IBM DB2 Information Integrator

Addendum to the DB2 Information Integrator Version 8.1 Release Notes

Version 8

IBM DB2 Information Integrator

Addendum to the DB2 Information Integrator Version 8.1 Release Notes

Version 8

Before using this information and the product it supports, be sure to read the general information under "Notices" on page 15.

Contents

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.

Addendum to the DB2 Information Integrator Version 8.1 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.

More information about DB2 Information Integrator support issues

For more information about DB2 Information Integrator and the latest support issues, see the DB2 Information Integrator support Web page at: http://www-3.ibm.com/software/data/integration/db2ii/support.html .

DB2 Information Integrator installation issues

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 in its name, the installation will fail. To avoid this problem, you take one of the following actions:

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

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

A Java problem that causes the license acceptance text on Russian Linux computers to be corrupted. For DB2 Information Integrator Version 8.1.2, the Russian license text that is displayed by the installer GUI 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.

Supported operating systems for DB2 Information Integrator

The supported operating systems shown in the following table supersede the list of supported operating systems in the ″Supported operating systems for DB2 Information Integrator″ topic in the *DB2 Information Integrator Installation Guide*:

Table 1. Supported operating systems for DB2 Information Integrator

Operating system	32-bit support	64-bit support
Windows NT	✔	
Windows 2000	✔	
Windows XP	✔	
Solaris Operating Environment 7, 8, 9	✔	✔
AIX 4.3	✔	
AIX 5	✔	✔
HP-UX, Version 11i	✔	✔
RedHat Linux, Version 7. See the DB2 for Linux Web page for the supported level of the Linux kernel at: http://www.ibm.com/ software/data/db2/ linux/validate	✔	

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

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.

Updated restrictions for the db2dj.ini file

The following restrictions on the db2dj.ini file in the topic, "Checking the data source environment variables," have been removed or updated:

- The restriction on the db2dj.ini file that prevents an environment variable value from containing white space has been removed.
- The restriction, "Every line must end with an ASCII carriage return (0x0D) or line feed (0x0A) character," should read, "The maximum length of any line in the db2dj.ini file is 1021 bytes. Data beyond that length is ignored."

Setting up Informix code page conversion

The topic, "Setting up Informix code page conversion," in the *DB2 Information Integrator Installation Guide*, contains the following text:

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

The text should be:

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

Setting up Oracle code page conversion

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

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

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

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

To set the NLS_LANG Oracle environment variable:

1. Open the sqllib/cfg/db2dj.ini file in a text editor. If the file does not exist, create a file named db2dj.ini.
2. In the db2dj.ini file, set the NLS_LANG environment variable to the Oracle locale that you want to use. For example:
   ```
   NLS_LANG=<Oracle_client_locale_value>
   <Oracle_client_locale_value> is the locale that you want to use.
   ```

Chinese code page GB 18030 updates

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, the Oracle wrapper sets the Oracle NLS_LANG environment variable as follows:
```
NLS_LANG=Simplified Chinese_China.UTF8
```

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

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

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

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

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

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

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

```
DBNLS=1
```

Updated data source environment variables

In addition to the data source environment variables that are listed in the topic, ″Checking the data source environment variables″ in the DB2 Information Integrator Installation Guide, the following optional environment variables are valid for Oracle data sources:

- NLS_LANG (optional)
- ORACLE_BASE (optional)

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

If you use the following configuration, you need to edit the genclntsh script and create the libclntsh.so file or your DB2 federated instance will fail when you attempt any remote operation that includes Oracle and the NET8 wrapper:

- The Oracle 9i client
- The Oracle NET8 wrapper
- A federated server that runs either the Linux, Solaris Operating Environment, or HP-UX operating system

The topics, ″Editing the Oracle genclntsh script before the federated server is set up to access Oracle data sources″ and ″Editing the Oracle genclntsh script after the federated server is set up to access Oracle data sources,″ should also include instructions for Linux operating systems. This topic gives the instructions for editing the genclntsh script on Linux. See the topics, ″Editing the Oracle genclntsh script before the federated server is set up to access

Oracle data sources″ and ″Editing the Oracle genclntsh script after the federated server is set up to access Oracle data sources,″ for instructions on editing the genclntsh script for the Solaris Operating Environment and the HP-UX operating system.

Prerequisites:

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

Procedure:

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

To prevent remote operations from failing on Linux before you install DB2 Information Integrator:
1. Edit the $ORACLE_HOME/bin/genclntsh script and add -Wl,-Bsymbolic to the link line. For example:
   ```
   LD="gcc -shared -Wl,-relax -L${ORACLE_HOME}/lib -Wl,-Bsymbolic"
   # shared library link command
   ```
2. From a command prompt, run the genclntsh script to create the libclntsh.so file.

To prevent remote operations from failing on Linux after you install DB2 Information Integrator:
1. Edit the $ORACLE_HOME/bin/genclntsh script and add -Bsymbolic to the link line. For example:
   ```
   LD="ld -v -G -b +s -L${OLIB}-Bsymbolic "     #shared library link command
   ```
2. From a command prompt, run the genclntsh script to create the libclntsh.so file.
3. From a command prompt, run the djxlinkOracle script to update the Oracle wrapper library.
4. Issue the db2iupdt command on each DB2 instance to enable federated access to the data sources.

New variable for Teradata in the db2dj.ini file

There is an additional variable, TERADATA_CHARSET, for the Teradata wrapper. This variable can be set in the db2dj.ini file and should be set to the suggested client character set. If you do not set this variable, DB2 Information Integrator detects the client character set based on the code pages of the database. However, if this variable is set, DB2 Information Integrator uses its value as the client character set. The value in the TERADATA_CHARSET

variable is not validated, but if it is not set to the correct value, the remote data source issues an error. The following values are correct values for the TERADATA_CHARSET variable:

For UNIX systems:
- KanjiEUC_0U
- HANGULKSC5601_2R4
- SCHGB2312_1T0
- TCHBIG5_1R0
- UTF8

For Windows systems:
- KanjiSJIS_0S
- HANGULKSC5601_2R4
- SCHGB2312_1T0
- TCHBIG5_1R0
- UTF8

Updated information for DB2 Universal Database documentation

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

Table 2. Updated file names for DB2 Universal Database documentation

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

Migration corrections and clarifications

For late-breaking migration issues, see the Migration Updates document available in English on the DB2 Information Integrator Web site at: www.ibm.com/software/data/integration/db2ii .

Migrating DB2 (Windows) topic updates

In the topic ″Migrating DB2 (Windows)″, the following text should be added after step 7:

On Windows, all existing instances will be migrated during the
DB2 Version 8 installation.

Space considerations for DB2 migration

The topic "Space considerations for DB2 migration", suggests that you double
the values of the `logfilsiz`, `logprimary`, and `logsecond` configuration
parameters.

However, when migrating from DataJoiner the log file space that you need
depends on the size of your database. The total size for all log files must be
twice the backup file size of the database. Use the following formula to
determine if you have sufficient log file space on your DataJoiner system:

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

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

For example, suppose that the backup file size for the DataJoiner database is
117.5 MB. Before you migrate to DB2, the total size for all log files needs to be
2 X 117.5 = 235 MB .

The `logfilsiz` parameter is 5000. The DataJoiner database has 6 `logprimary`
files and 4 `logsecond` files. Using these numbers, the total size for all log files
is:

```
( 4096 X 5000 X ( 6 + 4) ) = 204.8 MB
```

Because the total file size for all log files is less than twice the backup file size
of the DataJoiner database, migration will fail.

You can avoid this error by increasing the `logfilsiz` parameter or creating
additional log files before you migrate.

Migrating DataJoiner databases

In the topic "Migrating databases", when you migrate DataJoiner databases, if
the log file size is not large enough the following error is returned:

```
SQL1704N  Database migration failed.  Reason code "3".
```

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

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

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

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

After the migration is complete, reset the `logfilsiz`, `logprimary`, and `logsecond` parameters.

DB2 Control Center updates

db2updv8 command

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

Syntax:

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

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

```
[IBM][CLI Driver][DB2/NT] SQL0444N  Routine "GET_WRAP_CFG_C"
(specific name "SQL030325095829810") is implemented with code
in library or path "\GET_WRAP_CFG_C", function "GET_WRAP_CFG_C"
which cannot be accessed.  Reason code: "4".  SQLSTATE=42724
```

Federated client and server compatibility

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

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

If you use DB2 UDB Version 8.1 or DB2 UDB Version 8.1 FixPak 1 clients with a DB2 UDB Version 8.1.2 server, only some federated functionality is available in the DB2 Control Center. If you want to use all of the federated functionality in the DB2 Control Center, you must upgrade your clients to DB2 UDB Version 8.1.2.

Discovery window for XML data sources

The **Use streaming** option is not available in this release for the Discovery function for XML data sources. Additionally, in the corresponding online help

for this window, "Displaying objects to create nicknames and views for XML data sources", steps 7 and 8 (which refer to the **Use streaming** option) do not apply.

Federated replication restrictions

Microsoft SQL Server, Sybase, and Oracle replication targets

In this release, if you use Microsoft SQL Server and Sybase as replication targets, the longest character string that can be replicated is 255 characters. For an Oracle replication target the longest a character string that can be replicated is 4000 characters.

Federated limitations and restrictions

CREATE TABLE restrictions

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

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

The following options are valid CREATE TABLE options:
- REMOTE_SERVER
- REMOTE_SCHEMA
- REMOTE_TABNAME
- SQL_SUFFIX (DRDA wrapper only)

Avoid using invalid options on the CREATE TABLE statement.

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 work around for this problem.

Extended Search wrapper restrictions

Incorrect date format in Extended Search queries

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

The following example shows a sample of a query written using 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 using 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.

ODBC wrapper restriction

Positioned UPDATE and DELETE statements and certain searched UPDATE and DELETE statements on a nickname will fail with error SQL30090 reason code 21 if a unique index on non-nullable columns does not exist on the nickname or its corresponding remote table. To use character columns in these indexes, you must set the VARCHAR_NO_TRAILING_BLANKS column option. For more information about the VARCHAR_NO_TRAILING_BLANKS column option, see the ″Adding ODBC data sources to federated servers″ and ″Federated system INSERT, UPDATE, and DELETE restrictions″ topics in the DB2 Information Integrator library.

Teradata wrapper support

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

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

If you attempt 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.

Mapping to the federated GRAPHIC and VARGRAPHIC data types

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. These objects cannot be created because Teradata does not support GRAPHIC and VARGRAPHIC data types for the UTF-8 character set. See the Teradata documentation for more information.

Microsoft SQL Server wrapper support

MSSQL wrapper restrictions on COMMITs as CURSOR_CLOSE_ON_COMMIT set YES

The MSSQL wrapper does not support declare cursor with hold semantics. When a cursor is declared on MS SQL Server, any commits that are issued while the cursor is opened will result in the cursor being closed on the server. This results in errors on the next fetch or close cursor statement. The following scenario will return 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 close
  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

Microsoft SQL Server tables that contain triggers have insert, update, and delete restrictions. When you issue an insert, update, or delete command on a Microsoft SQL Server table through a nickname, you will receive one of the following error messages:

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

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

HMMER wrapper restrictions

In this release, the HMMER wrapper does not support data sources containing either ″Reference″ (RF) or ″Consensus structure″ (CS) annotation flags. If you receive SQL1822N ″XML Parser Error″ while using the HMMER wrapper, edit the HMMER profile database file and change all occurrences of:

```
RF     yes
CS     yes
```

to

```
RF     no
CS     no
```

Updates to the DB2 Information Integrator Data Source Configuration Guide

CREATE NICKNAME statement - Examples for XML wrapper topic updates

In the ″CREATE NICKNAME statement - Examples for XML wrapper″ topic, the example for the payment nickname view, the column name p.amount in the select clause should be p.number.

In the ″CREATE NICKNAME statement - Examples for XML wrapper″ topic, the example for the item nickname view, the columns named it.quantity and it.name in the select clause should be i.quantity and i.name.

Updates to administering Web applications

Application servers enable enterprises to develop, deploy, and integrate next-generation e-business applications. You can use application servers as tools to administer your Web applications.

DB2® Universal Database, Version 8.1.2, contains an embedded application server, referred to as the application server for DB2. Using the application server for DB2 means that you do not need to install a separate application server to run your DB2 provided Web applications on Windows®, Linux, AIX®, and Solaris. You can have only one DB2 application server in a system that has one or multiple DB2 instances.

You can install the application server for DB2 from the CD, *Java application development and Web administration tools supplement for DB2*. This CD is provided with the DB2 Universal Database™package. Issue the following command to install the application server for DB2:

```
DB2APPSERVERINSTALL
    -ASroot absolute path for the application server install
    -hostName host name
```

-**ASroot**
    The absolute path for the application server installation.

-**hostName**
    The name of the host system.

You can start and stop the application server for DB2 from the bin\ subdirectory of the application server for DB2 directory by using the following commands:

```
startServer <server> [options]
stopServer <server> [options]
```

The start and stop commands require the following parameter:

serverName
    The name of the configuration directory of the application server that you want to start.

To deploy a Java 2 Enterprise Edition (J2EE) application, the following components must be built for an enterprise e-business application:

WAR (Web Archive)
    The Web-related components (HTML, JavaScript, Java server pages (JSP))

JAR (Java Archive)
    The Java classes that make up the business logic components

EAR (Enterprise Archive)
    The JAR files plus WAR files that make up an enterprise solution

The minimum deployable unit in WebSphere Application Server 5.0 is a WAR file. If the application is developing Enterprise JavaBeans™ (EJBs), then a JAR file and an EAR file are necessary.

Notices

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

COPYRIGHT LICENSE:

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

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

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

Trademarks

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

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

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

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

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

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

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

IBM ®

Printed in USA