

IBM DB2 Information Integrator



# Release Notes

*Version 8*



IBM DB2 Information Integrator



# Release Notes

*Version 8*

Before using this information and the product it supports, be sure to read the general information under "Notices" on page 19.

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](http://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](http://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 1998 - 2003. All rights reserved.**

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

---

# Contents

<b>About the Release Notes</b> . . . . .	<b>v</b>	ODBC wrapper support . . . . .	11
<b>Known problems, limitations, and workarounds.</b> . . . . .	<b>1</b>	Oracle wrapper support . . . . .	13
Upgrading to a different edition of DB2 Information Integrator . . . . .	1	Sybase wrapper support . . . . .	13
Migration issues . . . . .	3	Teradata wrapper support . . . . .	14
Restoring DataJoiner configuration settings	3	XML wrapper support . . . . .	14
Restoring DB2 Version 7 federated systems configuration settings . . . . .	3	Accessing Microsoft Excel data sources . . . . .	15
Sybase migration issues . . . . .	3	Using the CREATE TYPE MAPPING statement . . . . .	15
Teradata migration issues . . . . .	3	DB2LOOK command limitations generating federated DDL statements . . . . .	15
ODBC migration issues . . . . .	3	The getstats tool . . . . .	16
Dropping mappings for a specific server type . . . . .	4	Control Center updates . . . . .	16
DB2 for z/OS and OS/390 application connection behavior. . . . .	4	MQ user-defined function message string updates . . . . .	17
Federated function and data source support issues . . . . .	5	Web services issues . . . . .	17
Data type restrictions . . . . .	5	Web services consumer user-defined functions . . . . .	17
Unicode UTF-8 issues . . . . .	7	Web services consumer messages . . . . .	17
LIKE predicate issues . . . . .	8	Additional documentation updates . . . . .	18
CHAR columns containing DBCS data . . . . .	9	DB2 Information Integrator Installation Guide . . . . .	18
DRDA wrapper support . . . . .	9	DB2 Information Integrator Data Source Configuration Guide . . . . .	18
BioRS wrapper support . . . . .	9	DB2 Information Integrator Federated Systems Guide . . . . .	18
Entrez wrapper support . . . . .	10	<b>Notices</b> . . . . .	<b>19</b>
Informix wrapper support . . . . .	10	Trademarks . . . . .	21
Microsoft SQL Server wrapper support . . . . .	10		



---

## About the Release Notes

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





---

## Known problems, limitations, and workarounds

The following information describes the currently known limitations, problems, and workarounds for DB2® Information Integrator, Version 8.1. The information in this section applies only to Version 8.1 of DB2 Information Integrator. Any limitations and restrictions might or might not apply to other releases of the product.

---

### Upgrading to a different edition of DB2 Information Integrator

Before you upgrade from one edition of DB2 Information Integrator to another, you need to remove the existing DB2 Information Integrator license key and add the new DB2 Information Integrator license key to your system.

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

The license files for DB2 Information Integrator are:

<b>DB2 Information Integrator Edition</b>	<b>License file name</b>
DB2 Information Integrator Replication Edition	db2iire.lic
DB2 Information Integrator Standard Edition	db2iise.lic
DB2 Information Integrator Advanced Edition	db2iiae.lic
DB2 Information Integrator Developer Edition	db2iiae.lic

#### **Prerequisites:**

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

#### **Procedure:**

To upgrade to a different edition of DB2 Information Integrator:

1. Log on to the system with a user ID that has Administrator authority on the DB2 instance.

2. Close all open programs so that the DB2 Information Integrator installation program can update files as required.
3. From a command prompt, change directories to the directory where DB2 Universal Database™ is installed.

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

- /usr/opt/db2\_08\_01 (AIX)
- /opt/IBM/db2/V8.1 (HP-UX, Linux, Solaris™ Operating Environment)
- \Program Files\IBM\SQLLIB (Windows)

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

```
db2licm -r db2ii
```

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

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

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

**UNIX® CD or network installations:**

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

```
./iiSetup.bin
```

The DB2 Information Integrator launchpad opens.

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

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

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

---

## Migration issues

### Restoring DataJoiner configuration settings

Before you migrate to DB2 Information Integrator, the migration tasks include steps to copy the DB2 DataJoiner® configuration settings. Once you migrate the instances and databases, you must restore these configuration settings.

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

### Restoring DB2 Version 7 federated systems configuration settings

Before you migrate to DB2 Information Integrator, the migration tasks include steps to copy the configuration settings stored in the `db2dj.ini` file. Once you migrate the instances and databases, you must restore these configuration settings.

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

### Sybase migration issues

Before you migrate to DB2 Information Integrator, make a copy of the `interfaces` file. This file is not automatically migrated to DB2 Information Integrator.

This file is typically located in the `$HOME/sqllib` directory.

Once you have completed the migration steps, restore the `interfaces` file back into the `$HOME/sqllib` directory of the DB2 federated instance.

### Teradata migration issues

After you migrate to DB2 Information Integrator, you cannot use the `ALTER NICKNAME` statement against Teradata® data sources.

### ODBC migration issues

If you use transparent DDL to create remote tables on ODBC data sources through DataJoiner, you might receive errors with `SELECT` statements after you migrate to DB2 Information Integrator. The DataJoiner `DATE` data type is

mapped to the ODBC DATETIME data type when you create a remote ODBC table through DataJoiner using transparent DDL.

To avoid errors with SELECT statements, after you migrate to DB2 Information Integrator change the local data type in the federated database system catalog from DATE to TIMESTAMP.

### **Dropping mappings for a specific server type**

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

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

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

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

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

---

## **DB2 for z/OS and OS/390 application connection behavior**

Applications connecting from DB2 for z/OS<sup>™</sup> and OS/390<sup>®</sup> to a DB2 federated enabled instance have the following behavior:

- CICS<sup>®</sup> applications accessing a DB2 for Linux, UNIX, and Windows federated database instance will be forced into a read-only mode and no updates will be allowed. The error -30090 is returned if an update is attempted.
- DB2 for z/OS and OS/390 coordinated transactions which access multiple resources in a single transaction, where one of the resources is a DB2 federated instance, results in either:
  - The DB2 federated instance is the only resource that can be updated within the transaction
  - A read-only connection to the DB2 federated instance

If the DB2 federated instance is the first resource updated, the update is allowed. If the DB2 federated instance is not the first resource updated, the connection to the DB2 federated instance becomes read-only.

- Applications prepared with CONNECT(1). These applications will remain one-phase commit transactions and are able to update one site, either only local data or only a single one-phase commit federated 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 contains any of these unsupported data types. Additionally, DB2 Information Integrator does not allow insert, update, and delete operations on data source objects that contain certain data types.

### Unsupported data types

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

*Table 1. Unsupported data source data types*

Data source	Unsupported data types
DB2 for iSeries™	VARG
Extended Search	DECIMAL
Microsoft® SQL Server	SQL_VARIANT
Oracle (NET8 wrapper only)	LONG LONG RAW 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 (fractional_seconds_precision) WITH TIME ZONE TIMESTAMP (fractional_seconds_precision) WITH LOCAL TIME ZONE
Sybase	unichar univarchar

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

### **Insert, update, and delete restrictions on data types**

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

*Table 2. Write operation restrictions on data types*

<b>Data source</b>	<b>Unsupported data types</b>
Informix™	BLOB CLOB TEXT
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
Oracle (SQLNET wrapper only)	BLOB CLOB INTERVAL DAY (day_precision) TO SECOND (fractional_seconds_precision) INTERVAL YEAR (year_precision) TO MONTH NCHAR NVARCHAR2 TIMESTAMP (fractional_seconds_precision) WITH TIME ZONE TIMESTAMP (fractional_seconds_precision) WITH LOCAL TIME ZONE

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

Data source	Unsupported data types
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 (length 32673–64000) varchar (length 32673–64000)

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

## Unicode UTF-8 issues

### UTF-8 code page data expansion

If the federated database uses the UTF-8 code page and the data source client converts to the UTF-8 code page, the conversion might result in data expansion. For example, a one-byte character at the remote data source might be stored as two 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, then use the ALTER NICKNAME statement to make the columns wider.

### Chinese code page GB 18030

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

For Oracle data sources, you might need to add a setting to your `sqllib/cfg/db2dj.ini` file so that the Oracle client correctly translates the GB 18030 data to Unicode. The setting to add is:

```
NLS_LANG=Simplified Chinese_China.AL32UTF8
```

For Informix data source, you might need to add several settings to your `sqllib/cfg/db2dj.ini` file so that the Informix client correctly translates the GB 18030 data to Unicode. The settings to add are:

```
CLIENT_LOCALE=zh_cn.UTF8  
DB_LOCALE=zh_cn.GB18030-2000  
GL_USEGLU=1
```

## LIKE predicate issues

The LIKE predicate might not function properly. Several examples of issues with the LIKE predicate are described in the following sections. A potential workaround for these problems is to set the PUSHDOWN server option into 'N'. Setting the PUSHDOWN server option into 'N' forces the LIKE portion of an SQL statement to be processed at the federated server. This portion of the SQL statement will not be pushed down to the data source. However, 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. The predicate is not pushed down because 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 return false in DB2 but true in the other data sources.

Additionally, for Microsoft SQL Server data sources, the LIKE predicate is performing 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, as described above.

### LIKE predicate pattern matching characters and varying length character columns in DBCS databases

Pattern matching characters used in LIKE predicates of a SELECT statement might return incorrect results when used with varying length character columns that are stored in DBCS databases. These errors occur because some data sources process the pattern matching characters differently than DB2, as shown in the following tables:

*Table 3. DB2 family data sources*

DB2 family data sources Pattern matching characters	Matches a DBCS character?	Matches a SBCS character?
DBCS percent	Yes	Yes
SBCS percent	Yes	Yes
DBCS underscore	Yes	No
SBCS underscore	No	Yes



Table 4. Sybase and Oracle data sources

Sybase and Oracle data sources Pattern matching characters	Matches a DBCS character?	Matches a SBCS character?
DBCS percent	Yes	Yes
SBCS percent	Yes	Yes
DBCS underscore	Yes	Yes
SBCS underscore	Yes	Yes

Table 5. Microsoft SQL Server and Informix data sources

Microsoft SQL Server and Informix data sources Pattern matching characters	Matches a DBCS character?	Matches a SBCS character?
DBCS percent	No	No
SBCS percent	Yes	Yes
DBCS underscore	No	No
SBCS underscore	Yes	Yes

## CHAR columns containing DBCS data

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

## DRDA wrapper support

When you create a remote table on DB2 for iSeries using transparent DDL that includes a federated VARGRAPHIC data type for a column, you will receive the following error:

```
SQL3324N Column xxx has a type of VARG which is not recognized.
```

The default reverse type mapping for the federated VARGRAPHIC data type is to the remote VARG data type. The VARG data type is not supported by the DRDA wrapper. You cannot use the federated transparent DDL feature to create the remote table through DB2 for Linux, UNIX, and Windows. You must create the remote table directly on the data source and create a nickname on the federated server for the remote table.

## BioRS wrapper support

To configure the BioRS wrapper, you must use the DB2 Control Center. Information about configuring the BioRS wrapper is available on the Web at <http://www.ibm.com/software/data/integration/library>.

## Entrez wrapper support

The Entrez wrapper allows 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.

## Informix wrapper support

To use the Informix wrapper on a Windows federated server, in addition to setting the environment variables in the federated configuration file (`sqllib/cfg/db2dj.ini`) you might need to perform one or both of the following steps:

- Use the Informix **setnet32** utility to set Informix environment variables
- Add the Informix environment variables to your Windows federated server environment variables

The Informix environment variables include the required environment variables, `INFORMIXDIR` and `INFORMIXSERVER`, as well as any optional environment variables that you want to set, for example, `INFORMIXSQLHOSTS`. You might need to re-boot the federated server in order for these changes to take effect.

## Microsoft SQL Server wrapper support

### 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 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 the DB2 federated database code page.

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

### Sample code page conversion configuration (UNIX and Linux):

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

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

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

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

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

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

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

## ODBC wrapper support

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

If setting **PUSHDOWN** server option value to 'Y' causes some queries to fail during remote statement generation, remove the server option or set the **PUSHDOWN** server option to 'N' to prevent this problem.

### ODBC wrapper **DB2\_ONE\_REQUEST\_PER\_CONNECTION** restriction

For ODBC drivers that allow only one active statement on a connection, you must set the **DB2\_ONE\_REQUEST\_PER\_CONNECTION** server option to 'Y' for the data source. When you set the **DB2\_ONE\_REQUEST\_PER\_CONNECTION** server option, you must also set the **DB2\_IUD\_ENABLE** server option to 'N'.

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

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

The ODBC wrapper does not allow **INSERT**, **UPDATE**, or **DELETE** statements when **DB2\_ONE\_REQUEST\_PER\_CONNECTION** is set to 'Y'.

To determine if your ODBC driver allows only one active statement on a connection, use the **SQLGetInfo** function of the ODBC driver and check the

values returned for the SQL\_ACTIVE\_STATEMENTS or SQL\_MAX\_CONCURRENT\_ACTIVITIES InfoTypes. If the values returned are one (1), then your ODBC driver allows only one active statement on a connection.

### **Unicode restriction**

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

### **Code page conversion requirements**

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

### **Using the ODBC wrapper instead an existing data source wrapper**

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

#### **DB2 for Linux, UNIX, and Windows data sources**

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

#### **Oracle data sources on AIX<sup>®</sup> federated servers**

When you use the ODBC wrapper to access Oracle data sources, you might receive an error message when you attempt to update a nickname using a character data type comparison in a WHERE clause. Use either the NET8 or SQLNET wrapper to access Oracle data sources.

#### **Informix data sources**

Do not attempt 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 using a pass-through session or using transparent DDL. To access Informix data sources, use 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 using the CREATE INDEX statement with the SPECIFICATION ONLY clause.

### **Transparent DDL and CHAR data types**

Do not use transparent DDL to create remote tables on an ODBC data source if the tables include CHAR data types. When you create a remote table using transparent DDL that includes a CHAR data type, the remote length for the CHAR column is defined as 1. Either create the remote table natively on the data source or use a pass-through session to create the remote table. Then create the nickname for the data source table.

## **Oracle wrapper support**

### **djxlinkOracle script errors**

The `djxlinkOracle` script will fail with linkage editor errors on AIX federated servers if you do not have the AIX Base Application Development Math Library installed. You can determine if the library is installed by issuing the following AIX command:

```
ls1pp -l bos.add.libm
```

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

### **NET8 64-bit wrapper support**

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

## **Sybase wrapper support**

### **Required version of Sybase Adaptive Server Enterprise**

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

### **Calculations against SMALLINT columns**

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

This problem occurs when using both the DBLIB and CTLIB wrappers.

### **Changing VARCHAR columns to CLOB or BLOB columns**

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

### **Changing to BIGINT columns**

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

This problem occurs when using the CTLIB wrapper.

### **Column names must be unique**

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

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

### **DBCS for federated server name**

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

### **Using SELECT statements against LOB columns**

To select a LOB column the data source table must have a unique index and a TIMESTAMP column. The DBLIB wrapper can only select one LOB column per SQL statement.

### **Changing DECIMAL or NUMERIC columns to INTEGER columns**

If you alter the local column type from DECIMAL or NUMERIC to INTEGER, the DBLIB wrapper cannot process a SELECT statement on that column.

## **Teradata wrapper support**

The default forward data type mapping for Teradata GRAPHIC (or VARGRAPHIC) columns is to federated GRAPHIC (or VARGRAPHIC) columns. If the federated database uses the UTF-8 code page, you must use the ALTER NICKNAME statement to alter the local column types in the federated database system catalog. Alter GRAPHIC column types to CHAR and VARGRAPHIC column types to VARCHAR. Altering the local column types is necessary because the Teradata client supports Unicode data only in the UTF-8 encoding. In a federated UTF-8 database, DB2 would expect GRAPHIC and VARGRAPHIC columns to contain UTF-16 data.

## **XML wrapper support**

Although documented in the *IBM DB2 Information Integrator Data Source Configuration Guide*, the STREAMING nickname option is currently not supported. If you create a nickname and set this option to 'YES', an error is returned indicating that this setting is not valid. The error message is:

```
SQL1882N The "nickname" option "STREAMING" cannot be
set to "<option_value>" for "<object_name>."
```

## Accessing Microsoft Excel data sources

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

## 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. However there are some limitations when this command attempts to generate federated statements.

### DBCS wrapper and server names on Windows federated servers

On Windows federated servers, DB2LOOK does not generate the correct string for DBCS wrapper and server names. This impacts both the CREATE WRAPPER and CREATE SERVER statements.

Use a wrapper and server name that is 63 characters or less.

### CREATE TYPE MAPPING statements

There are several issues with the CREATE TYPE MAPPING statements generated by DB2LOOK:

- The remote data types that are generated from the CREATE TYPE MAPPING statement are not enclosed with quotation marks ("). If the remote data types are in lower case, you must enclose the remote data types with quotation marks manually.
- Once the DB2LOOK command detects a CREATE TYPE MAPPING for a LOB data, DB2LOOK will not generate any subsequent CREATE TYPE MAPPING statements. You need to manually create the data type mapping statements that DB2LOOK omits.

### DDL transparency statements

The DB2LOOK command will not generate the correct DDL statements for federated objects that were created using transparent DDL. These federated objects will appear as CREATE NICKNAME statements. You must manually change the generated DDL statements from CREATE NICKNAME statements to CREATE TABLE statements using the correct syntax for transparent DDL.

### The getstats tool

Because the **runstats** utility is not supported for nicknames in this release, you can download a **getstats** tool that is available from the DB2 Information Integrator Web site. The **getstats** tool gathers key statistical information that is used by the DB2 query optimizer. The information gathered includes table cardinality, column cardinality, and index firstkeycard and fullkeycard. The **getstats** tool is a temporary solution that allows you to update nickname statistics without dropping and recreating the nickname.

The **getstats** tool can be used only with the relational data sources supported by DB2 Information Integrator. The **getstats** runs only on Windows NT<sup>®</sup> and AIX.

This sample tool is provided with no warranty, formal or implied. The **getstats** tool is provided for demonstration purposes with no guarantees, obligations, or commitments from the IBM<sup>®</sup> company.

You can download the **getstats** tool from the following Web site <http://www.ibm.com/software/data/integration>. See the documentation on the Web site for more information about the tool.

### Control Center updates

#### Differences between federated online help and Control Center object names

The window names documented in the online help do not reflect the current names for some of the Control Center windows.

The following table lists changes that have been made to the window names:

*Table 6. 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 documented in the online help do not reflect the current names for some of the controls on windows in the Control Center.

The following table lists changes that have been made to the control names:

Table 7. Changed control names

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

---

## MQ user-defined function message string updates

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

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

---

## Web services issues

### Web services consumer user-defined functions

Currently, the text in the Web services section of the *DB2 Information Integrator Developer's Guide* reads,

```
"This topic describes the Web services consumer standalone tool and
the WebsSphere Studio plug-in that IBM provides to convert WSDL to
DB2 SQL functions."
```

The text should read:

```
"This topic describes the Web services consumer user-defined
functions and the WebsSphere Studio plug-in that IBM provides to
convert WSDL to DB2 SQL functions."
```

This functionality is a set of user-defined functions, not a standalone tool.

### Web services consumer messages

When working with the Web services consumer user-defined functions, messages are returned in English. Messages are not returned in other languages.

---

## Additional documentation updates

### **DB2 Information Integrator Installation Guide**

In the topic "Checking the data source environment variables", one of the restrictions on the db2dj.ini file is:

The environment variable value cannot contain white space.

This restriction has been removed.

Another restriction states:

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

This restriction should read:

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

### **DB2 Information Integrator Data Source Configuration Guide**

In the "User mapping options for federated servers" appendix, the name of the accounting option should be ACCOUNTING instead of ACCOUNTING\_STRING.

### **DB2 Information Integrator Federated Systems Guide**

In the "User mapping options for federated servers" appendix, the name of the accounting option should be ACCOUNTING instead of ACCOUNTING\_STRING.

---

## 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:

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

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

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

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

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





Printed in U.S.A.