configure the BLAST
daemon on Windows servers.

Correction 3 of 6: Daemon executable file location on Windows
The daemon executable files, `db2blast_daemon.exe` and
`db2blast_daemon_svc.exe` can run in any directory on the
BLAST remote server. On federated servers that run Windows,
the daemon executable file is placed in the `\SQLLIB\bin`
directory when DB2 Universal Database is installed.

Correction 4 of 6: Sample daemon configuration file on Windows
A sample daemon configuration file named
BLAST_DAEMON.config is placed in the directory
`DB2PATH\samples\lifesci`, where `DB2PATH` is the directory in
which DB2 Universal Database is installed, usually the
`\SQLLIB` directory.

Correction 5 of 6: Update to database specification entry example
When you specify the GenBank BLAST data source on
Windows systems, add the following line to the daemon
configuration file:

`C:\vnr_data\genbank_nonest1.fasta`

Correction 6 of 6: Windows daemon configuration file example
The following example shows the contents of a sample
configuration file on Windows, with the required options and
BLAST data source specification for GenBank and
SWISS-PROT. Use the default file name,
BLAST_DAEMON.config.

```
=
DAEMON_PORT=4007
MAX_PENDING_REQUESTS=10
```

```
DAEMON_LOGFILE_DIR=.\
Q_SEQ_DIR_PATH=.\
BLAST_OUT_DIR_PATH=.\
BLASTALL_PATH=.\blastall.exe
genbank=c:\vnr_data\genbank_nonest1.fasta
swissprot=c:\vnr_data\swissprot
```

Topic title: Starting the BLAST daemon

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 –> Starting the
BLAST daemon

Correction: Procedure for starting the BLAST daemon on Windows systems

The BLAST daemon is available on federated servers that run
Windows. To manage the BLAST daemon on Windows, enter
the following command:

```
db2blast_daemon -a <action> -c <config_file>
 -d <debug_level> -u <user_id>
 -p <password>
```

<action>

Performs the specified activity. Valid actions are:

- *install*: Installs the service.
- *start*: Starts the service. The service is installed first
  if it is not already installed.
- *stop*: Stops the service.
- *remove*: Removes the service. The service is stopped
  first if it is still running.
- *status*: Shows the status of the service.

<config_file>

Instructs the service to use the specified configuration
file. If you do not specify the configuration file, the
daemon searches for the BLAST_DAEMON.config file.

<debug_level>

Sets the service debug level to the specified value. The
valid values are 1, 2, or 3.

<user_id>

Sets the service to run under the specified user ID.

<password>

Specifies the password for the specified user ID. The
password is valid and required only when you specify

the -u option. If the -p option is not specified when you set the -u option, the program prompts you for the password.

Options -c and -d are used only with the install and start actions. Options -u and -p are used only with the install action. Options -c, -d, -u, and -p are ignored if they are used for the other actions. The options that are specified with the start action affect only the current run of the daemon, and they override the values that are specified with the install action.

The following examples show daemon actions. These examples assume that the BLAST_DAEMON.config file is in the same directory as db2blast_daemon.exe.

- To check the status of the daemon:

  ```
  db2blast_daemon -a status
  ```

- To install the daemon service with debugging turned on:

  ```
  db2blast_daemon -a start install -d 3
  ```

- To start the daemon:

  ```
  db2blast_daemon -a start
  ```

- To stop the daemon:

  ```
  db2blast_daemon -a stop
  ```

- To remove the daemon service:

  ```
  db2blast_daemon -a remove
  ```

Topic title: Registering the server definitions for a DB2 data source

Location in the DB2 Information Center

Configuring federated systems and data sources–> Configuring data sources –> Configuring DB2 family data sources –> Adding DB2 family data sources to a federated server

Correction

DB2 Information Integrator no longer supports DB2 Universal Database for iSeries, Version 4, or DB2 Universal Database for Linux, UNIX, and Windows, Version 6 or version 7.1.

The supported versions are:

- DB2 for UNIX and Windows, Version 7.2, and version 8.1
- DB2 for z/OS and OS/390, Version 5 or later
- DB2 for iSeries, Version 5.1 or later

Topic title: What is Excel?

Location in the DB2 Information Center
>Configuring federated systems and data sources –>
>Configuring data sources –> Configuring Excel data sources
>–> Excel wrapper

Correction
>The Microsoft Excel versions that are supported have
>changed. The supported versions are Excel 97, Excel 2000, and
>Excel 2002.

Topic title: Configuring the HMMER daemon

Location in the DB2 Information Center
>Configuring federated systems and data sources –>
>Configuring data sources –> Configuring HMMER data
>sources –> Adding HMMER to a federated system –>
>Configuring the HMMER daemon

Correction 1 of 6: Limitation of the HMMER daemon for UNIX and Windows systems
>The HMMER daemon might not run properly if the HMMER
>executable file path or the databases paths contain spaces. For
>example, do not install the HMMER executable file in
>`C:\Program Files` on Windows systems.

Correction 2 of 6: Update to HMMER daemon and executable files available on Windows systems
>The HMMER daemon and executable files are available on
>federated servers that run the Windows operating system. The
>following three corrections describe how to configure the
>HMMER daemon on Windows.

Correction 3 of 6: Daemon executable file location on Windows
>The daemon executable files, `db2hmmer_daemon.exe` and
>`db2hmmer_daemon_svc.exe`, can run in any directory on the
>HMMER remote server. On federated servers that run
>Windows, the daemon executable file is placed in the
>`\SQLLIB\bin` directory when DB2 Universal Database is
>installed.

>If you did not install HMMER and the federated server on the
>same computer, you must copy the daemon executable files to
>the same computer where HMMER is installed. Copy all of
>the following files:
>- db2hmmer_daemon.exe
>- db2hmmer_daemon_svc.exe
>- HMMER_DAEMON.config
>- db2h2x.exe

- hmmpfam.exe (HMMER executable, not supplied by IBM)
- HMMER database files (not supplied by IBM)

Correction 4 of 6: RUNPFAM_PATH option
>This option is not used on Windows systems. This option is ignored if it is specified.

Correction 5 of 6: Update to the database specification entry example
>When you specify the MYHMMS profile database on Windows systems, add the following line to the daemon configuration file:
>
>```
>myhmms=C:\hmmer\tutorial\myhmms
>```

Correction 6 of 6: Windows daemon configuration file example
>The following example shows the contents of a sample configuration file on Windows, with the required options and HMMER data source specification for three different sequence databases.
>
>```
>=
>DAEMON_PORT=4098
>MAX_PENDING_REQUESTS=10
>DAEMON_LOGFILE_DIR=.\
>Q_SEQ_DIR_PATH=.\
>HMMER_OUT_DIR_PATH=.\
>HMMPFAM_PATH=c:\hmmer\bin\hmmpfam.exe
>H2X_PATH=.\db2h2x.exe
>myhmms=c:\hmmer\tutorial\myhmms
>globin=c:\hmmer\tutorial\tutorial\globin.hmm
>pfamseq=c:\hmmer\pfam\pfamseq
>```

Topic title: Starting the HMMER daemon

Location in the DB2 Information Center
>Configuring federated systems and data sources –> Configuring data sources –> Configuring HMMER data sources –> Adding HMMER to a federated system –> Starting the HMMER daemon

Correction: Procedure for starting the HMMER daemon Windows systems
>The HMMER daemon is available on Windows systems. To manage the HMMER daemon on Windows, enter the following command:
>
>```
>db2hmmer_daemon -a <action> -c <config_file> -d <debug_level>
> -u <user_id> -p <password>
>```
>
>*<action>*
>>Performs the specified activity. Valid actions are:
>>- *install*: Installs the service.

- *start*: Starts the service. The service is installed first if it is not already installed.
- *stop*: Stops the service.
- *remove*: Removes the service. The service is stopped first if it is still running.
- *status*: Shows the status of the service.

<config_file>
  Instructs the service to use the specified configuration file. If you do not specify the configuration file, the daemon searches for BLAST_DAEMON.config.

<debug_level>
  Sets the service debug level to the specified value. The valid values are 1, 2, or 3.

<user_id>
  Sets the service to run under the specified user ID.

<password>
  Specifies the password for the specified user ID. The password is valid and required only when you specify the -u option. If the -p option is not specified when you set the -u option, the program prompts you for the password.

Options -c and -d are used only with the install and start actions. Options -u and -p are used only with the install action. Options -c, -d, -u, and -p are ignored if they are used for the other actions. The options that are specified with the start action only affect the current run of the daemon, and they override the values that are specified with the install action.

The following examples show daemon actions. These examples assume that the HMMER_DAEMON.config file is in the same directory as the db2hmmer_daemon.exe file.

- To check the status of the daemon:

  ```
  db2hmmer_daemon -a status
  ```
- To install the daemon service with debugging turned on:

  ```
  db2hmmer_daemon -a install -d 3
  ```
- To start the daemon:

  ```
  db2hmmer_daemon -a start
  ```
- To stop the daemon:

  ```
  db2hmmer_daemon -a stop
  ```

- To remove the daemon service:

  ```
  db2hmmer_daemon -a remove
  ```

Developing applications

Topic title: Declaring and referencing parameters in the DADX file

Location in the DB2 Information Center

Developing applications –> Developing Web services and Web Applications –> Developing applications that use the Web services provider –> The DADX file

Correction

The table of supported XML schema and SQL types is updated to include the SQL types BIGINT, CLOB, LONG VARCHAR, and TINYINT. The following table includes the new SQL types:

Table 6. Supported XML Schema and SQL types

XML schema simple types	SQL type
string	CHAR, VARCHAR, CLOB, LONG VARCHAR
decimal	DECIMAL, NUMERIC
int	INTEGER
short	SMALLINT
float	FLOAT
double	REAL, DOUBLE PRECISION
date	DATE
time	TIME
timestamp	TIMESTAMP
long	BIGINT
byte	TINYINT

Reference information

Topic title: Wrapper_Utilities class

Location in the DB2 Information Center

Reference information –> Wrapper development reference –> Data classes –> C++ for developing wrappers

Correction

For this class, add the get_db2_release function. This function

returns the version of DB2 Universal Database, including the fix pack, that the wrapper currently runs on. The return value is one of the release identifiers defined in include file sql.h (for example, SQL_REL8105, which corresponds to DB2 Universal Database, Version 8.1, Fix Pack 5).

The syntax for this function is:

```
int get_db2_release (void)
```

This function does not have input arguments or output arguments.

Notices

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

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

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

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

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

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

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will

be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

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

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

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

COPYRIGHT LICENSE:

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

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

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

Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both:

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

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

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

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

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

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

IBM ®

Printed in USA