

IBM DB2 Information Integrator



Release Notes

Version 8.14

IBM DB2 Information Integrator



Release Notes

Version 8.14

Before using this information and the product it supports, be sure to read the general information under “Notices” on page 33.

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

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

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

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

© Copyright International Business Machines Corporation 2003. All rights reserved.

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

Contents

About the release notes	v	Sybase wrapper support	15
DB2 Information Integrator support on the Web	v	Extended Search wrapper support	15
Known problems, limitations, and workarounds	1	Teradata wrapper support	18
DB2 Information Integrator installation issues	1	BLAST and HMMER wrapper support	18
The DB2 Information Integrator installer fails when the directory name contains a blank character	1	Accessing Microsoft Excel data sources	19
Installing DB2 Information Integrator over versions of DB2 Universal Database or DB2 Connect higher than 8.1.2 (Linux only)	1	Using the CREATE TYPE MAPPING statement	19
DB2 Information Integrator installer has English license text on Russian Linux installers	3	DB2LOOK command limitations generating federated DDL statements	19
Changing library permissions	3	The get_stats and get_stats_nr tools	19
Installing life sciences user-defined functions	3	Error in left outer join queries to BLAST, Entrez, Extended Search, HMMER, and XML data sources	20
Updated information for DB2 Universal Database documentation	5	Informix data source connection errors (Solaris 64-bit Operating Environments only)	21
Migration issues	5	New features added to BioRS, Documentum, and Entrez wrapper custom functions	21
Accessing DB2 family remote data sources after migration	5	DB2 Control Center updates	22
Federated function and data source support issues	6	Differences between federated online help and DB2 Control Center object names	22
Data type restrictions	6	db2updv8 command	23
UTF-8 code page data expansion	9	Federated client and server compatibility	23
Materialized query table issues	9	Removing columns when creating a nonrelational nickname	23
LIKE predicate issues	10	MQ user-defined function message string updates	24
CHAR columns that contain DBCS data.	10	Corrections to the documentation	25
Entrez wrapper support	10	Reference information.	25
Microsoft SQL Server wrapper support	11	Notices	33
ODBC wrapper support	12	Trademarks	35
Oracle wrapper support	14		

About the release notes

The release notes contain the latest information about DB2[®] Information Integrator, Version 8.1.4.

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

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

If you installed a version of DB2 Information Integrator that is later than 8.1.2, or a version of DB2 Connect™ that is later than 8.1.2, and then install the DB2 Information Integrator relational wrappers or the DB2 Information Integrator nonrelational wrappers, the wrappers are not installed. A message 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 as a result. You must install the DB2 Information Integrator wrappers manually. You can use the `db2_install` script or the RedHat Package Manager to install the DB2 Information Integrator wrappers.

After you install the wrappers, you must reinstall the DB2 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 manually 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. For information about configuring DB2 Information Integrator to access remote data sources, see "Installing DB2 Information Integrator and setting up a federated server and database" in the *DB2 Information Integrator Installation Guide*.

Installing wrappers by using the `db2_install` script

To install wrappers using the `db2_install` script:

1. Log in as a user with root authority.
2. Mount the DB2 Information Integrator CD.
3. Change your directory to 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

To install wrappers using the RedHat Package Manager:

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

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

The `ComponentList.htm` file is in the Linux/*prod*/*db2*/*linux* directory on 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 to install the Signature for DB2 Information Integrator Nonrelational Wrappers component (filename: IBM_db2lssg81-8.1.0.16.i386.rpm) and the Structured Files Data Sources component (filename: IBM_db2lsstf81-8.1.0.16.i386.rpm) 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 8.1.4.

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 placed on your system:

Windows

sqllib\function\db2lsudfs.dll

AIX

- sqllib/function/libdb2lsudfs.a
- sqllib/function/libdb2lsгенewise.a

HP-UX

- sqllib/function/libdb2lsudfs.sl
- sqllib/function/libdb2lsгенewise.sl

Solaris Operating Environment

- sqllib/function/libdb2lsudfs.so
- sqllib/function/libdb2lsгенewise.so

Linux

- sqllib/function/libdb2lsudfs.so
- sqllib/function/libdb2lsгенewise.so

For the LSGeneWise user-defined function, the sqllib/lib/libdb2lsSTгенewise.a archive file is also installed on your system.

The following life sciences user-defined functions are installed with DB2 Information Integrator, Version 8.1.4:

- LSPep2AmbNuc
- LSPep2ProbNuc
- LSDefineParse
- 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 8.1.4. See the

DB2 Information Integrator Support site at [/www.ibm.com/software/data/integration/db2ii/support.html](http://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 8.1.4.
4. If you plan to use the LSGeneWise user-defined function, run the `djxlinkLSGeneWise` script to create the `sqlib/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
<i>CLI Guide and Reference, Volume 1,</i> SC09-4849-01	db211x80
<i>CLI Guide and Reference, Volume 2,</i> SC09-4850-01	db212x80
<i>Installation and Configuration Supplement,</i> GC09-4837-00	db2iyx80

Migration issues

Accessing DB2 family remote data sources after migration

If you need to access a DB2 family remote data source after you migrate to DB2 Information Integrator from either DB2 Version 7 or from DataJoiner[®] Version 2.1.1, you must 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 a -805 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 use the bind utility to reestablish a connection to a remote data source after migration:

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

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

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

3. Disconnect from the remote data source.

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
Microsoft® SQL Server	SQL_VARIANT

Table 2. Unsupported data source data types (continued)

Data source	Unsupported data types
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:

Table 3. Write operation restrictions on data types

Data source	Unsupported data types
DB2 family (DRDA [®])	BLOB CLOB DBCLOB
Informix [®]	BLOB CLOB
Microsoft SQL Server	image ntext text SQL_VARIANT

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

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

Mapping to the federated GRAPHIC and VARGRAPHIC data types

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

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

UTF-8 code page data expansion

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

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 called 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 issues

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

LIKE predicate applied to a CHAR column

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

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

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

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

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

CHAR columns that contain DBCS data

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

Entrez wrapper support

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

Microsoft SQL Server wrapper support

MSSQL wrapper restrictions on COMMIT statements as CURSOR_CLOSE_ON_COMMIT set YES

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

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

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

When you issue an insert, update, or delete 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 requirements

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

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

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_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, non-character (fixed length or varying length) columns does not exist on the nickname or on a corresponding remote table for the nickname. The failure error is SQL30090 reason code 21.

Oracle wrapper support

djxlinkOracle script errors

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

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

```
lspp -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 SQL 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 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.

Changing VARCHAR columns to CLOB or BLOB columns

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

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.

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

Incorrect format

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

Correct format

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

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

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

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

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

The following examples show queries that return the SQL0901N error.

A query without any predicate:

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

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

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

The NOT IN predicate is not supported using FULL SELECT:

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

The IN predicate is not supported using FULL SELECT:

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

A fixed column predicate:

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


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

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

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

FIELD_DATATYPE SMALLINT supported data types

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

Unsupported data types

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

Accessing a file system data source using the IBM Lotus Extended Search wrapper

This problem occurs with IBM Lotus Extended Search, Version 4.0.

One of the data sources that the Extended Search wrapper can access is 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

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. See the Teradata documentation for more information.

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 obtain 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 results in ODBC REMOTE stored as the server type in the catalog:

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

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

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

DB2LOOK command limitations generating federated DDL statements

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

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. 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 not created with the `DOCUMENT` option set to `FILE`.

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

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

db2_table can be a local table or a nickname.

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

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

Informix data source connection errors (Solaris 64-bit Operating Environments only)

If a deadlock or time out 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 time out error persists, contact IBM Service.

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 8.1.4.

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 three wrappers using DB2 Information Integrator, Version 8.1.4, then unexpected or 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. Find the create_function_mappings.ddl script file for the wrapper:

BioRS wrapper:

The file is in the \$DB2_INSTANCE/sqlllib/samples/lifesci/biors directory.

Documentum wrapper:

The file is in the \$DB2_INSTANCE/sqlllib/samples/lifesci/dctm directory.

Entrez wrapper:

The file is in the \$DB2_INSTANCE/sqlllib/samples/lifesci/entrez directory.

2. Using a text editor, 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/sqlllib/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.

does not belong in the list, you can remove the column before you create the nickname. Starting 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.

MQ user-defined function message string updates

The value '2pc', for two-phase commit, is not available in this release. The valid values are '0pc', '1pc', and 'all'. The messages should read:

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

Corrections to the documentation

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

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.

Reference information

Topic Title: Default forward data type mappings

Location in the DB2 Information Center

DB2 Information Integrator -> Reference information ->
Federated systems reference

Correction 1 of 2

For Microsoft SQL Server 2000, the remote data type bigint maps to the DB2 data type BIGINT. You can use the ALTER NICKNAME statement to change the DB2 data type of a numeric nickname column to the BIGINT data type.

Correction 2 of 2

Lengths for CHAR, FLOAT, RAW, and VARCHAR remote data types have been updated for Oracle NET8 and Oracle SQLNET data sources. LONG, LONG RAW, and TIMESTAMP remote data types have been added to the Oracle NET8 table.

The following table shows the Oracle NET8 forward default data type mappings.

Table 6. Oracle NET8 forward default data type mappings

REMOTE_TYPENAME	REMOTE_LOWER_LEN	REMOTE_UPPER_LEN	REMOTE_LOWER_SCALE	REMOTE_UPPER_SCALE	REMOTE_BIT_DATA	REMOTE_DATA_OPERATORS	FEDERATED_TYPENAME	FEDERATED_LENGTH	FEDERATED_SCALE	FEDERATED_BIT_DATA
BLOB	0	0	0	0	-	\0	BLOB	2147483647	0	Y
CHAR	1	254	0	0	-	\0	CHAR	0	0	N
CHAR	255	2000	0	0	-	\0	VARCHAR	0	0	N
CLOB	0	0	0	0	-	\0	CLOB	2147483647	0	N
DATE	0	0	0	0	-	\0	TIMESTAMP	0	0	N
FLOAT	1	126	0	0	-	\0	DOUBLE	0	0	N
LONG	0	0	0	0	-	\0	CLOB	2147483647	0	N
LONG RAW	0	0	0	0	-	\0	BLOB	2147483647	0	Y
MLSLABEL	0	0	0	0	-	\0	VARCHAR	255	0	N
NUMBER	1	38	-84	127	-	\0	DOUBLE	0	0	N
NUMBER	1	31	0	31	-	>=	DECIMAL	0	0	N
NUMBER	1	4	0	0	-	\0	SMALLINT	0	0	N
NUMBER	5	9	0	0	-	\0	INTEGER	0	0	N
RAW	1	254	0	0	-	\0	CHAR	0	0	Y
RAW	255	2000	0	0	-	\0	VARCHAR	0	0	Y
ROWID	0	0	0	NULL	-	\0	CHAR	18	0	N
TIMESTAMP ¹	-	-	-	-	-	-	TIMESTAMP	10	-	-
VARCHAR2	1	4000	0	0	-	\0	VARCHAR	0	0	N

Notes:

1. This type mapping is valid only for Oracle 9i (or later) client and server configurations.

The following table shows the Oracle SQLNET forward default data type mappings.

Table 7. Oracle SQLNET forward default data type mappings

REMOTE_TYPENAME	REMOTE_LOWER_LEN	REMOTE_UPPER_LEN	REMOTE_LOWER_SCALE	REMOTE_UPPER_SCALE	REMOTE_BIT_DATA	REMOTE_DATA_OPERATORS	FEDERATED_TYPENAME	FEDERATED_LENGTH	FEDERATED_SCALE	FEDERATED_BIT_DATA
CHAR	1	254	0	0	-	\0	CHAR	0	0	N
CHAR	255	2000	0	0	-	\0	VARCHAR	0	0	N
DATE	0	0	0	0	-	\0	TIMESTAMP	0	0	N
FLOAT	1	126	0	0	-	\0	DOUBLE	0	0	N
LONG	0	0	0	0	-	\0	CLOB	2147483647	0	N
LONG RAW	0	0	0	0	-	\0	BLOB	2147483647	0	Y
MLSLABEL	0	0	0	0	-	\0	VARCHAR	255	0	N
NUMBER	1	38	-84	127	-	\0	DOUBLE	0	0	N
NUMBER	1	31	0	31	-	>=	DECIMAL	0	0	N
NUMBER	1	4	0	0	-	\0	SMALLINT	0	0	N
NUMBER	5	9	0	0	-	\0	INTEGER	0	0	N
RAW	1	254	0	0	-	\0	CHAR	0	0	Y
RAW	255	2000	0	0	-	\0	VARCHAR	0	0	Y
ROWID	0	0	0	NULL	-	\0	CHAR	18	0	N
VARCHAR2	1	4000	0	0	-	\0	VARCHAR	0	0	N

Topic Title: Default reverse data type mappings

Location in the DB2 Information Center

DB2 Information Integrator -> Reference information ->
 Federated systems reference

Correction 1 of 3

For Microsoft SQL Server 2000, the DB2 data type BIGINT maps to the remote data type bigint.

Correction 2 of 3

The following changes were made to the Oracle SQLNET and Oracle NET8 data source default reverse data type mapping tables:

- The lengths of the CHAR federated data type were updated.
- A FLOAT federated data type row was added.

The following table shows the Oracle SQLNET reverse default data type mappings.

Table 8. Oracle SQLNET reverse default data type mappings

FEDERATED_TYPENAME	FEDERATED_LOWER_LEN	FEDERATED_UPPER_LEN	FEDERATED_LOWER_SCALE	FEDERATED_UPPER_SCALE	FEDERATED_BIT_DATA	FEDERATED_DATA_OPERATORS	REMOTE_TYPENAME	REMOTE_LENGTH	REMOTE_SCALE	REMOTE_BIT_DATA
BLOB	0	2147483647	0	0	Y	\0	LONG RAW	0	0	Y
CHARACTER	1	254	0	0	N	\0	CHAR	0	0	N
CHARACTER	1	254	0	0	Y	\0	RAW	0	0	Y
CLOB	0	2147483647	0	0	N	\0	LONG	0	0	N
DATE	0	4	0	0	N	\0	DATE	0	0	N
DECIMAL	0	0	0	0	N	\0	NUMBER	0	0	N
DOUBLE	0	8	0	0	N	\0	FLOAT	126	0	N
FLOAT	0	8	0	0	N	\0	FLOAT	126	0	N
INTEGER	0	4	0	0	N	\0	NUMBER	9	0	N
REAL	0	4	0	0	N	\0	FLOAT	63	0	N
SMALLINT	0	2	0	0	N	\0	NUMBER	4	0	N
TIME	0	3	0	0	N	\0	DATE	0	0	N
TIMESTAMP	0	10	0	0	N	\0	DATE	0	0	N
VARCHAR	1	2000	0	0	Y	\0	RAW	0	0	Y
VARCHAR	1	4000	0	0	N	\0	VARCHAR2	0	0	N

Table 8. Oracle SQLNET reverse default data type mappings (continued)

FEDERATED_TYPENAME	FEDERATED_LOWER_LEN	FEDERATED_UPPER_LEN	FEDERATED_LOWER_SCALE	FEDERATED_UPPER_SCALE	FEDERATED_BIT_DATA	FEDERATED_DATA_OPERATORS	REMOTE_TYPENAME	REMOTE_LENGTH	REMOTE_SCALE	REMOTE_BIT_DATA
--------------------	---------------------	---------------------	-----------------------	-----------------------	--------------------	--------------------------	-----------------	---------------	--------------	-----------------

Note: The DB2 Universal Database for Linux, UNIX, and Windows BIGINT data type is not available for transparent DDL. You cannot specify the BIGINT data type in a CREATE TABLE statement when you create a remote Oracle table.

The following table shows the Oracle NET8 reverse default data type mappings.

Table 9. Oracle NET8 reverse default data type mappings

FEDERATED_TYPENAME	FEDERATED_LOWER_LEN	FEDERATED_UPPER_LEN	FEDERATED_LOWER_SCALE	FEDERATED_UPPER_SCALE	FEDERATED_BIT_DATA	FEDERATED_DATA_OPERATORS	REMOTE_TYPENAME	REMOTE_LENGTH	REMOTE_SCALE	REMOTE_BIT_DATA
BLOB	0	2147483647	0	0	Y	\0	BLOB	0	0	Y
CHARACTER	1	254	0	0	N	\0	CHAR	0	0	N
CHARACTER	1	254	0	0	Y	\0	RAW	0	0	Y
CLOB	0	2147483647	0	0	N	\0	CLOB	0	0	N
DATE	0	4	0	0	N	\0	DATE	0	0	N
DECIMAL	0	0	0	0	N	\0	NUMBER	0	0	N
DOUBLE	0	8	0	0	N	\0	FLOAT	126	0	N

Table 9. Oracle NET8 reverse default data type mappings (continued)

FEDERATED_TYPENAME	FEDERATED_LOWER_LEN	FEDERATED_UPPER_LEN	FEDERATED_LOWER_SCALE	FEDERATED_UPPER_SCALE	FEDERATED_BIT_DATA	FEDERATED_DATA_OPERATORS	REMOTE_TYPENAME	REMOTE_LENGTH	REMOTE_SCALE	REMOTE_BIT_DATA
FLOAT	0	8	0	0	N	\0	FLOAT	126	0	N
INTEGER	0	4	0	0	N	\0	NUMBER	9	0	N
REAL	0	4	0	0	N	\0	FLOAT	63	0	N
SMALLINT	0	2	0	0	N	\0	NUMBER	4	0	N
TIME	0	3	0	0	N	\0	DATE	0	0	N
TIMESTAMP	0	10	0	0	N	\0	DATE	0	0	N
VARCHAR	1	4000	0	0	N	\0	VARCHAR2	0	0	N
VARCHAR	1	2000	0	0	Y	\0	RAW	0	0	Y

Note: The DB2 Universal Database for Linux, UNIX, and Windows BIGINT data type is not available for transparent DDL. You cannot specify the BIGINT data type in a CREATE TABLE statement when you create a remote Oracle table.

Correction 3 of 3

The VARCHAR federated data type rows of the Sybase default reverse data type mapping table have been updated.

The following table shows the Sybase CTLIB and DBLIB reverse default data type mappings.

Table 10. Sybase CTLIB and DBLIB default reverse data type mappings

FEDERATED_TYPENAME	FEDERATED_LOWER_LEN	FEDERATED_UPPER_LEN	FEDERATED_LOWER_SCALE	FEDERATED_UPPER_SCALE	FEDERATED_BIT_DATA	FEDERATED_DATA_OPERATORS	REMOTE_TYPENAME	REMOTE_LENGTH	REMOTE_SCALE	REMOTE_BIT_DATA
BIGINT	-	-	-	-	-	-	decimal	19	0	-
BLOB	-	-	-	-	-	-	image	-	-	-
CHARACTER	-	-	-	-	N	-	char	-	-	-
CHARACTER	-	-	-	-	Y	-	binary	-	-	-
CLOB	-	-	-	-	-	-	text	-	-	-
DATE	-	-	-	-	-	-	datetime	-	-	-
DECIMAL	-	-	-	-	-	-	decimal	-	-	-
DOUBLE	-	-	-	-	-	-	float	-	-	-
INTEGER	-	-	-	-	-	-	integer	-	-	-
REAL	-	-	-	-	-	-	real	-	-	-
SMALLINT	-	-	-	-	-	-	smallint	-	-	-
TIME	-	-	-	-	-	-	datetime	-	-	-
TIMESTAMP	-	-	-	-	-	-	datetime	-	-	-
VARCHAR ¹	1	255	-	-	N	-	varchar	-	-	-
VARCHAR ¹	256	32672	-	-	N	-	text	-	-	-
VARCHAR ²	1	16384	-	-	N	-	varchar	-	-	-
VARCHAR ²	16385	32672	-	-	N	-	text	-	-	-
VARCHAR ¹	1	255	-	-	Y	-	varbinary	-	-	-
VARCHAR ¹	256	32672	-	-	Y	-	image	-	-	-
VARCHAR ²	1	16384	-	-	Y	-	varbinary	-	-	-
VARCHAR ²	16385	32672	-	-	Y	-	image	-	-	-

Table 10. Sybase CTLIB and DBLIB default reverse data type mappings (continued)

FEDERATED_TYPENAME
FEDERATED_LOWER_LEN
FEDERATED_UPPER_LEN
FEDERATED_LOWER_SCALE
FEDERATED_UPPER_SCALE
FEDERATED_BIT_DATA
FEDERATED_DATA_OPERATORS
REMOTE_TYPENAME
REMOTE_LENGTH
REMOTE_SCALE
REMOTE_BIT_DATA

Notes:

1. This type mapping is valid only for DBLIB with all Sybase server versions and for CTLIB with Sybase server version 12.0 (or earlier).
 2. This type mapping is valid only for CTLIB with Sybase server version 12.5 (or later).
-

Notices

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

COPYRIGHT LICENSE:

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

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

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

Trademarks

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

IBM
AIX
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.



Printed in USA