IBM® DB2® Warehouse Manager

Installation Guide

Version 8

GC27-1122-00

IBM® DB2® Warehouse Manager

Installation Guide

Version 8

GC27-1122-00

Before using this information and the product it supports, be sure to read the general information under *Notices.*

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

To order DB2 publications from DB2 Marketing and Sales in the United States or Canada, call 1-800-IBM-4YOU (426-4968).

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.

Contents

About this book

This book provides the information that you need to install the following Warehouse Manager components: Information Catalog Center tools, warehouse agents, warehouse transformers, Warehouse Manager connectors, and Classic Connect drivers. Be sure to read the readme file that is included with DB2 Universal Database for important updates to product information.

Who should read this book

You should read this book if you are an Information Catalog Manager administrator or if you are responsible for installing Warehouse Manager components on workstations. You should be familiar with database concepts, client/server architectures, TCP/IP connectivity, and networking concepts.

Prerequisite publications

The following publications contain information that you need before and during the installation of Warehouse Manager components:

- *Data Warehouse Center Administration Guide*
- *Information Catalog Center Administration Guide*
- *Message Reference*

Chapter 1. Overview of the DB2 Warehouse Manager components

The DB2 Warehouse Manager package contains components that enhance the warehousing capability provided by the Data Warehouse Center, which is a part of DB2 Universal Database. Before you install any DB2 Warehouse Manager components, you should be familiar with the function each component provides, and the hardware and software prerequisites for each component.

The Warehouse Manager package includes the following components:

- Data Warehouse tools
  - Warehouse transformers
  - Warehouse agent
  - Classic Connect drivers
- Information Catalog Manager tools
  - Information Catalog Center
  - Information Catalog Center for the Web
  - Information Catalog Manager Samples
  - Manage Information Catalog Wizard

The Warehouse Manager connectors CD is included but the Warehouse Manager connectors are purchased separately.

Information Catalog Manager

The Data Warehouse Center is a metadata-driven system. Metadata, or information about your data, provides administrators and business users with descriptions of the data that is stored in the data warehouse. You can create information catalogs that describe business metadata in business terms, organize the metadata into subject areas, and customize it to your workgroup or enterprise's needs. Then, you can use the Information Catalog Manager to provide a graphical representation of data relationships and object definitions for warehouse steps.

The Information Catalog Manager provides a powerful, business-oriented solution to help users locate, understand, and access enterprise data. It enables business users to view aggregations, histories, data derivations, data sources, and descriptions of data.

Warehouse Manager overview

The Information Catalog Center is the graphical interface for the Information Catalog Manager functions. The Information Catalog Manager Tools consists of the following components: The Information Catalog Center, the Manage Information Catalog Wizard, the Information Catalog Manager samples, and the Information Catalog Center for the Web.

You must run the Manage Information Catalog wizard to create or migrate your information catalogs.

You can use the Information Catalog Center to enable metadata exchange and to keep your information catalog current with the warehouse control database. Information Catalog Center functions change depending on the user's authorities. A user can simply access the information available in the information catalog. A power user has the same authority as a user but can also define objects and update or delete objects that this person has already defined. An administrator has the same authority as a power user, but is also responsible for managing the content and use of an information catalog. Administrators can also grant levels of authority to users.

The Information Catalog Center includes utilities that extract descriptive data from JDBC data sources. The Information Catalog Center also includes a sample information catalog with data. See the *Information Catalog Center Tutorial* for more information about extractors and the sample information catalog.

You can use the Information Catalog Center for the Web from any browser to access information catalogs and obtain descriptions of available data, including format, currency, owner, and location.

Related tasks:

- "Installing the Information Catalog Manager components" on page 6
- "Installing the Information Catalog Center" on page 8
- "Installing the Information Catalog Center for the Web on a Windows NT Websphere IBM HTTP Web server" on page 9
- "Installing the Information Catalog Center for the Web on an AIX Websphere IBM HTTP Web server" on page 11

Warehouse agents

Warehouse agents manage the flow of data between the data sources and the target warehouses. Warehouse agents are available for the Solaris Operating Environment and on Windows® NT, Windows 2000, Windows XP, Linux, AIX, z/OS, and iSeries™ operating systems. Warehouse agents use Open Database Connectivity (ODBC) drivers or DB2® CLI to communicate with different

databases, including text files. Several warehouse agents can handle the transfer of data between sources and target warehouses.

Related tasks:

- "Preparing to install warehouse agents" on page 15
- "Installing a Windows warehouse agent" on page 24
- "Installing the z/OS warehouse agent" on page 26
- "Installing the iSeries warehouse agent" on page 47
- "Installing the AIX, Linux, and Solaris Operating Environment warehouse agents" on page 57

Warehouse transformers

Warehouse transformers are stored procedures and user-defined functions that you can use to transform data in a warehouse step. In the Data Warehouse Center, you use steps to create and maintain a warehouse. Steps control how data is transformed into meaningful business information. You can use warehouse transformers in a step to clean, invert, and pivot data; generate primary keys and period tables; and calculate various statistics.

Related concepts:

- "JDK installation for warehouse transformers" on page 72

Related tasks:

- "Updating the environment variables on Windows for warehouse transformers" on page 73
- "Updating the environment variables on AIX for warehouse transformers" on page 74
- "Installing warehouse transformers" on page 81
- "Setting up warehouse transformers on DB2 for z/OS" on page 86
- "Updating the environment variables on the Solaris Operating Environment for warehouse transformers" on page 76
- "Updating the environment variables on Linux for warehouse transformers" on page 77

Related reference:

- "National language support for warehouse transformers" on page 91

DB2 Warehouse Manager connectors

The DB2® Warehouse Manager connectors help you extract data and metadata from e-business repositories:

- DB2 Warehouse Manager Connector for SAP R/3
- DB2 Warehouse Manager Connector for the Web

With the DB2 Warehouse Manager Connector for SAP R/3, you can add the extracted data to a data warehouse, transform it using the Data Warehouse Center, or analyze it using DB2 tools or other vendors' tools. With the Connector for the Web, you can bring clickstream data from IBM® WebSphere® Site Analyzer into a data warehouse. The DB2 Warehouse Manager connectors are included in the Warehouse Manager package but are purchased separately.

Related concepts:

- "DB2 Warehouse Manager Connector for SAP R/3" on page 93
- "DB2 Warehouse Manager Connector for the Web" on page 95

Related tasks:

- "Installing the DB2 Warehouse Manager Connector for SAP R/3" on page 94
- "Installing the DB2 Warehouse Manager Connector for the Web" on page 95

Classic Connect drivers

Classic Connect provides read access to nonrelational data stored in Information Management Systems (IMS) databases and Virtual Storage Access Method (VSAM) data sets on z/OS. It provides communication, data access, and data mapping functions so you can read nonrelational data using relational queries

Related tasks:

- "Installing the Classic Connect Drivers component (Windows NT, Windows 2000, Windows XP)" in the *Data Warehouse Center Administration Guide*

Chapter 2. Installing Information Catalog Manager tools

This chapter describes the installation process for the Information Catalog Manager tools, which includes the Information Catalog Center, the Manage Information Catalog wizard, the Information Catalog Manager samples, and the Information Catalog Center for the Web.

Preparing to install the Information Catalog Manager components

You can install the following components of the Information Catalog Manager Tools:

- Information Catalog Center (required)
- Manage Information Catalog wizard (required to initialize or migrate catalogs)
- Information Catalog Manager samples (optional)
- Information Catalog Center for the Web (optional)

An information catalog administrator will typically want to install the Information Catalog Center, the Manage Information Catalog wizard, and the Information Catalog Manager samples. The Information Catalog Center for the web should be installed on a web server to allow business users who want a web interface to access the information catalog.

You need to install the Information Catalog Center on a workstation running Windows® NT, Windows 2000, Windows XP, AIX, Solaris Operating Environment, or Linux operating system that your administrators will use to create an information catalog or to migrate existing information catalogs. Then, you can choose to install the Information Catalog Center on additional workstations (for business users) running any Windows operating system (Windows 98, Windows ME, Windows NT, Windows 2000, or Windows XP) or Unix-based operating system (Linux, Solaris Operating Environment, or AIX). These workstations must have connectivity to the workstation where your DB2® Universal Database is located. Business users also have the option of using the Information Catalog Center for the Web to access the information catalog if they prefer.

After you install the Information Catalog Center, the remote database in which an information catalog is located must be cataloged on the local system.

Related concepts:

Installing Information Catalog Center components

- “Information Catalog Manager” on page 1

Related tasks:

- “Installing the Information Catalog Manager components” on page 6
- “Installing the DB2 Warehouse Manager Connector for the Web” on page 95

Related reference:

- “Environment structure for Information Catalog Manager components” on page 97

Installing the Information Catalog Manager components

The Information Catalog Manager tools include the Information Catalog Center, the Information Catalog Manager samples, the Manage Information Catalog wizard, and the Information Catalog Center for the Web.

Prerequisites:

You can install the Information Catalog Center on additional workstations running any Windows operating system (Windows 98, Windows ME, Windows NT, Windows 2000, or Windows XP) or Unix-based operating system (Linux, Solaris Operating Environment, or AIX). These workstations must have connectivity to the server where your DB2 Universal Database is located.

DB2 Warehouse Manager components require a DB2 Version 8 client or DB2 server, depending on the component. Unless otherwise mentioned, the DB2 Version 8 client is installed when you install DB2 Warehouse Manager components, if the DB2 client is not already installed.

Procedure:

To install the Information Catalog Manager tools:

1. Insert the Warehouse Manager CD-ROM into your CD-ROM drive. The launchpad opens.
2. Click **Install** from the launchpad.
3. In the Select Features window, clear all the check boxes except the **Information Catalog Manager Tools**. You can also expand the category to select specific Information Catalog Manager Tools components.
4. Click **Next** The next window allows you to initialize a new catalog or migrate an existing catalog to the version 8 format. If you skip this step, you can use the Manage Information Catalog Wizard later to initialize or

migrate information catalogs. The metadata in your information catalogs must be in the version 8 format to use the Information Catalog Manager Tools components.

5. Click **Next** to continue the installation program. When you have reached the last window, click **Finish**. The installation program completes.
6. Run the Manage Information Catalog wizard to prepare new information catalogs or migrate existing information catalogs to the version 8 format.

Related concepts:

- "Preparing to install the Information Catalog Manager components" on page 5
- "Information Catalog Manager" on page 1

Running the Manage Information Catalog wizard

Run the Manage Information Catalog wizard to create an information catalog or migrate existing information catalogs. The wizard is only required if if you did not prepare or migrate an information catalog during installation.

Procedure:

To run the Manage Information Catalog wizard:

1. Create a database in DB2 Universal Database or locate an existing database that contains an information catalog.
2. Click **Start —> Programs —> IBM DB2 —> Set-up tools—> Manage Information Catalog wizard**. The Manage Information Catalog Wizard opens.
3. Type the information required on each page of the Manage Information Catalog Wizard, review the information you have entered on the summary page, and click **Finish**. The Manage Information Catalog wizard closes, and the information catalog is successfully initialized or migrated.

Related concepts:

- "Information Catalog Manager" on page 1

Related tasks:

- "Installing the Information Catalog Manager components" on page 6
- "Installing the Information Catalog Center" on page 8

Installing the Information Catalog Center

You can install the Information Catalog Center separately from the Information Catalog Manager Tools (on the Warehouse Manager Installation CD) on any Windows operating system (Windows 98, Windows ME, Windows NT, Windows 2000, or Windows XP) or Unix operating system (AIX Version 4, AIX Version 5, Solaris Operating Environment, or Linux)

Prerequisites:

Information catalogs can also be hosted on Version 8 or later of the following DB2 family databases:

- DB2 Universal Database for Windows NT
- DB2 Universal Database for AIX
- DB2 Universal Database for Solaris Operating Environment
- DB2 Universal Database for Linux

or on earlier versions of the following DB2 family databases:

- DB2 Universal Database for iSeries (Version 5 Release 1 or later)
- DB2 Universal Database for OS/390 and z/OS (Version 7 or later)

You also need a connection to your LAN.

You will need to install the Manage Information Catalog wizard to prepare new information catalogs or to migrate existing information catalogs to the version 8 format.

Procedure:

To install the Information Catalog Center:

1. Close all Windows programs.
2. Insert the DB2 Warehouse Manager CD-ROM into your CD-ROM drive. The auto-run feature automatically starts the DB2 Setup launchpad.
3. Click **Install Products** from the launchpad.
4. Ensure that **DB2 Warehouse Manager** is selected, then click **Next**.
5. Proceed as prompted by the DB2 Setup wizard.
6. In the list of features, make sure that the features under the **Information Catalog Manager tools** are selected, then click **Next**.
7. Proceed as prompted by the DB2 Setup wizard. Online help is available to guide you through the remaining steps.

Related concepts:

- “Information Catalog Manager” on page 1

Related tasks:

- "Installing the Information Catalog Manager components" on page 6

Installing the Information Catalog Center for the Web on a Windows NT Websphere IBM HTTP Web server

You can install the Information Catalog Center for the Web on any operating system that supports a JDBC connection to DB2 Universal Database, and runs IBM WebSphere. A web server such as IBM HTTP server is also required.

Prerequisites:

Before you install the Information Catalog Center for the Web on a Windows NT Websphere IBM HTTP Web server:

- Ensure that the IBM HTTP Web server is installed.
- Ensure that IBM Websphere Application Server is installed on the server.
- Check the Web server port number. If the Web server has a port other than 80 (the usual default), append the port number to the host name in the Web address when you configure the domain name. For example: `http://hostname:portnumber/`

Procedure:

To install the Information Catalog Center for the Web:

1. Install the Information Catalog Center for the Web files onto your Web server:
   a. Insert the DB2 Universal Database CD-ROM into the CD-ROM drive on the web server workstation. The launchpad opens.
   b. Click **Install** from the launchpad.
   c. In the Select Products window, select the **DB2 Administration Client** check box. Then click **Next**.
   d. On the **Select Type of Install** page, click **Custom**. Click **Next**.
   e. On the **Select DB2 Components** page, clear the check boxes for all components except the **Information Catalog Manager tools**.
   f. Click **Subcomponents**.
   g. On the **Select Subcomponents** page, make sure the check box for the **Information Catalog Center for the Web** is selected and all other check boxes are cleared.
   h. Click **Continue**.
   i. Make sure the **Information Catalog Manager tools** remains selected.

j. Click **Next** to continue, and click **Finish** on the last page to complete the installation.

2. Open the WebSphere Application Server Administrative Console.
   a. Click the DB2 JDBC driver in the **Resources—>JDBC Drivers** folder. If the DB2 JDBC driver has not been installed, follow the IBM Websphere instructions for installing the driver.
   b. Click **Db2JdbcDriver**.
   c. Click the **Data sources** folder and click **New**. Enter the connection information for the Information Catalog Center catalog that is to be accessed via the web, and click **OK**.
   d. Click the **Nodes** folder, and click the folder with the name of the server.
   e. Select the **Enterprise Applications** folder, and click **Install**. The Application Installation wizard opens.
   f. Install the **icweb.ear** file located in `X:\sqllib\tools\icweb\`, where X is the drive on which DB2 is installed.
   g. Replace the **db2icweb.jar** file located in `X:\WebSphere\AppServer\installedApps\icweb.ear\Web-inf\lib\` with the **db2icweb.jar** file located in `X:\sqllib\tools\icweb\`. X is the drive on which DB2 is installed.
   h. Unzip the html.zip file located in `X:\sqllib\tools\icweb\` to the `X:\WebSphere\AppServer\installedApps\icweb.ear\icweb.war\` directory. Replace existing files of the same name.
   i. Unzip the icons.zip file located in `X:\sqllib\tools\icweb\` to the `x:\WebSphere\AppServer\installedApps\icweb.ear\icweb.war\images\` directory. Replace existing files of the same name.
3. Follow the steps in ″Customizing the Information Catalog Center for the Web after installation″ before using the Information Catalog Center for the Web.

Related concepts:

- “Information Catalog Manager” on page 1

Related tasks:

- “Installing the Information Catalog Manager components” on page 6
- “Installing the Information Catalog Center” on page 8
- “Installing the Information Catalog Center for the Web on an AIX Websphere IBM HTTP Web server” on page 11

Installing the Information Catalog Center for the Web on an AIX Websphere IBM HTTP Web server

You can install the Information Catalog Center for the Web on any operating system that supports a JDBC connection to DB2 Universal Database, and runs IBM WebSphere. A web server such as IBM HTTP server is also required.

Prerequisites:

Before you install the Information Catalog Center for the Web on an AIX Websphere IBM HTTP Web server:

- Ensure that the IBM HTTP web server is installed.
- Ensure that IBM Websphere Application Server is installed on the server.
- Check the Web server port number. If the Web server has a port other than 80 (the usual default), append the port number to the host name in the web address when you configure the domain name. For example: `http://hostname:portnumber/`

Procedure:

To install the Information Catalog Center for the Web:

1. Install the Information Catalog Center for the Web files onto your web server:
   a. Insert the DB2 Universal Database CD-ROM into the CD-ROM drive on the web server workstation. The launchpad opens.
   b. Click **Install** from the launchpad.
   c. In the Select Products window, select the **DB2 Administration Client** check box. Then click **Next**.
   d. On the **Select Type of Install** page, click **Custom**. Click **Next**.
   e. On the **Select DB2 Components** page, clear the check boxes for all components except the **Information Catalog Manager tools**.
   f. Click **Subcomponents**.
   g. On the **Select Subcomponents** page, make sure the check box for the **Information Catalog Center for the Web** is selected and all other check boxes are cleared.
   h. Click **Continue**.
   i. Make sure the **Information Catalog Manager tools** remains selected.
   j. Click **Next** to continue, and click **Finish** on the last page to complete the installation.
2. Open the WebSphere Application Server Administrative Console.

a. Click the DB2 JDBC driver in the **Resources—>JDBC Drivers** folder. If the DB2 JDBC driver has not been installed, follow the IBM Websphere instructions for installing the driver.
b. Click **Db2JdbcDriver**.
c. Click the **Data sources** folder and click **New**. Enter the connection information for the Information Catalog Center catalog that is to be accessed via the web, and click **OK**.
d. Click the **Nodes** folder, and click the folder with the name of the server.
e. Select the **Enterprise Applications** folder, and click **Install**. The Application Installation wizard opens.
f. Install the **icweb.ear** file located in `/usr/sqllib/tools/icweb/`.
g. Replace the **db2icweb.jar** file located in `/usr/WebSphere/AppServer/installedApps/icweb.ear/Web-inf/lib/` with the **db2icweb.jar** file located in `/usr/sqllib/tools/icweb/`.
h. Unzip the html.zip file located in `/usr/sqllib/tools/icweb/` to the `/usr/WebSphere/AppServer/installedApps/icweb.ear/icweb.war/` directory. Replace existing files of the same name.
i. Unzip the icons.zip file located in `/usr/sqllib/tools/icweb/` to the `/usr/WebSphere/AppServer/installedApps/icweb.ear/icweb.war/images/` directory. Replace existing files of the same name.

3. Follow the steps in ″Customizing the Information Catalog Center for the Web after installation″ before using the Information Catalog Center for the Web.

Related concepts:

- “Information Catalog Manager” on page 1

Related tasks:

- “Installing the Information Catalog Manager components” on page 6
- “Installing the Information Catalog Center” on page 8
- “Installing the Information Catalog Center for the Web on a Windows NT Websphere IBM HTTP Web server” on page 9

Customizing the Information Catalog Center for the Web after installation

You can customize the Information Catalog Center for the Web to be more easily accessible to your information catalog users.

Prerequisites:

Ensure that the Information Catalog Center for the Web is properly installed before you customize it.

Before you access information catalogs through the Information Catalog Center for the Web, ensure that the database manager is started on the database servers, that the web server is started, and that IBM WebSphere Application Server is started.

Notify users of their user IDs and passwords, and inform them of the web address that they can use to access the Information Catalog Center for the Web.

Procedure:

To customize the Information Catalog Center for the Web:

1. Edit the file `Webgui.conf`. Add a link for each information catalog that you want users to access from the Information Catalog Center for the Web. Update each line to reflect your system and catalog settings. Use the following format (with examples) to update your Information Catalog Center for the Web Settings:
   - `html=C:\WebSphere\AppServer\hosts\default_host\test2\web\html`

     The system directory where the Information Catalog Center for the Web HTML files are stored.
   - `servlet=http://hostname/webapp/iccweb/`

     The URL for the application set up in IBM WebSphere.
   - `databasename=ICMSAMP`

     The Information Catalog Center database name.
   - `catalogname=<catalog name>:<description>`

     The catalog name is followed by descriptive text that will be used as a link to the catalog. For example:

     ```
     catalogname=ICM:Sample 7.1 Information Catalog
     ```
   - `mri=C:\WebSphere\AppServer\hosts\default_host\test2\servlets\mri\en_US`

     The MRI directory based on the language desired.
2. Catalog the server node and databases by using the necessary DB2 client function. From the DB2 Command Line Processor, verify that the Web server can connect successfully to the databases.

Related concepts:

- “Information Catalog Manager” on page 1

Related tasks:

- "Installing the Information Catalog Manager components" on page 6
- "Installing the Information Catalog Center for the Web on a Windows NT Websphere IBM HTTP Web server" on page 9
- "Installing the Information Catalog Center for the Web on an AIX Websphere IBM HTTP Web server" on page 11

Chapter 3. Installing warehouse agents

When you install the warehouse server with DB2 Universal Database on Windows NT, Windows 2000, Windows XP, or AIX, a default agent is installed. If you require a warehouse agent that can accept commands from more than one warehouse server or if you require a warehouse agent that is remote from the warehouse server, you can install another warehouse agent using the DB2 Warehouse Manager.

The warehouse server and all of the warehouse agents must be at the same version and fixpak level. You must remove any earlier version warehouse agents that reside on UNIX.

Warehouse agents are available for the Windows NT, Windows 2000, Windows XP, Linux, AIX, iSeries, and z/OS operating systems, and for the Solaris Operating Environment. When you install warehouse agents on Windows NT, Windows 2000, Windows XP, AIX, Linux, and the Solaris Operating Environment, the warehouse ODBC drivers and driver manager are also installed.

Preparing to install warehouse agents

The following sections describe how to prepare for warehouse agent installation.

Preparing to install warehouse agents

Warehouse agent sites must have access to warehouse sources and warehouse targets, except when the warehouse sources were defined using Client Connect. In this case, your client must catalog the DB2 server. If you choose to access warehouse sources and targets through ODBC on Windows, Linux, AIX, and the Solaris Operating Environment, you must register these source and target databases with ODBC as system DSNs. On AIX, Linux, or the Solaris Operating Environment, you must specify warehouse sources and targets in the .odbc.ini file located in the home directory of the user ID that the agent runs under. If you install the warehouse server on AIX and you use the default agent to access warehouse sources and targets, you must also specify warehouse sources and targets to the server workstation for data that is accessed by the default agent. The warehouse sources and targets must be cataloged where the warehouse server and default agent are installed.

Prerequisites:

Warehouse agents are available for these operating systems:

- Microsoft Windows NT Workstation or Windows NT Server Version 4.0 with Service Pack 6 or later.
- Microsoft Windows 2000 or Windows XP.
- IBM AIX Version 4.3.3 FixPak 2 or later.
- IBM OS/390 Version 2.6 or later.
- Linux for 32-bit Intel processors. Linux distributions with the following levels are supported: kernel 2.4.7, glibc 2.2.4.
- IBM iSeries V4R5 or later, with any prerequisite software FixPaks and PTFs. For a current maintenance list, see the readme file. To determine what software is installed on your iSeries workstation, enter `DSPSFWRSC` at an iSeries command prompt.
- Solaris Operating Environment Version 2.6 or later.

Note: Future releases of the Linux platforms may or may not be supported. See the DB2 Universal Database for Linux website for more information.

If you are using the DB2 OLAP steps, you must have either DB2 OLAP server (Version 7 or later) or Essbase (Version 6 or later) installed. For iSeries, you must have a user ID that has ALLOBJ and JOBCTL authority. This level of authority is required for both the iSeries RSTLICPGM command and the STRVWD and ENDVWD commands. Also, the user profile that starts the warehouse agent daemon should also have *PGMR (bind capability) as the User Class.

Warehouse agents on AIX, Linux, Windows, and the Solaris Operating Environment require 50 MB each of fixed disk space.

The z/OS warehouse agent requires approximately 95 MB of HFS space.

Fixed disk space is not applicable to warehouse agents on iSeries.

Procedure:

Install TCP/IP on the warehouse agent site to establish connectivity between the warehouse agent site and the warehouse server.

To serve a remote client from a DB2 server, set the DB2COMM environment variable on the DB2 server:

```
db2set -i instance DB2COMM=TCPIP
```

This setting allows the server to accept TCP/IP requests from remote clients.

Related concepts:

- "Warehouse agents" on page 2

Related tasks:

- "Configuring a database connection using the Configuration Assistant (CA)" in the *Installation and Configuration Supplement*
- "Installing a Windows warehouse agent" on page 24
- "Installing the z/OS warehouse agent" on page 26
- "Installing the iSeries warehouse agent" on page 47
- "Installing the AIX, Linux, and Solaris Operating Environment warehouse agents" on page 57

Non-DB2 source database access using ODBC

You can access several non-DB2 warehouse sources from agent sites by using ODBC drivers. The Data Warehouse ODBC driver for AIX, Linux, the Solaris Operating Environment, and Windows (ODBC driver only) are installed when you install a warehouse agent on one of these operating systems. The Data Warehouse ODBC drivers for non-DB2 sources are the DataDirect ODBC drivers provided by DataDirect Technologies, Inc.

For NFS or SNA connectivity, you need additional connectivity software. The software that you need depends on the communication protocol that is used on the agent site.

Validating the connectivity of an ODBC data source for warehouse agents

The Warehouse Manager provides a test program, called odbctest, which you can run at your Windows and UNIX warehouse agent sites to validate connectivity from your Windows and UNIX warehouse agent sites to your ODBC data sources. This program attempts to connect to the database you specify and lists the contents of the database catalog. If the test program is able to connect to the database, then the connectivity is properly set up and the warehouse agent should be able to connect. If the test program is not able to connect to the source database, then the error code that was encountered will be displayed, as will any ODBC driver messages. These messages will help you to configure and fix the connection to the source database.

Prerequisites:

If your operating system is AIX, Linux, or the Solaris Operating Environment, ensure that your current home directory has the .odbc.ini file defined and that the source to which you are trying to connect is defined in that file. If your operating system is Windows ensure that the source is defined as an ODBC system DSN.

Procedure:

The location of the odbctest program varies with the operating system.

- Windows: ...\SQLLIB\bin
- AIX: /usr/opt/db2_08_01
- Solaris Operating Environment/Linux: /opt/IBM/db2/V8.1

To validate the connectivity of your ODBC data source for warehouse agents:

1. If you are running the program from a UNIX warehouse agent site, type the following:
   - On AIX, run this command: /usr/opt/db2_08_01/bin/IWH.environment.
   - On the Solaris Operating Environment and Linux, run this command: /opt/IBM/db2/V8.1/bin/IWH.environment
2. From a Windows command prompt or UNIX command line, type odbctest <dsn> <uid> <pw>, where:
   - <dsn> is the ODBC system (Windows) database you are attempting to connect to
   - <uid> is a valid user ID to connect to the <dsn> database
   - <pw> is a valid password for user <uid>
3. To verify the connection to the system ODBC data source (called in this example 'target'), type the following from a Windows command window or from an AIX, Linux, or Solaris Operating Environment terminal on the agent site:

   ```
   odbctest target labriejj my1pw
   ```

   If the connectivity is set up properly the following should be output to the screen:

```
Operation: Enter ODBC Test Program, RETCODE = 0
Operation: Completed Initialization, RETCODE = 0
Operation: SQLAllocEnv, RETCODE = 0
Operation: SQLAllocConnect, RETCODE = 0
target labriejj my1pw
Operation: SQLConnect, RETCODE = 0
Environment variable Files\SQLLIB\LOGGING/odbctest.set not defined
Operation: Environment settings written to $(VWS_LOGGING)/odbctest.set, RETCODE = 0
(1) Operation: SQLConnectOptions, RETCODE = -1
(1) SQLSTATE = S1092
(1) SQLCODE  = -99999
(1) Error Message:
(1)         [IBM][CLI Driver] CLI0133E  Option type out of range. SQLSTATE=S1092
Operation: SQLAllocStmt, RETCODE = 0
Operation: SQLTables, RETCODE = 0
Table type=TABLE, VIEW, SYSTEM TABLE, ALIAS
IWH.APPEND, type= TABLE, remarks=
IWH.EDITIONS, type= TABLE, remarks=
IWH.MARY2, type= TABLE, remarks=
```

```
IWH.NOMAPS, type= TABLE, remarks=
LABRIEJJ.DB2STAT, type= TABLE, remarks=
LABRIEJJ.KARL1, type= TABLE, remarks=
LABRIEJJ.KARL2, type= TABLE, remarks=
LABRIEJJ.MR_DB2STAT3, type= TABLE, remarks=
SYSCAT.ATTRIBUTES, type= VIEW, remarks=
SYSCAT.BUFFERPOOLNODES, type= VIEW, remarks=
SYSIBM.SYSDUMMY1, type= SYSTEM TABLE, remarks=
SYSSTAT.COLDIST, type= VIEW, remarks=
SYSSTAT.COLUMNS, type= VIEW, remarks=
SYSSTAT.FUNCTIONS, type= VIEW, remarks=
SYSSTAT.INDEXES, type= VIEW, remarks=
SYSSTAT.TABLES, type= VIEW, remarks=
SYSSTAT.TABLES, type= VIEW, remarks=
Operation: SQLFetch, RETCODE = 100
SQLSTATE = 00000
SQLCODE  = 0
Error Message:

Operation: SQLDisconnect, RETCODE = 0
Operation: Exit ODBC Test Program, RETCODE = 0
```

If there is something wrong with either the definition or the connectivity an error should be output to the screen. In this example the user was trying to verify connectivity to a data source called myTEXT:

```
Operation: Enter ODBC Test Program, RETCODE = 0
Operation: Completed Initialization, RETCODE = 0
Operation: SQLAllocEnv, RETCODE = 0
Operation: SQLAllocConnect, RETCODE = 0
myTEXT labriejj my1pw
Operation: SQLConnect, RETCODE = -1
SQLSTATE = IM002
SQLCODE  = 0
Error Message:
         [Microsoft][ODBC Driver Manager] Data source name not found and no
            default driver specified
Environment variable Files\SQLLIB\LOGGING/odbctest.set not defined
Operation: Environment settings written to $(VWS_LOGGING)/odbctest.set, RETCODE = 0
Operation: SQLConnectOptions, RETCODE = 0
Operation: SQLAllocStmt, RETCODE = -1
SQLSTATE = 08003
SQLCODE  = 0
Error Message:
         [Microsoft][ODBC Driver Manager] Connection not open
Table type=TABLE, VIEW, SYSTEM TABLE, ALIAS
Operation: SQLDisconnect, RETCODE = -1
SQLSTATE = 08003
SQLCODE  = 0
Error Message:
        [Microsoft][ODBC Driver Manager] Connection not open
Operation: Exit ODBC Test Program, RETCODE = 0
```

Attention: Please note:

- (1) An error may sometimes occur when setting connection options. This is normal as the ODBCTEST program is a generic program for all source types and some connection options may not be available for all database types. This error will not affect the operation of this test program.
- The Windows agent must have the database catalog as a system ODBC data source. Database connectivity cannot require the use of user environment variables.
- UNIX agents must have an entry for this source in the .odbc.ini file (located in the home directory of the user ID under which the UNIX agent is executing).

Related tasks:

- "Configuring the AIX, Linux, and Solaris Operating Environment warehouse agent environment" on page 61

Ensuring that UNICODE support is available for warehouse agents

If you are using a z/OS, AIX, Linux, or Solaris Operating Environment warehouse agent, you must ensure that UNICODE support is available on your system.

Prerequisites:

The computer on which the Linux warehouse agent runs must have glibc 2.1.3-21 or greater installed.

Procedure:

On z/OS, Iconv supports conversion to and from UCS-2 in all releases of z/OS that support the z/OS warehouse agent.

To verify that UNICODE support is available on AIX, enter:

```
lslpp -al | grep -i iconv
```

Ensure that the following file sets are listed:

bos.iconv.ucs.com
UNICODE base converters

bos.iconv.ucs.pc
UNICODE converters for PC code sets

You must also install the language converter for your language. For more information, go to http://www.austin.ibm.com/resource/aix_resource/Pubs.

To verify UNICODE support on the Solaris Operating Environment, enter:

```
pkginfo | grep -i iconv
```

If UNICODE support is available, you should see the following package:

```
SUNWuiu8 -- Iconv modules for UTF-8 Locale
```

If you are using a national language version of the warehouse agent, ensure that you installed the proper code set and Iconv module.

To verify UNICODE support on Linux, enter:

```
rpm -qf /usr/bin/iconv
```

If UNICODE support is available, you should see the following package:

```
glibc-2.1.3-21 or greater
```

Attention: On Linux, `iconv` belongs to the GNU C library package. If glibc-2.1.3-21 or greater is installed on the Linux machine then UNICODE is supported. You can verify this by entering `rpm -qf /usr/bin/iconv` and glibc-2.1.3-21 or greater will be displayed as the package name.

Related concepts:

- "Warehouse agents" on page 2

Related tasks:

- "Preparing to install warehouse agents" on page 15

Warehouse agent connectivity with warehouse sources and targets

You must establish connectivity from warehouse agent sites to your warehouse sources and targets. In order to access your warehouse sources and targets, you must have the Data Warehouse Center ODBC drivers or DB2® Connect, connectivity software (such as TCP/IP), and database clients on your warehouse agent sites.

Data Warehouse Center ODBC drivers for several non-DB2 databases are installed when you install a warehouse agent. Your DB2 Universal Database™ CD-ROM also includes ODBC drivers for DB2.

Related concepts:

- "Connectivity products for warehouse sources and targets" on page 22
- "Database clients on warehouse agent sites and warehouse servers" on page 22

Related tasks:

- "Verifying TCP/IP connectivity between the warehouse server and the warehouse agent" on page 23

Connectivity products for warehouse sources and targets

Install the connectivity products that are needed to access remote warehouse sources and targets. This might include TCP/IP or NetBIOS customization, SNA Client for Windows, or SNA Server for Windows.

You need to install the connectivity products on the warehouse server (if you are using a local warehouse agent) and on your warehouse agent sites. For example, you can access a DB2® family database through TCP/IP or NetBIOS. You can also access a DB2 z/OS™ database through TCP/IP or NetBIOS if a DB2 Connect™ gateway is accessible on the network with connectivity to the host database that you need. Client Connect users must establish connectivity to warehouse sources and targets on the client, since Client Connect uses JDBC and does not make use of the warehouse server or agent.

Related concepts:

- "Warehouse agent connectivity with warehouse sources and targets" on page 21
- "Database clients on warehouse agent sites and warehouse servers" on page 22

Related tasks:

- "Defining a warehouse target" in the *Data Warehouse Center Administration Guide*
- "Verifying TCP/IP connectivity between the warehouse server and the warehouse agent" on page 23
- "Defining DB2 warehouse sources" in the *Data Warehouse Center Administration Guide*

Related reference:

- "Supported non-DB2 data sources" in the *Data Warehouse Center Administration Guide*

Database clients on warehouse agent sites and warehouse servers

Install the client component of a remote non-DB2 database wherever you want to access a remote database (if you are using a local warehouse agent).

For the DB2® family of databases, the necessary DB2 client function is included with the DB2 Universal Database.

Related concepts:

- "Warehouse agent connectivity with warehouse sources and targets" on page 21

- "Connectivity products for warehouse sources and targets" on page 22

Related tasks:

- "Verifying TCP/IP connectivity between the warehouse server and the warehouse agent" on page 23

Verifying TCP/IP connectivity between the warehouse server and the warehouse agent

You can test to ensure that there is connectivity between workstations.

Procedure:

Test TCP/IP connectivity between the following workstations:

- Warehouse server to the warehouse agent sites
- Warehouse agent sites to the warehouse server
- Data Warehouse Center administrative client to the warehouse server
- Warehouse server to the Data Warehouse Center administrative client

To test the connectivity, enter `ping` *`hostname`* where *hostname* is the TCP/IP host name of the warehouse server, Data Warehouse Center administrative client, or warehouse agent site.

On AIX, Linux, and Windows operating systems, you will see several messages in the window that verify the TCP/IP connection, similar to those shown in the following example.

```
[C:\]ping dgntserv2.stl.ibm.com
PING dgntserv2.stl.ibm.com: 56 data bytes
64 bytes from 9.112.46.127: icmp_seq=1. time=0. ms
64 bytes from 9.112.46.127: icmp_seq=2. time=0. ms
64 bytes from 9.112.46.127: icmp_seq=3. time=0. ms
64 bytes from 9.112.46.127: icmp_seq=4. time=0. ms

----dgntserv2.stl.ibm.com PING Statistics----
5 packets transmitted, 4 packets received, 20% packet loss
round-trip (ms)  min/avg/max = 0/0/0
```

Figure 1. Sample response from PING command

On the Solaris Operating Environment, the ping command will return the following information:

```
host is alive
```

For a workstation with AIX or Linux warehouse agents, ensure that you can ping the workstation by using only the host name without the local domain

name. For example, enter `ping dgntserv2` instead of `ping dgntserv2.stl.ibm.com`. You might need to add an entry in the /etc/hosts file such as:

```
123.45.67.89       dgntserv2       dgntserv2.stl.ibm.com
```

Test ODBC connectivity between any non-DB2 database clients and servers. If you are using a Windows warehouse agent, verify that connectivity can be established as a Windows system process. You can use the odbc test program shipped with the warehouse agent to test proper connectivity and configuration.

Related concepts:

- "Warehouse agent connectivity with warehouse sources and targets" on page 21
- "Connectivity products for warehouse sources and targets" on page 22
- "Database clients on warehouse agent sites and warehouse servers" on page 22

Related tasks:

- "Setting up the Classic Connect ODBC driver and warehouse access" on page 37
- "Testing for bidirectional communication between the iSeries warehouse agent and the warehouse server" on page 52
- "Determining the source of connectivity errors for the iSeries warehouse agent" on page 53

Related reference:

- "Post-installation considerations for iSeries warehouse agents" on page 49

Installing the Windows warehouse agent

The following section describes how to install the Windows warehouse agent.

Installing a Windows warehouse agent

When you install the warehouse server with DB2 Universal Database on Windows NT, Windows 2000, or Windows XP, a default agent is installed. If you require a warehouse agent that is remote from the warehouse server, you can use the DB2 Warehouse Manager to install another warehouse agent.

Prerequisites:

Before you install a Windows warehouse agent, you must prepare for the installation.

Procedure:

To install a Windows warehouse agent:

1. Close all Windows programs.
2. Insert the DB2 Warehouse Manager CD-ROM into your CD-ROM drive. The auto-run feature automatically starts the DB2 Setup launchpad.
3. Click **Install Products** from the launchpad.
4. Ensure that **DB2 Warehouse Manager** is selected, then click **Next**.
5. Proceed as prompted by the DB2 Setup wizard.
6. In the list of features, select **Warehouse agent**, then click **Next**.
7. Proceed as prompted by the DB2 Setup wizard. Online help is available to guide you through the remaining steps.

After you install the warehouse agent, register your source and target databases as system DSNs.

Related concepts:

- "Warehouse agents" on page 2
- "Warehouse agent connectivity with warehouse sources and targets" on page 21

Related tasks:

- "Preparing to install warehouse agents" on page 15

Related reference:

- "Non-DB2 source database access using ODBC" on page 17

Installing and configuring the z/OS warehouse agent

The following sections describe how to install and configure the z/OS warehouse agent.

Overview of the z/OS warehouse agent

DB2® Warehouse Manager for z/OS™ includes a z/OS warehouse agent. You can use the agent to communicate between DB2 Universal Database™ for z/OS and other databases, including DB2 databases on other operating systems and non-DB2 databases. The warehouse agent can communicate with supported data sources that use an Open Database Connectivity (ODBC) connection. The warehouse agent runs on Unix Systems Services, requires OS/390® Version 2 Release 7 or later, and is compatible with DB2 for OS/390 Versions 5 and 6, and DB2 for z/OS and OS/390 Version 7.

The z/OS warehouse agent supports the following tasks:

- Copy data from a DB2 database source to a DB2 database target
- Produce sample contents from a table or file
- Run user-defined programs
- Access non-DB2 databases through DB2 Relational Connect
- Access VSAM or IMS™ data through Cross Access Classic Connect
- Run DB2 Universal Database for z/OS utilities
- Run the apply job for IBM® Data Propagator
- Execute warehouse-supplied transformers

Related tasks:

- "Installing the z/OS warehouse agent" on page 26
- "Starting the z/OS warehouse agent daemon" on page 31
- "Starting the agent daemon as a z/OS started task" on page 41
- "Running multiple warehouse agent daemons on one z/OS subsystem" on page 42

Related reference:

- "z/OS warehouse agent support for Trillium user-defined program steps" on page 46

Installing the z/OS warehouse agent

The *DB2 Warehouse Manager Program Directory* is included with the z/OS warehouse agent and contains information about requirements for storage, programming, and your driving and target systems.

Prerequisites:

The z/OS warehouse agent requires the following software:

- OS/390 Version 2.7 or later
- DB2 Universal Database for OS/390 Version 5 or later
- OS/390 UNIX System Services

The following table shows the APARs that you must apply to your DB2 subsystem before installing the z/OS agent.

Table 1. Requisite APARs

DB2 Universal Database Version	Requisite APAR
DB2 for OS/390 Version 5	PQ36586, PQ46261
DB2 for OS/390 Version 6	PQ36585, PQ46261
DB2 for z/OS and OS/390 Version 7	PQ36585, PQ46261

Restrictions:

The z/OS warehouse agent requires the following access privileges:

- User ID with Authorized Program Facility (APF) authority (to work with UNIX System Services)
- EXECUTE authority on your ODBC plan
- READ and WRITE authority to the logging and ODBC trace directories (if running with logging or ODBC trace on)

Procedure:

To install the z/OS warehouse agent:

1. Unload the code for the warehouse agent from the SMP/E tape. The program directory is supplied with the product.

   One of the target libraries that results from the SMP/E installation is a UNIX Systems Services directory, which contains the warehouse agent executables. The default name of this directory is /usr/lpp/DWC. This directory contains the executable files for the warehouse agent.
2. Set up a home directory within Unix Systems Services for each user ID that will start the agent daemon. The agent requires environment variables, which point to various DB2 libraries and output directories. The best way to set these variables is to put them in a `.profile` file in the home directory of the user who will start the agent daemon. That way, when the user signs on, the `.profile` file runs automatically and sets the environment variables.

   This example shows the contents of a sample .profile file:

   ```
   export VWS_LOGGING=/usr/lpp/DWC/logs
   export VWP_LOG=/usr/lpp/DWC/vwp.log
   export VWS_TEMPLATES=/usr/lpp/DWC
   export DSNAOINI=/u/userid/dsnaoini
   export LIBPATH=/usr/lpp/DWC/ :$LIBPATH
   export PATH=/usr/lpp/DWC/ :$PATH
   export STEPLIB=DSN710.SDSNEXIT:DSN710.SDSNLOAD
   ```
3. Set up kernel and daemon connections. To set up these connections, add the following lines to your `/etc/services`file or `TCPIP.ETC.SERVICESFILE`:

   ```
   vwkernel 11000/tcp
   vwd 11001/tcp
   vwlogger 11002/tcp
   ```

   To set up connections between the z/OS warehouse agent and databases, add any remote databases to your z/OS communications database.

   This example shows sample communications database inserts to connect a database that is running on Windows NT to a database on z/OS:

```
INSERT INTO SYSIBM.LOCATIONS
(LOCATION, LINKNAME, PORT)
VALUES
('NTDB','VWNT704','60002');
INSERT INTO SYSIBM.IPNAMES
(LINKNAME, SECURITY_OUT, USERNAMES, IPADDR)
VALUES
('VWNT704', 'P', 'O', 'VWNT704.STL.IBM.COM');
INSERT INTO SYSIBM.USERNAMES
(TYPE, AUTHID, LINKNAME, NEWAUTHID, PASSWORD)
VALUES
('O', 'MVSUID', 'VWNT704', 'NTUID', 'NTPW');
```

4. Bind ODBC to databases locally and remotely. Because the z/OS warehouse agent uses ODBC to communicate with DB2, you must bind your ODBC plan to all of the local and remote databases that your agent will access.

   The following example shows some sample bind package statements for a local DB2 for z/OS database. The example assumes that your local system is named *DWC6*, and that you want to use the name DWC6CLI for your ODBC package name on that subsystem:

   ```
   DSN SYSTEM(DWC6)
   BIND PACKAGE (DWC6CLI) MEMBER(DSNCLICS) ISO(CS)
   BIND PACKAGE (DWC6CLI) MEMBER(DSNCLINC) ISO(NC)
   BIND PACKAGE (DWC6CLI) MEMBER(DSNCLIRR) ISO(RR)
   BIND PACKAGE (DWC6CLI) MEMBER(DSNCLIRS) ISO(RS)
   BIND PACKAGE (DWC6CLI) MEMBER(DSNCLIUR) ISO(UR)
   BIND PACKAGE (DWC6CLI) MEMBER(DSNCLIQR) ISO(QR)
   BIND PACKAGE (DWC6CLI) MEMBER(DSNCLIMS)
   BIND PACKAGE (DWC6CLI) MEMBER(DSNCLIC1)
   BIND PACKAGE (DWC6CLI) MEMBER(DSNCLIC2)
   BIND PACKAGE (DWC6CLI) MEMBER(DSNCLIF4)
   ```

   The following example shows some sample bind package statements for a remote DB2 for z/OS database, where *REMLOC* is the location name of the remote DB2 for z/OS system that is registered in the Communications Database:

   ```
   DSN SYSTEM(DWC6)
   BIND PACKAGE (REMLOC.DWC6CLI) MEMBER(DSNCLICS) ISO(CS)
   BIND PACKAGE (REMLOC.DWC6CLI) MEMBER(DSNCLINC) ISO(NC)
   BIND PACKAGE (REMLOC.DWC6CLI) MEMBER(DSNCLIRR) ISO(RR)
   BIND PACKAGE (REMLOC.DWC6CLI) MEMBER(DSNCLIRS) ISO(RS)
   BIND PACKAGE (REMLOC.DWC6CLI) MEMBER(DSNCLIUR) ISO(UR)
   BIND PACKAGE (REMLOC.DWC6CLI) MEMBER(DSNCLIQR) ISO(QR)
   BIND PACKAGE (REMLOC.DWC6CLI) MEMBER(DSNCLIMS)
   BIND PACKAGE (REMLOC.DWC6CLI) MEMBER(DSNCLIC1)
   BIND PACKAGE (REMLOC.DWC6CLI) MEMBER(DSNCLIC2)
   BIND PACKAGE (REMLOC.DWC6CLI) MEMBER(DSNCLIF4)
   ```

   The following example shows some sample bind package statements for a DB2 Universal Database that is running on Windows NT:

```
BIND PACKAGE (NTDB.DWC6CLI) MEMBER(DSNCLICS) ISO(CS)
BIND PACKAGE (NTDB.DWC6CLI) MEMBER(DSNCLINC) ISO(NC)
BIND PACKAGE (NTDB.DWC6CLI) MEMBER(DSNCLIRR) ISO(RR)
BIND PACKAGE (NTDB.DWC6CLI) MEMBER(DSNCLIRS) ISO(RS)
BIND PACKAGE (NTDB.DWC6CLI) MEMBER(DSNCLIUR) ISO(UR)
BIND PACKAGE (NTDB.DWC6CLI) MEMBER(DSNCLIC1)
BIND PACKAGE (NTDB.DWC6CLI) MEMBER(DSNCLIC2)
BIND PACKAGE (NTDB.DWC6CLI) MEMBER(DSNCLIQR)
BIND PACKAGE (NTDB.DWC6CLI) MEMBER(DSNCLIF4)
BIND PACKAGE (NTDB.DWC6CLI) MEMBER(DSNCLIV1)
BIND PACKAGE (NTDB.DWC6CLI) MEMBER(DSNCLIV2)
```

Here is a sample bind statement to bind the ODBC packages for all locations together in a plan:

```
DSN SYSTEM(DWC6)
BIND PLAN(DWC6CLI) PKLIST(*.DWC6CLI.* )
```

When you complete binding the ODBC packages, verify that the DSNAOINI environment variable in your .profile file points to the ODBC initialization file that uses the ODBC plan that you just bound.

For example, if the ODBC plan is named DWC6CLI and the local system is named DWC6, the ODBC initialization file should contain the following information:

```
;SUBSYSTEM stanza
[DWC6]
MVSATTACH=CAF
PLANNAME=DWC6CLI
```

5. Set up your ODBC initialization file. A sample ODBC initialization file, INISAMP, is included in the usr/lpp/DWC/ directory. You can edit this file to work with your system, or you can create your own file.

   To ensure that the file works correctly, verify that it is properly configured:

   - The DSNAOINI environment variable must point to the initialization file.
   - The file must include CONNECTTYPE=2 in the common stanza.
   - The PLANNAME value in the subsystem stanza must be the name of the ODBC plan that you bound in the last step.
   - The file must include MVSATTACHTYPE=CAF in the subsystem stanza.

   Ensure that you have a data source stanza for your DB2 system. It must specify the location name of the local DB2 system.

   This is an example of a DSNAOINI file:

```
[COMMON]
MVSDEFAULTSSID=DWC6
CONNECTTYPE=2
;APPLTRACE=1
```

```
;APPLTRACEFILENAME=/usr/lpp/DWC/logs/application.CLITRACE
;TRACEFLUSH=1
;Example SUBSYSTEM stanza for V71A subsystem
[DWC6]
MVSATTACHTYPE=CAF
PLANNAME=DWC6CLI
;DATA SOURCE stanza for ISC710P1 data source
[ISC710P1]
```

To turn ODBC tracing on, remove the semicolons in the first column of the COMMON section.

6. The z/OS warehouse agent is a daemon process. Because the agent daemon uses the _password( ) function, define these agent executable programs to RACF program control:

   ```
   libtls4d.dll
   iwhcomnt.dll
   vwd
   ```

   To define the executable programs to RACF program control, change to the directory where the Data Warehouse Center executable files are stored, and run the following commands:

   ```
   extattr +p libtls4d.dll
   extattr +p iwhcomnt.dll
   extattr +p vwd
   ```

   To use the `extattr` command with the `+p` parameter, you must have at least READ access to the BPX.FILEATTR.PROGCTL FACILITY class.

   This example shows the RACF command that is used to give this permission to user ID JOEUSER:

   ```
   RDEFINE FACILITY BPX.FILEATTR.PROGCTL UACC(NONE)
   PERMIT BPX.FILEATTR.PROGCTL CLASS(FACILITY) ID(JOEUSER)
   ACCESS(READ)
   SETROPTS RACLIST(FACILITY) REFRESH
   ```

7. Start the agent daemon. Use Telnet to connect to OS/390 UNIX Systems Services through the z/OS and OS/390 host name and USS port.

Related concepts:

- "Sample contents of DB2 tables and flat files for z/OS and OS/390" on page 36
- "Warehouse agents" on page 2
- "Overview of the z/OS warehouse agent" on page 25

Related tasks:

- "Starting the z/OS warehouse agent daemon" on page 31
- "Starting the agent daemon as a z/OS started task" on page 41

Starting the z/OS warehouse agent daemon

In general, the user ID that starts the z/OS warehouse agent daemon should not be permitted to use the BPX.DAEMON FACILITY class profile. The agent daemon issues the `setuid()` command, which is a controlled function. When a user ID runs a program that issues a command that is a controlled function, the UNIX Systems Services kernel checks to see if the user ID has been permitted to use the BPX.DAEMON facility class profile. If it has, then the kernel checks whether all programs loaded into the address space have been defined to Program Control. If an uncontrolled program has been loaded, then the address space is marked dirty. In that case, the program cannot execute any controlled functions such as `setuid()`. When this is the case, an EMVSERR return code will be displayed with a JRENVIRTY reason code. The message: ″BPXP014I ENVIRONMENT MUST BE CONTROLLED FOR DAEMON (BPX.DAEMON) PROCESSING″ means that:

- The user ID that started the agent daemon has been permitted to use the BPX.DAEMON facility class profile
- A command was issued that was a controlled function
- At least one uncontrolled program has been loaded

Procedure:

To start the daemon in the foreground, enter `vwd` on a command line under a UNIX Systems Services shell.

To start the daemon in the background, enter

```
vwd>/usr/lpp/DWC/logs/vwd.log 2>&1 &
```

on a command line under a UNIX Systems Services shell, where `/usr/lpp/DWC/logs/vwd.log` is the name of the path and the file where you want to pipe the daemon output.

To verify that the z/OS warehouse agent daemon is running, enter `ps -e | grep vwd` on a UNIX shell command line.

Or, enter `D OMVS,a=all` on the z/OS console and search for the string `vwd`.

Related concepts:

- “Warehouse agents” on page 2
- “Overview of the z/OS warehouse agent” on page 25

Related tasks:

- “Installing the z/OS warehouse agent” on page 26
- “Starting the agent daemon as a z/OS started task” on page 41

DB2 Warehouse Manager user-defined programs

The DB2® Warehouse Manager for z/OS™ package includes the following user-defined programs:

- VWPFTP runs an FTP command file.
- VWPMVS submits a JCL jobstream.
- VWPRCPY copies a file using FTP.
- XTClient is a client trigger program.

In addition, you can create user-defined programs and stored procedures in the Data Warehouse Center. The z/OS warehouse agent supports any executable programs that run under UNIX® Systems Services.

A user-defined program is assigned to one or more steps. When you run a user-defined program, the following actions occur:

- The agent runs the user-defined program.
- The user-defined program returns a return code and a feedback file to the agent.
- The agent returns the results to the kernel.

Use the VWP_LOG environment variable to define a directory where the user-defined programs can write output.

To use a user-defined program to transfer a job with FTP, you must first create the JCL and data that you want to submit. The job name in the JCL must be USERID*X*, where *X* is a 1-character letter or number (example: JOEUSER*A*). The output class for the MSGCLASS and SYSOUT files that are contained in your JCL must specify a JES-held output class.

The maximum LRECL for the submitted job is 254 characters. JES scans only the first 72 characters of JCL.

Related concepts:

- "z/OS transformers" on page 84
- "Overview of the z/OS warehouse agent" on page 25

Related tasks:

- "Installing the z/OS warehouse agent" on page 26
- "Reducing the number of characters in the environment variable data set for the z/OS warehouse transformers" on page 88

Related reference:

- "z/OS warehouse agent support for Trillium user-defined program steps" on page 46

Scheduling warehouse steps with the trigger program (XTClient)

Use the trigger program to schedule warehouse steps from z/OS or OS/390. Either you or an OS/390 job scheduler can submit a job that triggers a step in the Data Warehouse Center. If the step is successful, the trigger step in the JCL returns a return code of 0.

Prerequisites:

You must have Java Development Kit (JDK) 1.1.8 or later installed on your Unix Systems Services to use the trigger program.

Procedure:

To start the trigger program, start XTServer on the computer where your warehouse server is running.

When XTServer is started on Windows, either start the XTClient on Unix System Services or by using batch JCL.

This is a sample JCL to start the trigger:

```
//DBA1A JOB 1,'XTCLIENT',CLASS=A,MSGCLASS=H,
// MSGLEVEL=(1,1),REGION=4M,NOTIFY=&SYSUID
//*********************************************
//* submit iwhetrig
//*********************************************
//BRADS EXEC PGM=BPXBATCH,
// PARM=('sh cd /usr/lpp/DWC/; java XTClient 9.317.171.133 1100x
// 9 drummond pw bvmvs2nt 1 1 100')
//STDOUT DD PATH='/tmp/xtclient.stdout',
// PATHOPTS=(OWRONLY,OCREAT),
// PATHMODE=SIRWXU
//STDERR DD PATH='/tmp/xtclient.stderr',
// PATHOPTS=(OWRONLY,OCREAT),
// PATHMODE=SIRWXU
//
```

The previous sample JCL shows that the parameters continue to a new line. Type the parameters up to column 71, put an `X` in column 72 and continue in column 16 on the next line. The first part of the parameter (`cd /usr/lpp/DWC/;`) is a statement that changes to the directory where the z/OS warehouse agent is installed.

The second part of the parameter starts XTClient and passes the following parameters:

- Your DWC server host name or IP address
- Your DWC server port (normally 11009)
- Your DWC user ID
- Your DWC password
- The name of the step to run
- A DWC server command, where:
  - 1 = populate the step
  - 2 = promote the step to test mode
  - 3 = promote the step to production mode
  - 4 = demote the step to test mode
  - 5 = demote the step to development mode
- The option whether to wait for the step completion, where 1= yes and 0 = no
- The maximum number of rows (use 0 or blank to fetch all rows)

Related concepts:

- "z/OS transformers" on page 84

Related tasks:

- "Installing the z/OS warehouse agent" on page 26

Related reference:

- "z/OS warehouse agent support for Trillium user-defined program steps" on page 46

Changing the Data Warehouse Center template for FTP support

The Data Warehouse Center installs a JCL template for transferring files using FTP. If you plan to use the FTP commands GET or PUT to transfer files from a z/OS host to another remote host, you need to change the account information in the JCL template for your z/OS system.

Procedure:

To change the Data Warehouse Center template for FTP support:

1. Log on with an ID that has authority to copy and update files in the `/usr/lpp/DWC` directory.
2. Find the `ftp.jcl` file and copy it with the new file name *systemname*`.ftp.jcl`, where *systemname* is the name of the z/OS system.
3. Create a copy of this file for each z/OS system on which you plan to run steps which submit JCL, such as, VWPMVS.

   For example, if you want to run a step which submits JCL on STLMVS1, create a copy of the file named *STLMVS1*`.ftp.jcl`.
4. Use a text editor to customize the JCL to meet your requirements.
5. Change the account information to match the standard account information for your OS/390 MVS system. Do not change any parameters that are contained in brackets, such as `[USERID]` and `[FTPFILE]`. The brackets are the hexadecimal characters X'AD' and X'BD', respectively. If your TSO terminal type is not set to 3278A in SPF Option 0, these values might display as special characters rather than as brackets. This is not a problem if you do not change the X'AD' or the X'BD', or any of the data that is between the characters.
6. Update the environment variable VWS_TEMPLATES to point to the directory of the copied template file.

   The Data Warehouse Center includes this sample JCL template:

```
//[USERID]A JOB , 'PUT/GET',
//* THE NAME OF THE JOB MUST BE THE NAME OF THE
//* MVS USER ID FOLLOWED BY ONE ALPHANUMERIC
//* CHARACTER.  THIS IS A RESTRICTION OF FTP/MVS.
//  CLASS=A,
//  USER=&SYSUID,
//  NOTIFY=&SYSUID,
//  TIME=(,30),
//  MSGCLASS=H
```

```
//STEP1 EXEC PGM=FTP,PARM='( EXIT'
//INPUT DD DSN=[FTPFILE],DISP=SHR
//OUTPUT DD SYSOUT=*
//SYSPRINT DD SYSOUT=*
```

Related concepts:

- "DB2 Warehouse Manager user-defined programs" on page 32
- "z/OS transformers" on page 84
- "Overview of the z/OS warehouse agent" on page 25

Related tasks:

- "Installing the z/OS warehouse agent" on page 26

Sample contents of DB2 tables and flat files for z/OS and OS/390

Using the z/OS™ warehouse agent, you can display a sample of the contents of DB2® tables or view the contents of any file (with or without column definitions). You can also produce sample contents of IMS™ or VSAM files with Classic Connect using the z/OS warehouse agent. Sampling contents of flat files is a limited capability. The warehouse displays only character data from flat files. In addition, when you set the file as a source, you must define the individual fields in the file definition so that the agent can display the data at the correct offset and length for each field.

The agent also has a limited capability to display Unix Systems Services and native z/OS sequential files. For these files, the warehouse displays only character data. In addition, you must define the individual fields in the file definition so that the agent can display the data at the correct offset and length for each field.

Related concepts:

- "DB2 Warehouse Manager user-defined programs" on page 32
- "Overview of the z/OS warehouse agent" on page 25

Related tasks:

- "Installing the z/OS warehouse agent" on page 26
- "Running DB2 for z/OS utilities" on page 39

Access to databases outside of the DB2 family with the z/OS warehouse agent

To access non-DB2 Universal Database systems, the z/OS™ warehouse agent uses DB2® Relational Connect. DB2 Relational Connect enables the warehouse agent to access a non-DB2 family database.

DB2 Relational Connect can access Oracle, Sybase, Microsoft® SQL Server, Teradata, and any other database that has an ODBC driver that runs on Windows®, AIX, Linux, or the Solaris Operating Environment.

Related concepts:

- "z/OS transformers" on page 84
- "Warehouse transformers" on page 3

Related tasks:

- "Installing warehouse transformers" on page 81

Access to IMS and VSAM on OS/390 or z/OS

The z/OS™ warehouse agent can access IMS™ and VSAM through the Classic Connect ODBC driver. With Classic Connect, you can set up a DB2-like definition of IMS and VSAM data sets, and then access them using ODBC. You must purchase and install Classic Connect separately from the warehouse agent.

The z/OS warehouse agent loads the correct ODBC driver based on whether a request is directed to Classic Connect or DB2. If you are accessing a DB2® source, the agent loads the DB2 ODBC driver. If you are accessing a VSAM or IMS source, the agent loads the Classic Connect ODBC driver. The agent request is then processed.

Related concepts:

- "z/OS transformers" on page 84
- "Warehouse transformers" on page 3

Related tasks:

- "Setting up warehouse transformers on DB2 for z/OS" on page 86
- "Setting up the Classic Connect ODBC driver and warehouse access" on page 37

Setting up the Classic Connect ODBC driver and warehouse access

With Classic Connect, you can view a single file or a portion of a file as one or more relational tables. You can map the IMS and VSAM data for Classic Connect to access it. You can map the data manually or use the Microsoft Windows Classic Connect nonrelational data mapper.

Procedure:

To set up the Classic Connect ODBC driver and warehouse access:

1. Install Classic Connect Data Server on OS/390.

2. Optional: Install the Classic Connect Data Mapper product on Windows.
3. Define the Classic Connect logical tables so that Classic Connect can access data relationally. You can use the data mapper to create the definitions for IMS and VSAM structures, or create the definitions manually.
4. Set up access to your warehouse:
   a. Create a Classic Connect `.ini` file. A sample Classic Connect application configuration file `cxa.ini` is in the `/usr/lpp/DWC/` directory:

      ```
      * national language for messages
      NL = US English
      * resource master file
      NL CAT = usr/lpp/DWC/v4r1m00/msg/engcat
      FETCH BUFFER SIZE = 32000
      DEFLOC = CXASAMP
      USERID = uid
      USERPASSWORD = pwd
      DATASOURCE = DJX4DWC tcp/9.112.46.200/1035
      MESSAGE POOL SIZE = 1000000
      ```

   b. Update the DATASOURCE line in the `.ini` file. This line contains a data source name and a protocol address. The data source name must correspond to a Query Processor name that is defined on the Classic Connect Data Server, which is located in the QUERY PROCESSOR SERVICE INFO ENTRY in the data server configuration file. The protocol address is in the same file in the TCP/IP SERVICE INFO entry. The USERID and USERPASSWORD in this file are used when defining a warehouse data source.
   c. In your .profile file, export the CXA_CONFIG environment variable to your Classic Connect ODBC driver program files, which are usually in the same directory as your `.ini` file.
   d. In your .profile file, update your LIBPATH environment variable to include the path to your Classic Connect ODBC driver program files, which are usually in the same directory as your `.ini` file.
   e. Optional: Verify the installation with the test program CXASAMP. Enter `cxasamp` from the directory that contains your `.ini` file. The location `/uid/pwd` is the data source *name/userid/userpassword* that is defined in your `.ini` file.
   f. Define a data source to the warehouse in the same way that you define any DB2 data source.

   You do not need to update your `dsnaoini` file because DB2 for z/OS and OS/390 does not have a driver manager. The driver manager for Classic Connect is built into the z/OS warehouse agent.

Related concepts:

- "z/OS transformers" on page 84
- "Access to IMS and VSAM on OS/390 or z/OS" on page 37

- "Warehouse transformers" on page 3

Related tasks:

- "Installing warehouse transformers" on page 81

Running DB2 for z/OS utilities

DSNUTILS is a DB2 for z/OS and OS/390 stored procedure that runs in a WLM and RRS environment. The DB2 Warehouse Manager provides an interface to DSNUTILS so that you can include DB2 utilities in Data Warehouse Center steps. You can use DSNUTILS to run any installed DB2 utilities by using the user-defined stored procedure interface. The DB2 for z/OS and OS/390 LOAD, REORG, and RUNSTATS utilities have property sheets that you can use to change how the utility runs. In addition, the UTILITY step has a property sheet that contains a list of the 41 DSNUTILS parameters with entry fields for each parameter, so that you can run any utility that DB2 for z/OS supports.

Procedure:

To run DB2 for z/OS utilities:

1. Run the DSNTIJSG job when you install DB2 to set and bind the DSNUTILS stored procedure. Ensure that the definition of DSNUTILS includes PARAMETER STYLE GENERAL.
2. Enable the WLM-managed stored procedures.
3. Set up your RRS and WLM environments.
4. Optional: Run the sample batch DSNUTILS programs supplied by DB2.
5. Bind the DSNUTILS plan with the DSNCLI plan so that the ODBC can call the stored procedure:

   ```
   BIND PLAN(DSNAOCLI) PKLIST(*.DSNAOCLI.*, *.DSNUTILS.*)
   ```

6. Set up a step using the Data Warehouse Center and run the step. The population type must be APPEND. If it is not, the Data Warehouse Center deletes everything in the table before it runs the utility.

Related concepts:

- "z/OS transformers" on page 84
- "Warehouse transformers" on page 3

Related tasks:

- "Installing the z/OS warehouse agent" on page 26

Related reference:

- "z/OS warehouse agent support for Trillium user-defined program steps" on page 46

Copying data between DB2 for z/OS tables using the LOAD utility

When you want to copy a table by unloading it into a flat file and then loading the flat file to a different table, you normally must unload the data, edit the load control statements that unload produces, then load the data. Using the z/OS warehouse agent, you can specify that you want to reload data to a different table without stopping between steps and manually edit the control statements.

The following procedure will work for any DB2 for z/OS source and target tables on the same or different DB2 subsystems. You can specify the table name with a schema (:TABLE:DBVW.INVENTORY) or without a schema (:TABLE:INVENTORY) to use a default schema.

Procedure:

To copy data between DB2 for z/OS tables using the LOAD utility:

1. Create a step that unloads a file using the UNLOAD utility or the REORG TABLESPACE utility. Both of these utilities produce two output data sets, one with the table data and one with the utility control statement that can be added to the LOAD utility.

   This is an example of the DSNUTILS parameters you might use for the Reorg Unload step:

   ```
   UTILITY_ID REORGULX
   RESTART NO
   UTSTMT REORG TABLESPACE DBVW.USAINENT UNLOAD EXTERNAL
   UTILITY_NAME REORG TABLESPACE
   RECDSN DBVW.DSNURELD.RECDSN
   RECDEVT SYSDA
   RECSPACE 50
   PNCHDSN DBVW.DSNURELD.PNCHDSN
   PNCHDEVT SYSDA
   PNCHSPACE 3
   ```

2. Use the DB2 for z/OS Utility interface to create a load step. The DSNUTILS utility statement parameter specifies a utility control statement. The warehouse utility interface allows you to include a file name in the utility statement field. You can specify the file that contains the valid control statement using the keyword :FILE:, and the name of the table that you want to load using the keyword :TABLE:.
3. To use the LOAD utility to work with the output from the previous example, apply the following parameter values in the LOAD properties:

```
UTILITY_ID LOADREORG
RESTART NO
UTSTMT :FILE:DBVW.DSNURELD.PNCHDSN:TABLE:[DBVW].INVENTORY
UTILITY_NAME LOAD
RECDSN DBVW.DSNURELD.RECDSN
RECDEVT SYSDA
```

4. In the UTSTMT field, type either a load statement or the name of the file that was produced from the REORG utility with the UNLOAD EXTERNAL option. The previous example will work for any DB2 for z/OS source table or target table, whether these tables are on the same or different DB2 subsystems. The control statement flat file can be either HFS or native MVS files.

.

Related concepts:

- "z/OS transformers" on page 84

Related tasks:

- "Running DB2 for z/OS utilities" on page 39

Starting the agent daemon as a z/OS started task

Setting up the agent daemon as a started task allows you to start the daemon from the z/OS console.

Prerequisites:

- You must define a user ID to associate with the procedure. The user ID must be able to write to STDOUT and STDERR and to read from STDENV.
- The user ID and its group must have an OMVS segment defined.

Procedure:

To start the agent daemon as a z/OS started task:

1. Create a file which contains the environment variables for the agent daemon. For example, create an environment file which contains the environment variables (without the `export` command. Add `/bin` to the PATH environment variable and add the environment variables _BPX_BATCH_SPAWN and _BPX_SHAREAS as shown. The environment variable file might be named /u/USERID/BPXprofile and contain:

```
STEPLIB=DSN710.SDSNEXIT:DSN710.SDSNLOAD:$STEPLIB
LIBPATH=/usr/lpp/DWC
PATH=/bin:/usr/lpp/DWC
CXA_CONFIG=/usr/lpp/DWC/cxa.ini
VWS_LOGGING=/u/USERID/logs
VWP_LOG=/u/USERID/logs/vwp.log
```

```
DSNAOINI=/u/USERID/dsnaoini
IBM_MIXED_MODE_THRESHOLD=0
_BPX_BATCH_SPAWN=YES
_BPX_SHAREAS=NO
```

2. Create a procedure in a procedure library. You may need to contact your system programmer to determine what procedure library to use. The procedure will use BPXBATCH to start the agent daemon. Point the STDENV DD card to the environment variable file created in step 1. The following is an example of a procedure to start the agent daemon.

```
//*************************************************************
//*  Start the 390 agent daemon in batch
//*************************************************************
//VWDPROC EXEC PGM=BPXBATCH,PARM='pgm /usr/lpp/DWC/vwd',
//             REGION=8M
//STDOUT   DD  PATH='/tmp/stdout',
//             PATHOPTS=(OWRONLY,OCREAT),
//             PATHMODE=(SIRWXU,SIRWXG,SIRWXO)
//STDERR   DD  PATH='/tmp/stderr',
//             PATHOPTS=(OWRONLY,OCREAT),
//             PATHMODE=(SIRWXU,SIRWXG,SIRWXO)
//STDENV   DD  PATH='/u/USERID/BPXprofile',
//             PATHOPTS=ORDONLY
```

To start the agent daemon from a console, enter `S vwdproc`, where vwdproc is the name of the procedure you created.

To stop the agent daemon from a console, enter `C vwdproc`, where vwdproc is the name of the procedure you created.

Related concepts:

- "Sample contents of DB2 tables and flat files for z/OS and OS/390" on page 36
- "Overview of the z/OS warehouse agent" on page 25

Related tasks:

- "Installing the z/OS warehouse agent" on page 26
- "Setting up Java stored procedures on your DB2 subsystem for the warehouse transformers" on page 85
- "Running DB2 for z/OS utilities" on page 39
- "Running multiple warehouse agent daemons on one z/OS subsystem" on page 42

Running multiple warehouse agent daemons on one z/OS subsystem

The default well-known port that the agent daemon listens to is 11001. There can only be one daemon listening at port 11001, and its environment can only point to one DSNAOINI file. A DSNAOINI file can only point to one DB2

subsystem. That means that a particular z/OS agent can only make a LOCAL connection to one DB2 subsystem per instance of z/OS. However, the agent can connect to many DB2 subsystems via the use of data sharing or through DRDA. You can also have multiple daemons running on one z/OS system, each locally connected to a DB2 subsystem.

Procedure:

To run multiple warehouse agent daemons on one z/OS system:

1. Add a new entry to /etc/services or TCPIP.ETC.SERVICES on z/OS. For example:

   ```
   vwda             11014/tcp      <--- use a new port number
   ```

2. In the .profile that is used to execute this new daemon, add:

   ```
   export vwdPortName=vwda    <---matches the name in services file
   ```

3. On the warehouse server, change the entry in `c:\winnt\system32\drivers\etc\services` for Windows or `/etc/services` on UNIX to match the new port number:

   ```
   vwd               11014/tcp
   ```

Related concepts:

- "Overview of the z/OS warehouse agent" on page 25

Related tasks:

- "Installing the z/OS warehouse agent" on page 26
- "Starting the z/OS warehouse agent daemon" on page 31
- "Setting up Java stored procedures on your DB2 subsystem for the warehouse transformers" on page 85
- "Starting the agent daemon as a z/OS started task" on page 41

Using the z/OS warehouse agent to automate DataPropagator replication apply steps

You can use the z/OS warehouse agent to automate your DataPropagator replication apply steps. Replication requires a source database, a control database, and a target database (these can be different databases or the same). Replication control tables must be installed on these databases. A capture job reads the DB2 log to determine which of the rows in the source database are added, updated, or deleted. The job then writes the changes to a change-data table (also known as a CD table). An apply job is then run to apply the changes to a target database. The DB2 Warehouse Manager package can automate the apply job by creating a replication step. Use the Data Warehouse Center to define a step for the apply job to run and when to run it. The warehouse agent runs the Apply program when the warehouse step is run.

In order to use DataPropagator Version 8, you must update the STEPLIB and PATH environment variable to point to the replication library and directory of the executables. DataPropagator Version 8 does not require a JCL template (a JCL template was required for DataPropagator Version 7 or earlier).

When using DataPropagator Version 7 (or earlier), you must use a JCL template (included in the DB2 Warehouse Manager package) for replication support. If you plan to use the z/OS warehouse agent to run the Apply program, you need to change the account and data set information in this template for your OS/390 or z/OS system.

Procedure:

To use the z/OS warehouse agent to automate DataPropagator Version 7 (or earlier) steps:

1. Log on with an ID that has authority to copy and update files in the `/usr/lpp/DWC/` directory.
2. Find the `apply.jcl` file and copy it as *systemname*`.apply.jcl`, where systemname is the name of the OS/390 MVS system.

   For example, on STLMVS1, create a copy of the file named `STLMVS1.apply.jcl`.
3. Use a text editor to customize the JCL to meet your requirements. Change the account information to match the standard account information, and change the data set for STEPLIB DD and MSGS DD for your OS/390 MVS system.
4. If necessary, change the program name on the EXEC card.

   Do not change any parameters that are contained in brackets, such as `[USERID]` and `[APPLY_PARMS]`. The brackets are the hexadecimal characters X'AD' and X'BD', respectively. If your TSO terminal type is not set to 3278A in SPF Option 0, these values might display as special characters rather than as brackets. This is not a problem if you do not change the XAD or the X BD, or any of the data that is between the characters.
5. Update the environment variable VWS_TEMPLATES to point to the directory of the copied template file.

   This is the JCL template that is included with the Data Warehouse Center:

```
//[USERID]A JOB ,MSGCLASS=H,MSGLEVEL=(1,1),
// REGION=2M,TIME=1440,NOTIFY=&SYSUID
//* DON'T CHANGE THE FIRST LINE OF THIS TEMPLATE.
//* THE REMAINING JCL SHOULD BE MODIFIED FOR YOUR SITE.
//*********************************************
//* RUN APPLY/MVS ON OS/390 DB2 6.1 *
//*********************************************
//ASNARUN EXEC PGM=ASNAPV66,REGION=10M,
// [APPLY_PARMS]
//STEPLIB DD DISP=SHR,DSN=DPROPR.V6R1M0.SASNLINK
```

```
// DD DISP=SHR,DSN=DSN610.SDSNLOAD
//MSGS DD DSN=DPROPR.V2R1M0A.MSGS,DISP=SHR
//ASNASPL DD DSN=&ASNASPL
,DISP=(NEW,DELETE,DELETE),
// UNIT=SYSDA,SPACE=(CYL,(10,1)),
// DCB=(RECFM=VB,BLKSIZE=6404)
//SYSTERM DD SYSOUT=*
//SYSTSPRT DD SYSOUT=*
//SYSPRINT DD SYSOUT=*
//
```

Related concepts:

- "Overview of the z/OS warehouse agent" on page 25

Related tasks:

- "Installing the z/OS warehouse agent" on page 26
- "Starting warehouse agent logging" on page 45

Starting warehouse agent logging

Many DB2 Warehouse Manager components such as the server, the logger, agents, and some Data Warehouse Center programs write logs to the logging directory, which is specified in the VWS_LOGGING environment variable. These log files are plain text.

Procedure:

To start agent logging from the Data Warehouse Center:

1. From the left pane, right-click **Warehouse**, and click **Properties**.
2. On the Trace Level page, change the settings to the trace level that you want.

   The agent trace supports levels 0-4:

   - Level 1 - entry/exit tracing
   - Level 2 - Level 1 plus input/output parameters and return codes
   - Level 3 - Level 1 plus Level 2 plus warning codes plus data
   - Level 4 - Level 1 plus Level 2 plus Level 3 plus more data buffers and staging data

When trace is set higher than level 1, performance will decrease. Turn on tracing only for debugging purposes. The tracing information is stored in the file `AGNTxxx.LOG`. Environment information is stored in the file `AGNTxxx.SET`.

Related concepts:

- "Trace files for the DB2 for iSeries Load utilities" in the *Data Warehouse Center Administration Guide*

- "Warehouse agents" on page 2
- "Component trace data" in the *Data Warehouse Center Administration Guide*
- "Start error trace files" in the *Data Warehouse Center Administration Guide*

Related tasks:

- "Viewing trace files for the DB2 for iSeries Load utilities" in the *Data Warehouse Center Administration Guide*
- "Running a Data Warehouse Center component trace" in the *Data Warehouse Center Administration Guide*
- "Tracing errors created by the Apply program" in the *Data Warehouse Center Administration Guide*

Related reference:

- "Warehouse agent environment structures" on page 97

z/OS warehouse agent support for Trillium user-defined program steps

The z/OS warehouse agent supports the Trillium Batch System User-Defined Program created by the Data Warehouse Center Import Trillium Metadata window. You can start the JCL with the z/OS warehouse agent.

When you create the Trillium Batch System User-Defined Program step using the Import Trillium Metadata window, you must always select **Remote host** for the z/OS warehouse agent, even when the JCL is on the same system as the agent. In addition, all parameters for **Remote host** must be entered.

After you create the Trillium User-Defined Program step, you must change the agent site in the Properties notebook of the Trillium Batch System step to the z/OS warehouse agent site that you want to use.

The Import Trillium Metadata window requires you to enter a JCL file name and an output error file name. If either of these names contains quotation marks or parentheses, you must enclose it in quotation marks.

Related concepts:

- "Sample contents of DB2 tables and flat files for z/OS and OS/390" on page 36
- "Overview of the z/OS warehouse agent" on page 25

Related tasks:

- "Installing the z/OS warehouse agent" on page 26
- "Running DB2 for z/OS utilities" on page 39
- "Running multiple warehouse agent daemons on one z/OS subsystem" on page 42

Installing and configuring the iSeries warehouse agent

The following sections describe how to install and configure the iSeries warehouse agent.

Installing the iSeries warehouse agent

You can install a warehouse agent on an iSeries (AS/400) work station.

Prerequisites:

Before you install an iSeries warehouse agent:

- Remove any previously installed iSeries warehouse agent.
- Define security and privileges for your warehouse.
- Ensure that the following software is installed:
  - IBM DB2 Warehouse Manager Version 8
  - IBM AS/400 V4R5 or later
- Ensure that you have the following authority levels:
  - *ALLOBJ (to use the RSTLICPGM command to install the product)
  - *JOBCTL (to use the STRVWD and ENDVWD commands)
  - *SECOFR (optional for installation but necessary if you want to verify installation). You can also run the DSPSFWRSC command.
  - *USE object authority (to access all commands). Additional authority is required to create a library.

Procedure:

To install the iSeries warehouse agent:

1. Insert the DB2 Warehouse Manager CD-ROM into your iSeries CD-ROM drive.
2. Log on to the iSeries system.
3. At an iSeries command prompt, enter:

   ```
   RSTLICPGM LICPGM(5765F42) DEV(OPT01)
   ```

   where *OPT01* is the CD-ROM drive. The /QIBM/UserData/IWH subdirectory and the /QIBM/ProdData/IWH subdirectory are created by the install process. Traces are written to the /QIBM/UserData/IWH subdirectory.
4. On the Work with Licensed Programs screen, verify that you successfully installed the iSeries warehouse agent. From the iSeries command prompt, enter `GO LICPGM` to use Option 10 (Display Licensed Installed Programs).

Related concepts:

- “Warehouse agents” on page 2

Related tasks:

- “Preparing to install warehouse agents” on page 15
- “Removing iSeries warehouse agents” on page 48

Related reference:

- “Non-U.S. English installations of the iSeries warehouse agent” on page 48
- “Post-installation considerations for iSeries warehouse agents” on page 49

Removing iSeries warehouse agents

Before you install any iSeries agent, you must remove any previously installed iSeries warehouse agent.

Procedure:

To remove iSeries warehouse agent code, at an iSeries command prompt, enter the following command `DLTLICPGM LICPGM(`*`PRODUCT ID`*`)`, where *PRODUCT ID* is any iSeries or AS/400 version previously installed.

This command performs the following actions:

- Removes the SRVTBLE entries
- Removes the /QIBM/ProdData/IWH subdirectory
- Stops the daemon
- Deletes the product

Related concepts:

- “Warehouse agents” on page 2

Related tasks:

- “Installing the iSeries warehouse agent” on page 47
- “Troubleshooting for iSeries warehouse agents” on page 50

Non-U.S. English installations of the iSeries warehouse agent

Only one language version of the iSeries warehouse agent can be present on the iSeries system. The iSeries warehouse agent is language independent. All data is represented internally in Unicode.

All traces and messages created by the STRVWD and ENDVWD commands are displayed in U.S. English on non-U.S. English systems. Messages the agent creates are interpreted on the workstation; messages are displayed in the language that your version of DB2 Warehouse Manager supports.

The user profile that you use to start the daemon must be set to the correct locale for the CCSID that is being used. If the correct locale is not selected, the agent might stop with the error IWH9149, particularly if you are using a DBCS computer.

You can obtain a list of all the available *LOCALE files on your system by entering the command `WRKOBJPDM QSYS *ALL *LOCALE`.

Scroll down until you find the name of the locale that corresponds to the CCSID under which you are running the iSeries warehouse agent. To configure your user profile for the locale that you selected, use the CHGUSRPRF command. On the Change User Profile screen, press F10 to display more options. Scroll down until you find the LOCALE option. Enter the fully qualified locale path using the following syntax:

`/QSYS.LIB/`*`your_locale`*`.LOCALE`

where *your_locale* is the locale that corresponds to the CCSID that you selected for your user profile. For example, the locale path for CCSID 500 is: `/QSYS.LIB/FR_CA.LOCALE`.

Use the following command to install the U.S. English iSeries warehouse agent on your non-US English iSeries system:

`RSTLICPGM LICPGM(5765F42) DEV(`*`OPT01`*`) LNG(2924)`

where *OPT01* is the name of the CD-ROM drive.

Related tasks:

- “Preparing to install warehouse agents” on page 15
- “Installing the iSeries warehouse agent” on page 47
- “Removing iSeries warehouse agents” on page 48

Post-installation considerations for iSeries warehouse agents

All parameters that the Data Warehouse Center passes to warehouse programs are passed as string parameters. If you have applications that require packed decimal parameters, wrap the applications in a CL script that converts the input parameters to a non-string format.

SQL-driven processes are not supported for the flat-file source types. Because there is no equivalent to the ODBC flat-file driver on the iSeries system, you cannot use SQL to retrieve data from, for example, comma-delimited files. Most files on the iSeries system are DB2 tables. However, many applications produce fixed-field or comma-delimited text files. To address the need to load flat file data into a DB2/400 table, see the instructions for working with flat

files on an iSeries system. You can complete the transfer using the warehouse programs that are included with the iSeries warehouse agent.

You cannot use DBCS step names with the iSeries warehouse agent. Also, DB2 for iSeries does not support DBCS object names, including column names, table names, and schema names.

Related tasks:

- "Preparing to install warehouse agents" on page 15
- "Installing the iSeries warehouse agent" on page 47
- "Removing iSeries warehouse agents" on page 48
- "Working with flat files on the iSeries system" on page 56

Related reference:

- "Non-U.S. English installations of the iSeries warehouse agent" on page 48

iSeries warehouse agent security

The QIWH library, where the DB2® Warehouse Manager product is installed, contains a file named FTPCMD. The DB2 Warehouse Manager uses this file for the sample programs VWPFTP and VWPRCPY. The *QSYS user owns this file, and grants *PUBLIC and *CHANGE authority to this file so that any user profile can run the DB2 Warehouse Manager programs. However, using the VWPRCPY program increases your risk for a security exposure. During the operation of VWPRCPY, the program creates a temporary command file that contains the user ID and password of the remote system that you are connecting to.

Related tasks:

- "Installing the iSeries warehouse agent" on page 47
- "Removing iSeries warehouse agents" on page 48

Related reference:

- "Post-installation considerations for iSeries warehouse agents" on page 49

Troubleshooting for iSeries warehouse agents

Use the information contained in this topic to diagnose problems with the iSeries agent.

Procedure:

To troubleshoot the iSeries warehouse agent:

1. Look up the error message.

Every error message contains return codes. The RC2 (Return Code 2) error field is a multipurpose field with different meanings in different contexts. It can store a Data Warehouse Center error code (which starts with *DWC*), a system code, or an application return code, depending on what the primary error (RC 1) was.

- When there is an error during normal agent processing, the warehouse server will always issue an RC 1 = 7356 error. When the server issues a 7356 error, check the RC2 field, which will contain the error number that was reported by the agent.
- When a failure occurs while a warehouse program is running, RC 1 = 8410 and RC 2 = the application return code.
- When there is a communications failure or a failure to run some system function, the Data Warehouse Center error codes are in the 6000 and 9000 ranges. RC 2 usually contains the error number returned by the iSeries system C/C++ language environment. These error numbers (″errno″ values) can help you to pinpoint the cause of an error.

2. Look at the appropriate DB2 Warehouse Manager trace.
3. Check the iSeries job log.

After you locate the relevant error message, use the information to determine the source of the problem and the course of action that is required to fix it. The information about your error message should be the first information that you provide to IBM Software Support when you report a problem.

Related concepts:

- “iSeries warehouse agent security” on page 50

Related tasks:

- “Preparing to install warehouse agents” on page 15
- “Installing the iSeries warehouse agent” on page 47
- “Removing iSeries warehouse agents” on page 48
- “Reading iSeries warehouse agent trace files” on page 54
- “Testing for bidirectional communication between the iSeries warehouse agent and the warehouse server” on page 52
- “Determining the source of connectivity errors for the iSeries warehouse agent” on page 53

Related reference:

- “Post-installation considerations for iSeries warehouse agents” on page 49
- “iSeries agent trace files” on page 54

Testing for bidirectional communication between the iSeries warehouse agent and the warehouse server

While using the iSeries warehouse agent, you might get this error message:

Return Code = 7183 (Method = VWRemoteAgent::Initialize; Secondary Code = 9117) Message: *The warehouse server tried to spawn an agent but did not receive a valid start up acknowledgement from either the agent or the daemon.*

The most common cause of RC7183 is improper configuration of TCP/IP connectivity between the warehouse server and the iSeries warehouse agent. Communication between the server and the agent is bidirectional; the server sends messages to the agent, and the agent sends messages back to the server. Ensure that the warehouse server is connected to the iSeries warehouse agent and vice versa.

Procedure:

To test for bidirectional communication between the iSeries warehouse agent and the warehouse server:

1. Ping the TCP/IP host name. Your host name is specified on the Parameters page for the iSeries warehouse agent site definition, which you can locate from your warehouse server.

   If the ping fails, check that:

   - The iSeries system is registered with your domain name server or that there is an entry for the system in the TCP/IP HOSTS file in the `\winnt\system32\drivers\etc` directory.
   - The iSeries system is running.
   - The network is active.
2. Ping the fully qualified TCP/IP host name for the warehouse server from the iSeries command line. You must use the fully qualified name (hostname.domain), for example `yourmachine.yourcompany.com`. The fully qualified host name is the return address that the server gives the agent.

   If the ping fails, be sure that:

   - The warehouse server is registered with your domain name server or has a host table entry on the iSeries server. Use the CFGTCP command to check that you are using the correct fully qualified TCP/IP host name.
   - The warehouse server is running.
   - The network is active.

If both of the ping attempts were successful, verify that the numeric IP address returned by the ping is actually the IP address of the workstation that you are trying to connect to.

Related tasks:

- "Installing the iSeries warehouse agent" on page 47
- "Working with flat files on the iSeries system" on page 56

Related reference:

- "Non-U.S. English installations of the iSeries warehouse agent" on page 48
- "Post-installation considerations for iSeries warehouse agents" on page 49

Determining the source of connectivity errors for the iSeries warehouse agent

If you get an error message when connecting to a local database from an iSeries database, use this procedure to determine the source of the connection error.

Procedure:

To determine the source of connectivity errors for the iSeries warehouse agent:

1. Add your local database in RDBDIRE on iSeries. Be sure to use the correct port number. To find the correct port number, go to the `C:\WINNT\system32\drivers\etc` directory and open the Services folder. Look for the port that is defined for DB2 instance.
2. Check that your iSeries database is in the Client Configuration Assistant. If not, add it.
3. To test the connection, open the Interactive SQL Assistant and try to connect to your local database using the correct user name and password.

Related concepts:

- "iSeries warehouse agent security" on page 50
- "Warehouse agent connectivity with warehouse sources and targets" on page 21

Related tasks:

- "Installing the iSeries warehouse agent" on page 47
- "Verifying TCP/IP connectivity between the warehouse server and the warehouse agent" on page 23
- "Troubleshooting for iSeries warehouse agents" on page 50
- "Validating the connectivity of an ODBC data source for warehouse agents" on page 17

Related reference:

- "Post-installation considerations for iSeries warehouse agents" on page 49
- "iSeries agent trace files" on page 54

Reading iSeries warehouse agent trace files

Many Data Warehouse Center trace files are stored in the iSeries Integrated File System. To read these trace files, you can either use FTP to move these files to the workstation or use Client Access for iSeries.

Procedure:

To read Data Warehouse Center trace files:

1. Add your system connection in the Client Access Operations Navigator.
2. Expand your system name.
3. Expand the File Systems tree.
4. Right-click on **Integrated File System** and click **Properties**.
5. Select all file extensions, for example `.log, .iwh4msgq, .trc` and click **Add**.
6. Click **OK**.

Related concepts:

- "Warehouse agents" on page 2

Related tasks:

- "Installing the iSeries warehouse agent" on page 47

Related reference:

- "iSeries agent trace files" on page 54

iSeries agent trace files

Many Data Warehouse Center trace files are stored in the iSeries Integrated File System. The following tables give information about the trace files that are produced by the iSeries warehouse agent.

Table 2. Agent trace file

Agent trace	Shows all the information that is passed to the agent from the warehouse server and the command steps that are completed
Location	/QIBM/UserData/IWH directory (root of IFS).
File name	AGNTxxxxx.LOG, where xxxxx is the process ID of the agent instance.

Table 2. Agent trace file (continued)

When to use	Check this trace file first if the error message does not give enough information to solve the problem.
How to use	You can view the trace with any text editor. Scroll down until you find the operation that failed. The parameters sent from the warehouse server for that operation appear below the line that contains the failing operation. This information can help you solve most problems.

Table 3. Message queue trace

Message queue trace	Traces the functions of the message queue process and prints dumps of received message buffers both before and after code page conversion.
Location	/QIBM/UserData/IWH directory (root of IFS).
File name	VWxxxxxx.IWH4MSGQ, where xxxxxx is the process ID of the process that started the message queue process.
When to use	Useful for communications problems that cannot be diagnosed with the error code information alone. Can also be used for Cancel and Get Row Count problems.
How to use	This trace is started automatically when tracing is turned on for either the agent or the daemon. To read this trace, use Microsoft Wordpad or any other unicode enabled editor. This trace might contain non-printable characters and will not format correctly on a basic text editor like Notepad.

An additional trace file can be produced by the message queue process. The msgq_err.log file is a cumulative trace file that records all nonrecoverable message queue errors. This file is useful for tracking down terminations of the message queue process that cannot be recorded in the regular message queue trace file.

Table 4. Daemon trace

Daemon trace	A cumulative trace that records operations that the daemon could not complete successfully.
Location	/QIBM/UserData/IWH directory (root of IFS).
File name	vwd_err.log
When to use	Most useful for diagnosing problems where the agent terminated or abended before it could begin its own trace. Also useful for diagnosing nonrecoverable errors in the daemon.
How to use	This trace is automatically appended when the daemon encounters any error condition during processing.

Table 5. Sample Data Warehouse Center program traces

Sample warehouse program traces	Functional traces of the major events during the execution of a sample iSeries warehouse program.
Location	/QIBM/UserData/IWH
File name	VWxxxxxx.yyyyyyyy, where xxxxxx is the process ID under which the warehouse program was started and yyyyyyyy is the name of the warehouse program. If the warehouse program is started by the agent process, it will run in the same job as the agent process, so it will share the same process ID. The message queue trace, agent trace, and the warehouse program trace will all share the same xxxxxx value.
When to use	Use this trace to determine the cause of a non zero return code.

Related tasks:

- "Reading iSeries warehouse agent trace files" on page 54
- "Working with flat files on the iSeries system" on page 56
- "Troubleshooting for iSeries warehouse agents" on page 50

Working with flat files on the iSeries system

Use the following steps to define a step to load a local flat file into a DB2/400 table. This process will start the warehouse program AS/400 Load with Replace (iSeries CPYFRMIMPF utility).

Restrictions:

You cannot use SQL statements on flat files. SQL is not supported in the current version of the iSeries warehouse agent. You also cannot use sample contents of a flat file using the iSeries warehouse agent. The data must be transferred using the AS/400 Load with Replace warehouse program.

Procedure:

To define a process to load a local file into a local table:

1. Define a flat-file warehouse source for your source file. In the **File name** field, type the fully qualified file name (see iSeries Load with Replace documentation for naming rules).
2. Create a step with the warehouse-supplied AS/400 Load with Replace program.

3. Select your flat-file source, and add the source file to the step.
4. Select your target table from warehouse target and connect with the step.
5. Promote the step to test mode and run it. The target table now contains all the source data from your flat file.

Related concepts:

- "iSeries warehouse agent security" on page 50

Related tasks:

- "Installing the iSeries warehouse agent" on page 47

Related reference:

- "iSeries agent trace files" on page 54

Using the warehouse FTP utility with the iSeries warehouse agent

The iSeries™ warehouse agent provides two FTP utilities:

- Copy File using FTP
- Run FTP Command File

Both of these utilities can be used to copy a file from a remote system to the iSeries warehouse agent system. The Run FTP Command File utility allows you to create an FTP script on the iSeries system and use it on the Parameters page of the step properties notebook.

Related tasks:

- "Installing the iSeries warehouse agent" on page 47
- "Working with flat files on the iSeries system" on page 56
- "Troubleshooting for iSeries warehouse agents" on page 50

Related reference:

- "Post-installation considerations for iSeries warehouse agents" on page 49

Installing the AIX, Linux, and Solaris Operating Environment warehouse agents

The following sections describe how to install and configure the iSeries warehouse agent.

Installing the AIX, Linux, and Solaris Operating Environment warehouse agents

You can install warehouse agents on AIX, Linux, or the Solaris Operating Environment.

Prerequisites:

Before you install the AIX or Solaris Operating Environment warehouse agents, you must remove any earlier AIX or Solaris Operating Environment agents that are located on that workstation.

Procedure:

Use the db2setup utility to install a warehouse agent on AIX, Linux, or the Solaris Operating Environment.

Before you install the warehouse agent, you should create and configure at least one DB2 instance.

To install the AIX, Linux, and Solaris Operating Environment warehouse agents using the db2setup utility:

1. Log in as a user with root authority.
2. Insert and mount the DB2 Warehouse Manager CD-ROM.
3. Change to the directory where the CD-ROM is mounted by entering `cd /cdrom`, where *cdrom* is the mount point of your product CD-ROM.
4. Change to one of the following directories:

AIX	/cdrom/db2/aix
Linux	/cdrom/db2/linux
Solaris	/cdrom/unnamed_cdrom/db2/solaris

5. Enter the `./db2setup` command. After a few moments, the DB2 Setup wizard opens.
6. Click **Install Products** once the DB2 launchpad opens.
7. Ensure that the **DB2 Warehouse Manager** is selected, then click **Next**.
8. Proceed as prompted by the DB2 Setup wizard.
9. In the list of features, select **Warehouse agent**, then click **Next**.
10. Proceed as prompted by the DB2 Setup wizard. Online help is available to guide you through the remaining steps.
11. Log out.

When the installation process completes, your warehouse agent is installed in the following directory:

- On AIX, the `/usr/opt/db2_08_01` directory
- On Linux and the Solaris Operating Environment, the `/opt/IBM/db2/V8.1/` directory

If the installation failed, contact IBM Software Support.

Related concepts:

- "Connectivity products for warehouse sources and targets" on page 22

Related tasks:

- "Configuring the AIX, Linux, and Solaris Operating Environment warehouse agent environment" on page 61
- "Removing DB2 Version 8 warehouse agents" on page 60

Related reference:

- "Database connections for the AIX, Linux, and Solaris Operating Environment warehouse agents" on page 68

Removing a previously installed AIX warehouse agent

Remove previously installed AIX agents before you install AIX warehouse agents.

Procedure:

To remove a previously installed AIX agent:

1. Log on to AIX as *root*.
2. Enter `smit`.
   The System Management window opens.
3. Select **Software Installation and Maintenance**.
   The Software Installation and Maintenance window opens.
4. Select **Maintain Installed Software**.
   The Maintain Installed Software window opens.
5. Select **Remove Software Products**.
   The Remove Software Products window opens.
6. Click the arrow next to the **Software Name** field.
7. Select **db2_07_01.dwa**.
8. Click **OK**.
   The Remove Software Products configuration window opens.
9. Type `No` in the **Preview only?** field.
10. Click **OK**.
    A confirmation window opens.
11. Click **OK** to confirm your selection.

After the AIX agent is removed, the Output window opens. The Results column at the bottom of the Output window displays the status of the removal process. If the removal process fails, contact IBM Software Support.

Related tasks:

- “Removing iSeries warehouse agents” on page 48
- “Removing a previously installed Solaris Operating Environment warehouse agent” on page 60
- “Removing DB2 Version 8 warehouse agents” on page 60

Removing a previously installed Solaris Operating Environment warehouse agent

Remove previously installed Solaris Operating Environment agents before you install Solaris Operating Environment agents.

Procedure:

To remove a previously installed Solaris Operating Environment agent, enter the following command at a command prompt:

```
usr/sbin/pkgrm db2dwa71
```

To confirm the removal of the Solaris Operating Environment agent, enter yes in the confirmation display.

Related tasks:

- “Removing iSeries warehouse agents” on page 48
- “Removing a previously installed AIX warehouse agent” on page 59
- “Removing DB2 Version 8 warehouse agents” on page 60

Removing DB2 version 8 warehouse agents

The following section describes how to remove version 8 warehouse agents.

Removing DB2 Version 8 warehouse agents

The procedure for removing DB2 version 8 warehouse agents varies by platform.

To remove the z/OS warehouse agent, use SMP/E.

Related tasks:

- “Removing iSeries warehouse agents” on page 48
- “Removing DB2 products on UNIX” in the *Installation and Configuration Supplement*

Chapter 4. Configuring the AIX, Linux, and Solaris Operating Environment warehouse agent environments

After you install an AIX, Linux, or Solaris Operating Environment warehouse agent, you must configure the warehouse agent environment and catalog the DB2 nodes and databases. This chapter describes how to configure the warehouse agent environment.

Configuring the AIX, Linux, and Solaris Operating Environment warehouse agent environment

You must set certain environment variables in your IWH.environment file to configure the warehouse agent environment. All the environment variables that you must change are located between the START and END comment lines.

Prerequisites:

The following rules apply to both ODBC and DB2 CLI users:

1. You must set the DB2INSTANCE variable in all situations.
2. For DB2 CLI access, you need to set only the variables in the DB2 environment section.

   To access warehouse sources or targets in DB2 databases, you must configure the variables in the DB2 environment section.
3. You must remove the comment characters (#) for the variables that you set.

The following rules apply to ODBC users:

1. For ODBC access, you must set the IS_ODBC variable to the ODBC install path. This variable must not have comments.

   On AIX, the default path is */usr/opt/db2_08_01/odbc*. On Linux and Solaris Operating Environment, the default path is */opt/IBM/db2/V8.1/odbc*.
2. You must set the ODBC environment variables for any databases that you selected.

When you change settings in the IWH.environment file, you must restart the warehouse agent daemon.

Procedure:

To configure the warehouse agent environment:

1. Set the connection type:
   a. Log on as `root`.
   b. Change to the bin subdirectory in the install directory by entering one of the following commands:
      - On AIX, enter `cd /usr/opt/db2_08_01/bin`
      - On Linux and the Solaris Operating Environment, enter `cd /opt/IBM/db2/V8.1/bin`
   c. Establish the correct soft link to IWH2AGNT. To do this, run the IWH.agent.db.interface command. The syntax for this command is:

      ```
      IWH.agent.db.interface [ odbc | db2cli ]
      ```

      The warehouse agent executable file has two versions:
      - The IWH2AGNT.db2cli file is the DB2 CLI link. Use this link to access DB2 family and heterogeneous data through DB2 Relational Connect (Oracle, Sybase, Microsoft SQL Server, Informix, VSAM, and IMS).

        This is the default link.
      - The IWH2AGNT.ivodbc file is the ODBC link. Use this link to access DB2 family, Oracle, Sybase, and Informix databases, and data files.
2. To access DB2 or DB2 Relational Connect source or target data, set the following variables in your IWH.environment file. On AIX, the file is located in the */usr/opt/db2_08_01/bin* directory. On Linux and the Solaris Operating Environment, the file is located in the */opt/IBM/db2/V8.1/bin* directory. You need to set these variables whether you use a double-byte character set (DBCS) or a single-byte character set (SBCS).

DB2INSTANCE	The name of the DB2 instance on which you plan to run the warehouse agent. If you have multiple DB2 instances, you must choose one instance to be the master DB instance.
LANG	Your local locale.
LC_ALL	Your local locale. The LANG and LC_ALL environment variables direct the warehouse agent to set the locale to your local locale.
DB2CODEPAGE	The code page of the client database. The DB2CODEPAGE variable overrides the system code page. This variable is used to communicate with servers that cannot interpret the default code page from the

client. The DB2CODEPAGE variable sets the default code page to a code page that the server supports.

3. To access ODBC data using the Data Warehouse Center ODBC driver, set the following variables in your IWH.environment file:
   - IS_ODBC = the base directory where you installed the Data Warehouse Center ODBC driver. On AIX, the default path is */usr/opt/db2_08_01/odbc*. On Linux and Solaris Operating Environment, the default path is */opt/IBM/db2/V8.1/odbc*
   - To access Oracle data, set the following environment variables:
     – ORACLE_HOME
     – TNS_ADMIN
     – ORACLE_SID
   - To access Sybase data, set the following environment variables:
     – SYBASE
     – DSQUERY
   - To access Informix data, set the following environment variables:
     – INFORMIXDIR
     – INFORMIXSERVER
4. To access Essbase or DB2 OLAP Server™ data, set the ARBORPATH environment variable in your IWH.environment file. If you will be using the Data Warehouse OLAP programs, update the PATH environment variable so that it includes the executables needed by Essbase or DB2 OLAP.
5. Optional: Customize your environment further by setting the following variables in your IWH.environment file:
   - VWS_LOGGING = the directory where the warehouse agent trace files, program output logs, and environment output files are stored.

     The default directory is /var/IWH.
   - VWSPATH = the warehouse agent base install directory.
   - VWA_LOCALE_TABLE = locale conversion table.

     If you are using a national language version of the Solaris Operating Environment warehouse agent, you might need to choose a user-defined conversion table that the agent can use for string conversions. To do this, add the environment variable VWA_LOCALE_TABLE to the IWH.environment file and set its value to the code set that you want to use (for example, 8859-1 for en_US).
6. If you are using a shared /etc/services file in the NFS or AFS environment, add the following services to the shared /etc/services file.

   If these port numbers are already in use in your operating environment, you can change them in the /etc/services file on your system. However,

you must also change the same port numbers on the workstation where the Data Warehouse Center administrative client and warehouse server are installed, and on all the associated warehouse agent sites.

You must add these services because the install process updates only your local /etc/services file.

Port Name	Port Number
vwd	11001/tcp

7. For these changes to take effect, stop and restart the warehouse agent daemon:
   a. Determine the daemon process ID by entering:
      ```
      ps -e | grep vwdaemon
      ```
   b. Stop the daemon by entering:
      ```
      kill pid
      ```
      where `pid` is the daemon process ID.

   After you stop the daemon, it will restart automatically. If you installed the AIX warehouse server, you must also restart the server daemons after changing the IWH.environment file using the command `db2vwsvr`.

Related tasks:

- "Cataloging the DB2 nodes and databases after installing the AIX, Linux, or Solaris Operating Environment warehouse agent" on page 67
- "Installing the AIX, Linux, and Solaris Operating Environment warehouse agents" on page 57
- "Using the Data Warehouse Center external trigger" on page 69
- "Starting and stopping the warehouse server and logger (AIX)" in the *Data Warehouse Center Administration Guide*

Related reference:

- "Database connections for the AIX, Linux, and Solaris Operating Environment warehouse agents" on page 68
- "Sample IWH.environment file" on page 64

Sample IWH.environment file

The IWH.environment file contains environment variables that you must set when you configure your warehouse agent environment.

```
#!/bin/bsh
##############################################################################
#
# Licensed Materials - Property of IBM
```

```
#
# 5648-B90
# (C) COPYRIGHT International Business Machines Corp. 1993, 1999
#
# 5648-B91
# (C) COPYRIGHT International Business Machines Corp. 1993, 1999
#
# 5648-B95
# (C) COPYRIGHT International Business Machines Corp. 1993, 1999
#
# 5648-B97
# (C) COPYRIGHT International Business Machines Corp. 1993, 1999
#
# 5648-B99
# (C) COPYRIGHT International Business Machines Corp. 1993, 1999
#
# All Rights Reserved
# US Government Users Restricted Rights - Use, duplication or
# disclosure restricted by GSA ADP Schedule Contract with IBM Corp.
#
################################################################################
#
# NAME: IWH.environment
#
################################################################################

DB2DIR="/usr/opt/db2_08_01"

# Visual Warehouse Daemon environment setup script.
# Please Update this script to setup environment variables for other databases.

# NOTE :
# After adding new setup, the daemon process named : vwdaemon should
#  be terminated using command: "kill <PID>" for changes to
#  take effect.
# DO NOT EXECUTE vwdaemon. It will be re-executed immediately by init process.

# !!!!!!!!!!!!!!!!!!!!!!!!!! ATTENTION !!!!!!!!!!!!!!!!!!!!!!!!!!
# !!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
# !     START OF MUST SET ENVIRONMENT VARIABLES SECTION        !
# !!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
# !!!!!!!!!!!!!!!!!!!!!!!!!! ATTENTION !!!!!!!!!!!!!!!!!!!!!!!!!!

# DB2 environment section. Please set DB2 instance id
DB2INSTANCE=db2inst1 # make sure this is correct!

# These are the defaults.  To change any of them, uncomment it and make
# the change.
#LANG=en_US
#VWS_LOGGING=/var/IWH
#VWA_LOCALE_TABLE=8859-1
#LC_ALL=en_US
#DB2CODEPAGE=819
#VWSPATH=${DB2DIR?}
#VWS_TEMPLATES=${VWSPATH?}/templates
```

```
# If using InterSolv odbc drivers, set the IS_ODBC
# environment variable to the path for Intersolv odbc
# directory.
#
# IS_ODBC=<..../intersolv/..../odbc> # Should have path for Intersolv odbc dir

# For essbase udp, set the following path :
#
# ARBORPATH=<..../essbase>
# export ARBORPATH

# If using Oracle InterSolv driver, set the following paths:
#
# ORACLE_HOME=<.... your Oracle client HOME dir ....>
# TNS_ADMIN=<.... your TNS_ADMIN dir ....>
# ORACLE_SID=<.... your ORACLE_SID ....>
# export ORACLE_HOME TNS_ADMIN ORACLE_SID

# If using Sybase InterSolv driver, set the following paths:
#
# SYBASE=<.... your Sybase client dir ....>
# DSQUERY=<.... your DSQUERY value ....>
# export SYBASE DSQUERY

# If using Informix InterSolv driver, set the following paths:
#
# INFORMIXDIR=<.... your Informix client dir ....>
# INFORMIXSERVER=<.... your Informix Server ....>
# export INFORMIXDIR INFORMIXSERVER

# !!!!!!!!!!!!!!!!!!!!!!!!!! ATTENTION !!!!!!!!!!!!!!!!!!!!!!!!!!
# !!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
# !       END OF MUST SET ENVIRONMENT VARIABLES SECTION        !
# !!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
# !!!!!!!!!!!!!!!!!!!!!!!!!! ATTENTION !!!!!!!!!!!!!!!!!!!!!!!!!!

# Do not modify the following lines.

# For the following, uncomment out the appropriate section(s)

# if [ "${INFORMIXDIR}" != "" ]; then
#     PATH=${INFORMIXDIR?}/bin:${PATH?}
#     LIBPATH=${LIBPATH:-""}:${INFORMIXDIR?}/lib
# else
#     INFORMIXDIR=${IS_ODBC?}
#     export INFORMIXDIR
# fi
#
# if [ "${SYBASE}" != "" ]; then
#     PATH=${SYBASE?}/bin:${PATH}
#     LIBPATH=${LIBPATH:-""}:${SYBASE?}/lib
# fi
#
# if [ "${ORACLE_HOME}" != "" ]; then
```

```
#     PATH=${ORACLE_HOME?}/bin:${PATH?}
#     LIBPATH=${LIBPATH:-""}:${ORACLE_HOME?}/lib
# fi
#
# if [ "${ARBORPATH}" != "" ]; then
#     PATH=$PATH:${ARBORPATH?}/bin
#     LIBPATH=${LIBPATH:-""}:${ARBORPATH?}/api/lib
# fi
#
# if [ "${IS_ODBC}" != "" ]; then
#     # for intersolv odbc
#     LIBPATH=${LIBPATH:-""}:${IS_ODBC?}/lib
#     IV_GLS_LCDIR=${IS_ODBC?}/gls/lc11
#     IV_GLS_REGISTRY=${IS_ODBC?}/gls/cm3/registry
#     export IV_GLS_LCDIR IV_GLS_REGISTRY
# fi
```

Cataloging the DB2 nodes and databases after installing the AIX, Linux, or Solaris Operating Environment warehouse agent

After you install the AIX, Linux, or Solaris Operating Environment warehouse agent and configure the warehouse agent environment, you must configure the DB2 nodes and databases.

Procedure:

To catalog the DB2 nodes and databases after installing the AIX, Linux, or Solaris Operating Environment warehouse agent:

1. Log on to the DB2 instance ID.
2. Catalog the DB2 nodes and databases that you plan to access.

   If you have multiple DB2 instances on the same workstation, note the value of the DB2INSTANCE variable in the IWH.environment file. Any databases that are not created on this instance must be cataloged as remote.
3. For ODBC users: When you define the DB2 data source in the appropriate .odbc.ini file, ensure that the DSN is the true name of the DB2 database name or database alias name. The .odbc.ini file must reside in the home directory of the user ID that is used on the agent site.

   Tip: If you are using Data Warehouse Center ODBC driver, see the sample .odbc.ini file. On AIX, this file is located in the `/usr/opt/db2_08_01/odbc` directory. On Linux and the Solaris Operating Environment, this file is located in the `/opt/IBM/db2/V8.1/odbc` directory. The name of the file is `odbc.ini.intersolv`.

Related tasks:

- "Configuring the AIX, Linux, and Solaris Operating Environment warehouse agent environment" on page 61
- "Using the Data Warehouse Center external trigger" on page 69

Related reference:

- "Database connections for the AIX, Linux, and Solaris Operating Environment warehouse agents" on page 68

Database connections for the AIX, Linux, and Solaris Operating Environment warehouse agents

When you install the AIX, Linux, or Solaris Operating Environment warehouse agent, two ODBC connection types are available through the agent:

DB2 CLI

This connection type accesses the following sources:

- DB2 family (through ODBC)
- Heterogeneous data sources through federated support provided by DB2 Relational Connect, including:
  - Oracle
  - Sybase
  - Microsoft SQL Server
  - Informix
  - VSAM
  - IMS™

If you use the DB2 CLI connection type, you must have DB2 Relational Connect for heterogeneous data access. VSAM and IMS source access also requires DB2 Classic Connect.

Data Warehouse ODBC

This connection type accesses the following sources:

- DB2 family (through ODBC)
- Oracle
- Sybase
- Microsoft SQL Server
- Informix
- Flat file

Both connection types support the following warehouse target databases:

- DB2 Universal Database Enterprise Server Edition

- DB2 Universal Database Enterprise Edition
- DB2 Universal Database Enterprise - Extended Edition
- DB2 Universal Database for iSeries
- DB2 for z/OS
- DB2 Relational Connect
- Oracle (through DB2 Relational Connect)
- Sybase (through DB2 Relational Connect)
- Informix (through DB2 Relational Connect)

Related tasks:

- "Configuring the AIX, Linux, and Solaris Operating Environment warehouse agent environment" on page 61
- "Cataloging the DB2 nodes and databases after installing the AIX, Linux, or Solaris Operating Environment warehouse agent" on page 67
- "Installing the AIX, Linux, and Solaris Operating Environment warehouse agents" on page 57
- "Using the Data Warehouse Center external trigger" on page 69

Using the Data Warehouse Center external trigger

If you will be using the Data Warehouse Center external trigger to externally start steps, or to promote or demote steps to or from development, test, or production status, you must reserve port number 11004 so that the Data Warehouse external trigger can communicate with the Data Warehouse Center server.

Related tasks:

- "Configuring the AIX, Linux, and Solaris Operating Environment warehouse agent environment" on page 61
- "Cataloging the DB2 nodes and databases after installing the AIX, Linux, or Solaris Operating Environment warehouse agent" on page 67
- "Installing the AIX, Linux, and Solaris Operating Environment warehouse agents" on page 57

Related reference:

- "Database connections for the AIX, Linux, and Solaris Operating Environment warehouse agents" on page 68

Chapter 5. Preparing for and installing warehouse transformers

This chapter describes how to install and enable warehouse transformers.

Preparing your environment for warehouse transformers

Warehouse transformers are stored procedures and user-defined functions that you can use to transform data in a warehouse step. To use warehouse transformers, you must install them at the agent site and enable them for use at the target warehouse. Prepare your environment before you enable warehouse transformers.

Procedure:

To prepare your environment for warehouse transformers:

1. Install a warehouse agent.
2. Install the JDK for your operating system. If your target database is DB2 Universal Database Version 7.2 or earlier, the Java Developer's Kit (JDK) Version 1.1.8 or later must be installed on the database where you plan to use the warehouse transformers. If your target database is DB2 Universal Database Version 8, the Java Developer's Kit (JDK) Version 1.3.0 or later must be installed on the database where you plan to use the warehouse transformers. The DB2 Application Development Client includes the JDK for AIX, the Solaris Operating Environment, Linux, and Windows.
3. Update the environment variables.
4. Update the database manager configuration for the target DB2 instance.
5. Update the database configuration for the target database.

Related concepts:

- "JDK installation for warehouse transformers" on page 72
- "z/OS transformers" on page 84
- "Warehouse transformers" on page 3

Related tasks:

- "Updating the environment variables on Windows for warehouse transformers" on page 73
- "Updating the environment variables on AIX for warehouse transformers" on page 74

- "Updating the database manager configuration before installing warehouse transformers" on page 79
- "Updating the database configuration for the target database" on page 80
- "Installing warehouse transformers" on page 81
- "Enabling warehouse transformers" on page 82
- "Setting up warehouse transformers on DB2 for z/OS" on page 86
- "Updating the environment variables on the Solaris Operating Environment for warehouse transformers" on page 76
- "Updating the environment variables on Linux for warehouse transformers" on page 77

Related reference:

- "National language support for warehouse transformers" on page 91

JDK installation for warehouse transformers

Ensure that you install the JDK for your system before you install warehouse transformers (for the Linux transformer, you must use the IBM® JDK for Linux).

The DB2® Application Development Client includes the JDK for AIX, Linux, and Windows:

- When you install the DB2 Application Development Client on AIX® and Linux, the JDK is installed if an earlier version of the JDK is not detected.
- When you install the DB2 Application Development Client on Windows, the JDK is always installed.

To obtain and install the JDK for the Solaris Operating Environment, go to http://www.sun.com/java/products/platform.html and follow the instructions provided.

Related concepts:

- "z/OS transformers" on page 84
- "Warehouse transformers" on page 3

Related tasks:

- "Updating the environment variables on Windows for warehouse transformers" on page 73
- "Updating the environment variables on AIX for warehouse transformers" on page 74
- "Updating the database manager configuration before installing warehouse transformers" on page 79

- "Updating the database configuration for the target database" on page 80
- "Installing warehouse transformers" on page 81
- "Enabling warehouse transformers" on page 82
- "Preparing your environment for warehouse transformers" on page 71
- "Setting up warehouse transformers on DB2 for z/OS" on page 86
- "Updating the environment variables on the Solaris Operating Environment for warehouse transformers" on page 76

Related reference:

- "National language support for warehouse transformers" on page 91

Updating environment variables for warehouse transformers

The following sections describe how to update environment variables for warehouse transformers.

Updating the environment variables on Windows for warehouse transformers

You must update the environment variables on Windows before you enable warehouse transformers.

Procedure:

To update the environment variables on Windows for warehouse transformers:

1. Open the Windows Control Panel.
2. Double-click **System**.
3. Click the **Environment** tab to view the variables.
4. Click the **Path** system variable and add the path of the JDK bin subdirectory to the end of the path. The default path is `c:\sqllib\java\jdk`.
5. Click **Set**.
6. Click the **CLASSPATH** system variable. Add the path where the transformers are installed, and the path of the directory that contains the DB2 Java support classes. Separate the path names with a semicolon. If the CLASSPATH variable does not exist, you must define it.

   For example, you might add `C:\SQLLIB\FUNCTION;C:\SQLLIB\JAVA\DB2JAVA.ZIP;` to your CLASSPATH variable.
7. Click **Set**.
8. Click **OK** to close the window.

Related concepts:

- "JDK installation for warehouse transformers" on page 72
- "z/OS transformers" on page 84
- "Warehouse transformers" on page 3

Related tasks:

- "Updating the environment variables on AIX for warehouse transformers" on page 74
- "Updating the database manager configuration before installing warehouse transformers" on page 79
- "Updating the database configuration for the target database" on page 80
- "Installing warehouse transformers" on page 81
- "Enabling warehouse transformers" on page 82
- "Preparing your environment for warehouse transformers" on page 71
- "Setting up warehouse transformers on DB2 for z/OS" on page 86
- "Updating the environment variables on the Solaris Operating Environment for warehouse transformers" on page 76
- "Updating the environment variables on Linux for warehouse transformers" on page 77

Related reference:

- "National language support for warehouse transformers" on page 91

Updating the environment variables on AIX for warehouse transformers

Update the environment variables before you install warehouse transformers.

Procedure:

To update the environment variables on AIX for warehouse transformers:

1. Ensure that the current directory is in the CLASSPATH variable by entering the following command:

   ```
   $ echo $CLASSPATH
   ```

   For example:

   ```
   .:/INSTHOME/sqllib/java/db2java.zip
   ```

   where */INSTHOME* is the home directory of the instance owner.
2. Enter the following commands to change the CLASSPATH variable:

   ```
   CLASSPATH=.:$CLASSPATH
   export CLASSPATH
   ```

3. Ensure that *INSTHOME/sqllib/function* is in the CLASSPATH variable by entering the following command:

   ```
   $ echo $CLASSPATH
   ```

4. Enter the following commands to change the CLASSPATH variable:

   ```
   CLASSPATH=INSTHOME/sqllib/function:$CLASSPATH
   export CLASSPATH
   ```

5. Make sure that the PATH variable includes /sqllib/bin. For example:

   ```
   /INSTHOME/sqllib/bin
   ```

   where */INSTHOME* is the home directory of the instance owner.

6. Make sure that the LIBPATH variable includes /sqllib/lib. For example:

   ```
   /INSTHOME/sqllib/lib
   ```

   where */INSTHOME* is the home directory of the instance owner.

Related concepts:

- "JDK installation for warehouse transformers" on page 72
- "z/OS transformers" on page 84
- "Warehouse transformers" on page 3

Related tasks:

- "Updating the environment variables on Windows for warehouse transformers" on page 73
- "Updating the database manager configuration before installing warehouse transformers" on page 79
- "Updating the database configuration for the target database" on page 80
- "Installing warehouse transformers" on page 81
- "Enabling warehouse transformers" on page 82
- "Preparing your environment for warehouse transformers" on page 71
- "Setting up warehouse transformers on DB2 for z/OS" on page 86
- "Updating the environment variables on the Solaris Operating Environment for warehouse transformers" on page 76
- "Updating the environment variables on Linux for warehouse transformers" on page 77

Related reference:

- "National language support for warehouse transformers" on page 91

Updating the environment variables on the Solaris Operating Environment for warehouse transformers

Update the environment variables before you install warehouse transformers.

Procedure:

To update the environment variables on the Solaris Operating Environment for warehouse transformers:

1. Update the db2profile file, located in the sqllib subdirectory (for example, /home/db2inst1/sqllib), to include *jdk_path*/lib/sparc/native_threads in the LD_LIBRARY_PATH variable, where *jdk_path* is the subdirectory where the JDK is installed.
2. Make sure that the current directory is in the CLASSPATH variable by entering the following command:

   ```
   $ echo $CLASSPATH
   ```

   For example:

   ```
   .:/INSTHOME/sqllib/java/db2java.zip
   ```

   where */INSTHOME* is the home directory of the instance owner.
3. Enter the following commands to change the CLASSPATH variable:

   ```
   CLASSPATH=.:$CLASSPATH
   export CLASSPATH
   ```
4. Make sure that the PATH variable includes sqllib/bin. For example:

   ```
   /INSTHOME/sqllib/bin
   ```

   where */INSTHOME* is the home directory of the instance owner.
5. Ensure that *INSTHOME*/*sqllib*/*function* is in the CLASSPATH variable by entering the following command:

   ```
   $ echo $CLASSPATH
   ```
6. Enter the following commands to change the CLASSPATH variable:

   ```
   CLASSPATH=INSTHOME/sqllib/function:$CLASSPATH
   export CLASSPATH
   ```

Related concepts:

- "JDK installation for warehouse transformers" on page 72
- "z/OS transformers" on page 84
- "Warehouse transformers" on page 3

Related tasks:

- "Updating the environment variables on Windows for warehouse transformers" on page 73

- “Updating the environment variables on AIX for warehouse transformers” on page 74
- “Installing warehouse transformers” on page 81
- “Enabling warehouse transformers” on page 82
- “Preparing your environment for warehouse transformers” on page 71
- “Setting up warehouse transformers on DB2 for z/OS” on page 86
- “Updating the environment variables on Linux for warehouse transformers” on page 77

Related reference:

- “National language support for warehouse transformers” on page 91

Updating the environment variables on Linux for warehouse transformers

Update the environment variables before you install warehouse transformers.

Prerequisites:

To run warehouse transformers (Java stored procedures or user-defined functions) on DB2 Universal Database for Linux, the Linux runtime linker must be able to access certain Java shared libraries. You can either add the location of the Java shared libraries to `/etc/ld.so.conf`, or create symbolic links to the libraries in the `/usr/lib` directory. If you decide to create symbolic links for the libraries in `/usr/lib`, the list of libraries to link to is different for different versions of the IBM Developer Kit for Java.

For IBM Developer Kit for Java Version 1.1.8 (which you use for warehouse transformers installed on a DB2 Universal Database Version 7.2 or earlier database), you must have symbolic links pointing to: `libjava.so libjitc.so libmath.so libzip.so`. For Version 1.2.2 or 1.3 of the IBM Developer Kit for Java, you must have symbolic links pointing to: `libjava.so libjvm.so libhpi.so`.

Once you either add the location of the Java shared libraries to `/etc/ld.so.conf` or create the symbolic links, you must refresh the runtime linker cache by running the following command as root: `bash# ldconfig`. To run with version 1.2.2 or 1.3 of the IBM Developer Kit for Java, you also need to enter the following commands: `bash$ db2set DB2_USE_JDK12=true bash$ db2stop bash$ db2start`. Please note that you should only use IBM JDK 1.3 for DB2 Universal Database Version 8 and higher.

Procedure:

To update the environment variables on Linux for warehouse transformers:

1. Make sure that the current directory is in the CLASSPATH variable by entering the following command:

   ```
   $ echo $CLASSPATH
   ```

   For example:

   ```
   .:/INSTHOME/sqllib/java/db2java.zip
   ```

   where /INSTHOME is the home directory of the instance owner.
2. Enter the following commands to change the CLASSPATH variable:

   ```
   CLASSPATH=.:$CLASSPATH
   export CLASSPATH
   ```
3. Ensure that INSTHOME/sqllib/function is in the CLASSPATH variable by entering the following command:

   ```
   $ echo $CLASSPATH
   ```
4. Enter the following commands to change the CLASSPATH variable:

   ```
   CLASSPATH=INSTHOME/sqllib/function:$CLASSPATH
   export CLASSPATH
   ```
5. Make sure that the PATH variable includes sqllib/bin.

   For example:

   ```
   /INSTHOME/sqllib/bin
   ```

   where /INSTHOME is the home directory of the instance owner.
6. Make sure that the library path variable (LD_LIBRARY_PATH or LIBPATH as applicable) includes `/sqllib/lib`.

   For example:

   ```
   /INSTHOME/sqllib/lib
   ```

   where /INSTHOME is the home directory of the instance owner.

Related concepts:

- "JDK installation for warehouse transformers" on page 72
- "z/OS transformers" on page 84
- "Warehouse transformers" on page 3

Related tasks:

- "Updating the environment variables on Windows for warehouse transformers" on page 73
- "Updating the environment variables on AIX for warehouse transformers" on page 74
- "Installing warehouse transformers" on page 81
- "Enabling warehouse transformers" on page 82
- "Preparing your environment for warehouse transformers" on page 71

- "Setting up warehouse transformers on DB2 for z/OS" on page 86
- "Updating the environment variables on the Solaris Operating Environment for warehouse transformers" on page 76

Related reference:

- "National language support for warehouse transformers" on page 91

Updating the database manager configuration for warehouse transformers

The following sections describe how to update the database manager configuration for warehouse transformers.

Updating the database manager configuration before installing warehouse transformers

Use the DB2 Command Line Processor to update the database manager configuration for the target DB2 instance before you install warehouse transformers.

Procedure:

To update the database manager configuration before installing warehouse transformers:

- Set the JDK path parameter (JDK_PATH) to the subdirectory where the JDK is installed:

  ```
  UPDATE DATABASE MANAGER CONFIGURATION USING JDK_PATH path
  ```

  where *path* is the subdirectory where the JDK is installed.

  Note: Do not include the bin directory in jdk_path. If you are using a DB2 Universal Database Version 7.2 (or earlier) target that has transformers installed, you should update JDK11_PATH as opposed to JDK_PATH.
- Set the Java Virtual Machine heap size parameter (JAVA_HEAP_SZ) to 4096:

  ```
  UPDATE DATABASE MANAGER CONFIGURATION USING JAVA_HEAP_SZ 4096
  ```

Related concepts:

- "Warehouse transformers" on page 3

Related tasks:

- "Updating the environment variables on Windows for warehouse transformers" on page 73
- "Updating the environment variables on AIX for warehouse transformers" on page 74

- "Updating the database configuration for the target database" on page 80
- "Installing warehouse transformers" on page 81
- "Enabling warehouse transformers" on page 82
- "Preparing your environment for warehouse transformers" on page 71
- "Setting up warehouse transformers on DB2 for z/OS" on page 86

Related reference:

- "National language support for warehouse transformers" on page 91

Updating the database configuration for the target database

Update the database configuration for the target database before you install warehouse transformers.

Procedure:

To update the database configuration for the target database, use the DB2 Command Line Processor to set the default application heap parameter (APPLHEAPSZ) for this database to 1024:

```
UPDATE DATABASE CONFIGURATION FOR warehouse_database_name USING APPLHEAPSZ 1024
```

Related concepts:

- "JDK installation for warehouse transformers" on page 72
- "z/OS transformers" on page 84
- "Warehouse transformers" on page 3

Related tasks:

- "Updating the environment variables on Windows for warehouse transformers" on page 73
- "Updating the environment variables on AIX for warehouse transformers" on page 74
- "Updating the database manager configuration before installing warehouse transformers" on page 79
- "Installing warehouse transformers" on page 81
- "Enabling warehouse transformers" on page 82
- "Preparing your environment for warehouse transformers" on page 71
- "Setting up warehouse transformers on DB2 for z/OS" on page 86
- "Updating the environment variables on the Solaris Operating Environment for warehouse transformers" on page 76
- "Updating the environment variables on Linux for warehouse transformers" on page 77

Related reference:

- "National language support for warehouse transformers" on page 91

Installing and enabling warehouse transformers

The following sections describe how to install and enable warehouse transformers.

Installing warehouse transformers

Warehouse transformers are stored procedures. Before you enable the warehouse transformers, certain authorization levels are needed for the database in which the target warehouse resides.

Prerequisites:

The authorization ID of the statement must hold at least one of the following authorities or privileges:

- SYSADM or DBADM authority
- IMPLICIT_SCHEMA authority on the database, if the implicit or explicit schema name of the procedure does not exist
- CREATEIN privilege on the schema, if the schema name of the procedure refers to an existing schema

To create a fenced stored procedure, no additional authorities or privileges are required.

To create a not-fenced stored procedure, the authorization ID of the statement must also hold one of the following authorities:

- CREATE_NOT_FENCED authority on the database
- SYSADM or DBADM authority

Procedure:

To install warehouse transformers:

1. Insert the DB2 Warehouse Manager CD-ROM into your CD-ROM drive. The auto-run feature automatically starts the DB2 Setup launchpad.
2. Click **Install Products** from the launchpad.
3. Ensure that **DB2 Warehouse Manager** is selected, then click **Next**.
4. Proceed as prompted by the DB2 Setup wizard.
5. In the list of features, select **Warehouse transformers**, then click **Next**.
6. Proceed as prompted by the DB2 Setup wizard. Online help is available to guide you through the remaining steps.

Related concepts:

- “JDK installation for warehouse transformers” on page 72
- “z/OS transformers” on page 84
- “Warehouse transformers” on page 3

Related tasks:

- “Updating the environment variables on Windows for warehouse transformers” on page 73
- “Updating the environment variables on AIX for warehouse transformers” on page 74
- “Updating the database manager configuration before installing warehouse transformers” on page 79
- “Updating the database configuration for the target database” on page 80
- “Enabling warehouse transformers” on page 82
- “Preparing your environment for warehouse transformers” on page 71
- “Setting up warehouse transformers on DB2 for z/OS” on page 86
- “Updating the environment variables on the Solaris Operating Environment for warehouse transformers” on page 76
- “Updating the environment variables on Linux for warehouse transformers” on page 77

Related reference:

- “National language support for warehouse transformers” on page 91

Enabling warehouse transformers

Warehouse transformers are Java stored procedures. A single database can contain several warehouse targets; however, you must enable the warehouse transformers for each warehouse target. If you have multiple warehouse targets in a single database and you change these settings for one warehouse target, the other warehouse targets in that database will be affected.

For example, if you drop the warehouse transformers after enabling them for several warehouse targets, a message is displayed, indicating that this change will affect other warehouse targets defined to the database. The names of the other warehouse targets will be listed in the message. You can choose to cancel the drop procedure or to continue. If you continue, any processes or steps in those warehouse targets that use warehouse transformers will not run, unless you recreate the warehouse transformers.

Prerequisites:

Before you enable the warehouse transformers, you must:

- Install a warehouse agent and warehouse transformers.
- Install the JDK for your operating system.
- Update the environment variables.
- Update the database manager configuration for the target DB2 instance.
- Update the database configuration for the target database.

Procedure:

To enable warehouse transformers:

1. Start the Data Warehouse Center.
2. In the tree view on the left, locate the warehouse target for which you want to enable warehouse transformers, and open the Properties notebook.
3. On the Database page, in the **Target database** area, click either **Create transformers and register as fenced** or **Create transformers and register as unfenced**.
4. In the **Warehouse target** area, click **Enable target for transformers**.
5. Click **OK**.
6. Close the notebook.

If you change the warehouse transformers from fenced to unfenced, or from unfenced to fenced, and there are multiple warehouse targets defined to the database, a message is displayed, indicating that this change will affect other warehouse targets defined to the database. The names of the other warehouse targets will be listed in the message. You can choose to cancel the change or to continue. If you continue, the transformers will be dropped, then recreated as specified.

Related concepts:

- "JDK installation for warehouse transformers" on page 72
- "z/OS transformers" on page 84
- "Warehouse transformers" on page 3

Related tasks:

- "Updating the environment variables on Windows for warehouse transformers" on page 73
- "Updating the environment variables on AIX for warehouse transformers" on page 74
- "Updating the database manager configuration before installing warehouse transformers" on page 79
- "Updating the database configuration for the target database" on page 80
- "Installing warehouse transformers" on page 81

- "Preparing your environment for warehouse transformers" on page 71
- "Setting up warehouse transformers on DB2 for z/OS" on page 86
- "Updating the environment variables on the Solaris Operating Environment for warehouse transformers" on page 76
- "Updating the environment variables on Linux for warehouse transformers" on page 77

Related reference:

- "National language support for warehouse transformers" on page 91

z/OS transformers

The Data Warehouse Center includes transformers, which are Java™ stored procedures that provide some basic data transformations. To run transformers, you must first set up Java stored procedures on your DB2® subsystem.

You can run the following transformers with the z/OS™ agent:

- IWH.CLEAN
- IWH.PERIODTABLE
- IWH.KEYTABLE
- IWH.CHISQUARE
- IWH.CORRELATION
- IWH.STATISTICS
- IWH.INVERTDATA
- IWH.PIVOTDATA
- IWH.REGRESSION
- IWH.ANOVA
- IWH.SUBTOTAL
- IWH.MOVINGAVERAGE

Related concepts:

- "JDK installation for warehouse transformers" on page 72
- "Warehouse transformers" on page 3

Related tasks:

- "Updating the environment variables on Windows for warehouse transformers" on page 73
- "Updating the environment variables on AIX for warehouse transformers" on page 74
- "Updating the database configuration for the target database" on page 80
- "Installing warehouse transformers" on page 81

- "Enabling warehouse transformers" on page 82
- "Preparing your environment for warehouse transformers" on page 71
- "Setting up Java stored procedures on your DB2 subsystem for the warehouse transformers" on page 85
- "Setting up warehouse transformers on DB2 for z/OS" on page 86
- "Reducing the number of characters in the environment variable data set for the z/OS warehouse transformers" on page 88
- "Updating the environment variables on the Solaris Operating Environment for warehouse transformers" on page 76
- "Updating the environment variables on Linux for warehouse transformers" on page 77

Related reference:

- "National language support for warehouse transformers" on page 91

Setting up Java stored procedures on your DB2 subsystem for the warehouse transformers

The following instructions provide a brief version of how to set up Java stored procedures.

Procedure:

To set up Java stored procedures on your DB2 subsystem for the warehouse transformers:

1. Install Visual Age for Java 2.0 or later on your z/OS system.
2. Install Java Database Connectivity (JDBC) and bind the JDBC packages in your DB2 subsystem.
3. Set up the Revised Report on Scheme (RRS) and DB2 Work Load Manager (WLM) stored procedures for your DB2 subsystem.
4. Set up Java stored procedures for your DB2 subsystem. This includes creating a Java WLM startup procedure for the Java stored procedures address space.
5. Under WLM, associate your Java WLM startup procedure with a WLM environment name.
6. Specify the WLM application environment name for the WLM_ENVIRONMENT option on CREATE or ALTER PROCEDURE to associate a stored procedure or user-defined function with an application environment.
7. Ensure that the owner of your DB2 subsystem's started tasks has access to the libraries in the Java WLM startup procedure.

Related concepts:

- "JDK installation for warehouse transformers" on page 72
- "z/OS transformers" on page 84
- "Warehouse transformers" on page 3

Related tasks:

- "Updating the environment variables on Windows for warehouse transformers" on page 73
- "Updating the environment variables on AIX for warehouse transformers" on page 74
- "Updating the database configuration for the target database" on page 80
- "Installing warehouse transformers" on page 81
- "Enabling warehouse transformers" on page 82
- "Preparing your environment for warehouse transformers" on page 71
- "Setting up warehouse transformers on DB2 for z/OS" on page 86
- "Reducing the number of characters in the environment variable data set for the z/OS warehouse transformers" on page 88
- "Updating the environment variables on the Solaris Operating Environment for warehouse transformers" on page 76
- "Updating the environment variables on Linux for warehouse transformers" on page 77

Related reference:

- "National language support for warehouse transformers" on page 91

Setting up warehouse transformers on DB2 for z/OS

The following instructions describe how to set up the warehouse transformers on DB2 for z/OS and OS/390.

Procedure:

To set up warehouse transformers on DB2 for z/OS:

1. Define your transformer sources to the warehouse.
2. To define the transformers to DB2 for z/OS, go to the directory that contains the required SQL statement and apply the SQL statement. The location of the SQL statement varies with the version of DB2.

Table 6. SQL statement location according to DB2 Universal Database version

DB2 Universal Database Version	SQL statements location

Table 6. SQL statement location according to DB2 Universal Database version (continued)

DB2 for OS/390 Version 5	Use commented SQL statements in the/usr/lpp/DWC/createXfSQL directory (comment out all of the CREATE PROCEDURE statements; then remove comments and use the INSERT INTO SYSIBM.SYSPROCEDURES statements to define the transformers to DB2 for OS/390 Version 5)
DB2 for OS/390 Version 6	Use SQL statements in the /usr/lpp/DWC/createXfSQL directory
DB2 for z/OS and OS/390 Version 7	Use SQL statements in the /usr/lpp/DWC/createXfSQLV7 directory

When you set up Java stored procedures, use WLM to associate the Java WLM startup procedure with a WLM environment name. The environment name is specified in the WLM ENVIRONMENT option of the CREATE PROCEDURE statement. DSNWLMJ is the WLM environment name included with the transformer definitions. You can either add a WLM association name of DSNWLMJ, or change the WLM ENVIRONMENT option for each transformer definition to a name that is already associated with your startup procedure.

3. Set up links from UNIX Systems Services to the transformer load modules in IWH810.SIWHLOAD:
   a. Use telnet to connect to UNIX Systems Services on your OS/390 or z/OS host system.
   b. Change to the directory where you installed the z/OS warehouse agent. The default installation directory is /usr/lpp/DWC.
   c. If you are using DB2 V7, go to step 3e. If you are using DB2 Version 5 or Version 6, edit the trlinks data set in the installed directory.
   d. Make this line a comment by putting a pound sign (#) in column 1. For example: #ln -e IWHXF xf.jll;. Remove the comment from this line by removing the pound sign(#) in column 1. For example: ln -e IWHXFV7 xf.jll;. Save your changes.
   e. Type trlinks and press Enter. An xf.jll link is created in the /usr/lpp/DWC/com/ibm/data directory. The link will direct the agent to load either the IWHXF or IWHXFV7 modules.
4. APF-authorize IWH810.SIWHPDSE, then add it to the STEPLIB concatenation in your DB2 Java stored procedures startup procedure.
5. Add the directory where the xf.jll link is (the default is /usr/lpp/DWC) to the CLASSPATH and LIBPATH environment variables in the WLM environment data set. To find the WLM environment data set, look in the DB2 Java stored procedures startup procedure. The WLM environment data set is the one that your JAVAENV DD card points to.

6. Start the stored procedures, then create and run your warehouse steps.

Related concepts:

- "JDK installation for warehouse transformers" on page 72
- "z/OS transformers" on page 84
- "Warehouse transformers" on page 3

Related tasks:

- "Updating the environment variables on Windows for warehouse transformers" on page 73
- "Updating the environment variables on AIX for warehouse transformers" on page 74
- "Updating the database configuration for the target database" on page 80
- "Installing warehouse transformers" on page 81
- "Enabling warehouse transformers" on page 82
- "Preparing your environment for warehouse transformers" on page 71
- "Setting up Java stored procedures on your DB2 subsystem for the warehouse transformers" on page 85
- "Updating the environment variables on the Solaris Operating Environment for warehouse transformers" on page 76
- "Updating the environment variables on Linux for warehouse transformers" on page 77

Related reference:

- "National language support for warehouse transformers" on page 91

Reducing the number of characters in the environment variable data set for the z/OS warehouse transformers

The environment variable data set cannot contain more than 244 characters.

Prerequisites:

When you install JDBC, you run the installVAJDLL command. This command creates an sqlj.jll link to the JDBC driver. The name of the directory where you installed this sqlj.jll link must be added to the CLASSPATH and LIBPATH statements in your environment variable data set.

SQLj is included in a Java package named /COM/ibm/db2os390/. For example, if the sqlj.jll link is in the /usr/lpp/db2710/COM/ibm/db2os390/ directory, you need to add /usr/lpp/db2710/ to your CLASSPATH and LIBPATH statements.

Procedure:

To reduce the number of characters in the environment variable data set:

1. Create a directory such as /u/userid/links. Include both the JDBC link and the transformer link in the /u/userid/links directory:

   ```
   /u/userid/links/com/ibm/data/xf.jll
   /u/userid/links/COM/ibm/db2os390/sqlj.jll
   ```

2. Change the LIBPATH statement to LIBPATH=/u/userid/links.
3. Change the CLASSPATH statement to /u/userid/links.

If you move the links, you must execute authorities again for the .jll files so that the extended bit is on.

To determine whether the extended bit is on:

1. Go to the directory that contains the links, and enter the command ls -lt.

   The permission bit contains e in the high-order bit position. For example, erwxrwxrwx.

2. If the high order bit is l (for example, lwxrwxrwx), and you are running DB2 for OS/390 Version 5 or Version 6, you need to create the link again using the command:

   ```
   ln -e DSNAQJLL sqlj.jll
   ln -e IWHXF xf.jll
   ```

   If the high order bit is l, and you are running DB2 for z/OS and OS/390 Version 7, create the link again using the command:

   ```
   ln -e IWHXFV7 xf.jl
   ```

Java objects in the signature of a stored procedure are supported only in DB2 for z/OS and OS/390 Version 7. In DB2 for OS/390 Version 5 and Version 6, the transformers do not support null values in their parameters. In these versions, if you pass a null parameter, it is considered a zero. DB2 Version 5 and Version 6 transformers treat zero parameters like null strings.

DB2 supports the COMMIT SQL statement in stored procedures only in DB2 for z/OS and OS/390 Version 7. The INVERTDATA stored procedure drops and recreates a table within the stored procedure; therefore, it requires a COMMIT statement. IWH.INVERTDATA is not supported in DB2 for OS/390 Version 5 or Version 6.

DB2 for z/OS and OS/390 does not support Java user-defined functions, so IWH.FORMATDATE is not supported on the S/390 platform.

This is a sample startup procedure for Java stored procedures:

```
//DSNWLMJ PROC DB2SSN=DSN,NUMTCB=5,APPLENV=DSNWLMJ
//*********************************************
//* THIS PROC IS USED TO START THE WLM-ESTABLISHED SPAS *
//* ADDRESS SPACE FOR THE DSNWLMJ APPLICATION ENVIRONMENT *
//* V WLM,APPLENV=DSNWLMJ,RESUME *
//*********************************************
//DSNWLMJ EXEC PGM=DSNX9WLM,TIME=1440,REGION=0M,
// PARM='&DB2SSN, &NUMTCB, &APPLENV'
// DD DSN=IWH810.SIWHPDSE,DISP=SHR
// DD DSN=DSN.HPJSP.PDSE.JDBC,DISP=SHR
// DD DSN=SYS1.PP.PDSELINK,DISP=SHR
// DD DSN=DSN710.SDSNEXIT,DISP=SHR
// DD DSN=DSN710.SDSNLOAD,DISP=SHR
// DD DSN=SYS1.SCEERUN,DISP=SHR
// DD DSN=DSN.PDSE,DISP=SHR
//JAVAENV DD DSN=DSN.WLMENVJ.JSPENV,DISP=SHR
//CEEDUMP DD SYSOUT=A
//DSSPRINT DD SYSOUT=A
//JSPDEBUG DD SYSOUT=A
//SYSABEND DD SYSOUT=A
//SYSPRINT DD SYSOUT=A
```

In this example, the IWH810.SIWHPDSE library contains the transformer load modules. DSN.HPJSP.PDSE.JDBC contains the High Performance Java (HPJ) DLLs from HPJ setup. In DB2 for z/OS Version 7, this library is named DSN710.SDSNLOD2. SYS1.PP.PDSELINK and contains the HPJ run time libraries. DSN.PDSE contains HPJ setup information. DSN.WLMENVJ.JSPENV contains the environment variables.

This is an environment variable data set:

```
ENVAR("TZ=PST07",
"DB2SQLJPROPERTIES=/usr/lpp/db2/jdbc/db2710/classes/db2sqljjdbc.properties",
"LIBPATH=/usr/lpp/DWC",
"VWSPATH=/usr/lpp/DWC",
"CLASSPATH=/usr/lpp/db2/jdbc/db2710/classes:/usr/lpp/DWC:/usr/lpp/hpj/lib"),
MSGFILE(JSPDEBUG)
```

Related concepts:

- "JDK installation for warehouse transformers" on page 72
- "z/OS transformers" on page 84
- "Warehouse transformers" on page 3

Related tasks:

- "Updating the environment variables on Windows for warehouse transformers" on page 73
- "Updating the environment variables on AIX for warehouse transformers" on page 74
- "Updating the database configuration for the target database" on page 80

- “Installing warehouse transformers” on page 81
- “Enabling warehouse transformers” on page 82
- “Preparing your environment for warehouse transformers” on page 71
- “Setting up Java stored procedures on your DB2 subsystem for the warehouse transformers” on page 85
- “Setting up warehouse transformers on DB2 for z/OS” on page 86
- “Updating the environment variables on the Solaris Operating Environment for warehouse transformers” on page 76
- “Updating the environment variables on Linux for warehouse transformers” on page 77

Related reference:

- “National language support for warehouse transformers” on page 91

National language support for warehouse transformers

Most messages produced by the z/OS warehouse agent are sent to the warehouse server system to be interpreted, so in most cases the message language is dependent on how DB2 Universal Database was installed.

Transformers are an exception. For transformers, the message language is not dependent on how the warehouse server was installed. The z/OS warehouse agent includes the following message files for the transformers:

Table 7. National language support message files for transformers

File name	Language
Xf.properties_Fi_FI	Finnish
Xf.properties_No_NO	Norwegian
Xf.properties_Ru_RU	Russian
Xf.properties_Zh_CN	Simplified Chinese
Xf.properties_Zh_TW	Traditional Chinese
Xf.properties_Da_DK	Danish
Xf.properties_De_DE	German
Xf.properties_En_US	U.S. English
Xf.properties_Es_ES	Spanish
Xf.properties_Fr_FR	French
Xf.properties_It_IT	Italian
Xf.properties_Ja_JP	Japanese
Xf.properties_Ko_KR	Korean
Xf.properties_Pt_BR	Brazilian Portuguese

Table 7. National language support message files for transformers (continued)

Xf.properties_Sv_SE	Swedish

If your transformer messages are in a language other than English, select one of the files in Table 7 on page 91 and copy its contents to the `Xf.properties` file.

Related concepts:

- "JDK installation for warehouse transformers" on page 72
- "z/OS transformers" on page 84
- "Warehouse transformers" on page 3

Related tasks:

- "Updating the environment variables on Windows for warehouse transformers" on page 73
- "Updating the environment variables on AIX for warehouse transformers" on page 74
- "Updating the database configuration for the target database" on page 80
- "Installing warehouse transformers" on page 81
- "Enabling warehouse transformers" on page 82
- "Preparing your environment for warehouse transformers" on page 71
- "Setting up warehouse transformers on DB2 for z/OS" on page 86
- "Updating the environment variables on the Solaris Operating Environment for warehouse transformers" on page 76
- "Updating the environment variables on Linux for warehouse transformers" on page 77

Chapter 6. DB2 Warehouse Manager connectors

The DB2 Warehouse Manager provides the following Connectors to help you to extract data and metadata from e-business repositories:

- DB2 Warehouse Manager Connector for SAP® R/3®
- DB2 Warehouse Manager Connector for the Web

For installation prerequisites, see the sections in this chapter for each of the specific Connectors. All of the Connectors require DB2 Warehouse Manager.

DB2 Warehouse Manager Connector for SAP R/3

SAP R/3 is an enterprise-wide business repository from SAP AG that integrates business processes, such as those for Enterprise Resource Planning or Customer Relationship Management.

With the DB2® Warehouse Manager Connector for SAP R/3 you can bring SAP business objects that are stored in SAP R/3 systems into a DB2 data warehouse. Business objects and business components provide an object-oriented view of R/3 business functions. Business Objects encapsulate the R/3 business functions that are used to extract data from the R/3 system. You can then use the power of DB2 and DB2 Warehouse Manager for data analysis, data transformation, or data mining.

When you define an SAP source, you see all the metadata about the SAP business object, including key fields, parameter names, data types, precision, scale, length, and mandatory parameters. You also see all basic and detailed parameters that are associated with the SAP business object. You define the data extraction step from the Data Warehouse Center by first clicking and dropping an SAP business object onto the Process Modeler window, then clicking and dropping the SAP Data Extract step icon onto the process modeler, then adding a data link between the business object and the step icon, and finally by adding the required information to the step properties notebook.

The DB2 Warehouse Manager Connector for SAP R/3 runs on Microsoft® Windows® NT, Windows 2000, and Windows XP. The SAP R/3 server can be on any supported platform.

Related concepts:

• "DB2 Warehouse Manager Connector for the Web" on page 95

- “DB2 Warehouse Manager connectors” on page 4

Related tasks:

- “Installing the DB2 Warehouse Manager Connector for SAP R/3” on page 94

Installing the DB2 Warehouse Manager Connector for SAP R/3

With the DB2 Warehouse Manager Connector for SAP R/3, you can add the extracted data to a data warehouse, transform it using the DB2 Data Warehouse Center, or analyze it using DB2 tools or other vendors’ tools.

Prerequisites:

The DB2 Warehouse Manager Connector for SAP R/3 supports the following operating systems:

- Microsoft Windows NT Workstation or Windows NT Server Version 4.0 with Service Pack 5 or later
- Microsoft Windows 2000 (with Service Pack 2) or Windows XP

Before you install the DB2 Warehouse Manager Connector for SAP R/3, complete the following tasks:

- Install RFC Runtime modules from the the Presentation CD, SAP Release 4.6D Compilation 4 (included with SAP clients) on the same workstation on which you plan to install the Connector for SAP. The file that contains the RFC runtime modules is libRfc32.dll.
- Optionally, install SAPGUI for troubleshooting.

The DB2 Warehouse Manager Connector for SAP R/3 works with SAP R/3 Release 4.6C servers with the English language option installed. The DB2 Warehouse Manager Connector for SAP R/3 connects to an R/3 server with SAP logon language ’E’.

Restrictions:

The DB2 Warehouse Manager Connector for SAP R/3 supports only English-language installations of the SAP R/3 system.

Procedure:

To install the DB2 Warehouse Manager Connector for SAP R/3, insert the DB2 Warehouse Manager Connectors CD-ROM into your CD drive, and follow the on-screen instructions.

Related concepts:

- "DB2 Warehouse Manager Connector for SAP R/3" on page 93
- "DB2 Warehouse Manager Connector for the Web" on page 95
- "DB2 Warehouse Manager connectors" on page 4

Installing the DB2 Warehouse Manager Connector for the Web

With the DB2 Warehouse Manager Connector for the Web, you can bring clickstream data from IBM WebSphere Site Analyzer into a data warehouse.

Prerequisites:

Before you install the DB2 Warehouse Manager Connector for the Web, complete the following tasks:

- Install and configure WebSphere Site Analyzer version 4 (with Fixpak 1) or later.

Your programs must run on the same workstation where the warehouse agent is installed.

Procedure:

To install the DB2 Warehouse Manager Connector for the Web, insert the DB2 Warehouse Manager Connectors CD-ROM into your CD drive, and follow the instructions that are displayed.

Related concepts:

- "DB2 Warehouse Manager Connector for SAP R/3" on page 93
- "DB2 Warehouse Manager Connector for the Web" on page 95
- "DB2 Warehouse Manager connectors" on page 4

DB2 Warehouse Manager Connector for the Web

IBM® WebSphere® Site Analyzer (WSA) is part of the IBM WebSphere family of Web servers and application servers. WSA helps you analyze the traffic to and from your Web site.

The Connector for the Web allows you to extract data from a WebSphere Site Analyzer database, or webmart, into a data warehouse. The Connector for the Web provides a polling step that checks whether WSA copied Web traffic data from its data imports (log files, tables, and other sources of clickstream data) to the webmart. After this check is successful, an SQL step could copy the Web traffic data from the webmart into a warehouse target. You can then use

the power of DB2® and DB2 Warehouse Manager for data analysis, data transformation, or data mining. You can also incorporate WebSphere Commerce data with the Web traffic data for a more complete analysis of your Web site.

After defining a Web traffic source, you can define the Web traffic polling step from the Data Warehouse Center by simply clicking and dropping a Web Traffic Polling step onto the process modeler.

The Connector for the Web runs on the same platform as the warehouse agent: Windows® NT, Windows 2000, Windows XP, AIX, or the Solaris Operating Environment.

Related concepts:

- "DB2 Warehouse Manager Connector for SAP R/3" on page 93
- "DB2 Warehouse Manager connectors" on page 4

Related tasks:

- "Installing the DB2 Warehouse Manager Connector for the Web" on page 95

Appendix.

Environment structure for Information Catalog Manager components

When you install the Information Catalog Center, entries are added or modified for the several Windows environment variables and user variables. In the following tables, the drive letter C indicates a local drive.

Table 8. Environment variable updates for Information Catalog Manager components

The environment variable:	**Is added to, or modified, to include:**
PATH	C:\SQLLIB\BIN
INCLUDE	C:\SQLLIB\LIB
VWS_TEMPLATES	C:\SQLLIB\TEMPLATES

Table 9. User variable updates for the Information Catalog Manager

The user variable:	**Is added to, or modified, to include:**
LOCPATH	%LOCPATH%
LIB	%LIB%
INCLUDE	%INCLUDE%

Warehouse agent environment structures

The information in this appendix describes the structure of the AIX, Linux, and Solaris Operating Environment warehouse agent environments. Use this information to help you configure the AIX, Linux, and Solaris Operating Environment warehouse agent environments.

The following tables show the directory structures for the AIX, Linux, and Solaris Operating Environment warehouse agents.

Table 10. Directory structures for the AIX warehouse agent

This directory:	**Contains:**
/usr/opt/db2_v8_01/bin/IWH.startup	Agent daemon and server startup file.
/usr/opt/db2_08_01/bin	Agent, step executables, user-defined executable files, and the environment setup file. This is the default install directory.

Table 10. Directory structures for the AIX warehouse agent (continued)

This directory:	Contains:
/usr/opt/db2_08_01/odbc	Sample .odbc.ini file and Data Warehouse Center ODBC driver.
/var/IWH	Trace files. This is the default directory for the daemon traces, agent traces, and user-defined program output files.

Table 11. Directory structures for the Solaris Operating Environment and Linux warehouse agents

This directory:	Contains:
/usr/opt/db2_v8_01/bin/IWH.startup	Agent daemon startup file.
/opt/IBM/db2/V8.1/	Default install base directory.
/opt/IBM/db2/V8.1/bin	Agent, step executables, user-defined executable files, and the environment setup file.
/opt/IBM/db2/V8.1/odbc	Sample .odbc.ini file and Data Warehouse Center ODBC driver.
/opt/IBM/db2/V8.1/lib	Solaris Operating Environment and Linux run-time library.
/var/IWH	Trace files. This is the default directory for the daemon traces, agent traces, and user-defined program output files.

Related concepts:

- "Communications between Data Warehouse Center clients and servers" on page 101
- "Communications between Data Warehouse Center servers and warehouse agents" on page 102

Identifying ports for Data Warehouse Center components if your system configuration uses a firewall

If you install the Data Warehouse Center server, client, and agent components across a firewall, you need to identify the ports that these components will use to communicate with each other. To identify these ports to your system, you will update the `.../etc/services` file on each of the systems where a Data Warehouse Center component is installed. You will also add service information for each system in the `.../etc/services` file, and you will make this information available to the firewall administrator.

If you plan to use firewall protection between any Data Warehouse Center component, you must set up the ports on all of the systems where a Data Warehouse Center component is installed, not just on the systems where communication across the firewall will occur. Failure to do this might cause a

Data Warehouse Center component to allocate a port outside the firewall port range and cause a firewall security alert.

Procedure:

To support a firewall, you must specify a set of ports to be used by the Data Warehouse Center and the firewall administrator. These ports will enable communications between the various Data Warehouse Center components (server, client, and warehouse agents). You must identify a range of ports that the various Data Warehouse Center components can use to communicate with each other. This must be done on any system on which the component is installed.

To identify ports for Data Warehouse Center components when your system configuration uses a firewall:

1. On client systems, open the `.../etc/services` file and add the following lines:

Service Name	Service Port
DWC_MIN	nnnnn
DWC_MAX	nnnnn + x

   where:

   - nnnnn is the lowest port number that you want the Data Warehouse Center client to use when it communicates with a Data Warehouse Center server across the firewall
   - nnnnn + x is the highest port number that you want the Data Warehouse Center client to use when it communicates with a Data Warehouse Center server across the firewall

   You will need to make sure that the range between the value of DWC_MIN and DWC_MAX is sufficient to support two ports for each client that might run concurrently on that system. For example, say that you install the Data Warehouse Center client on a Solaris Operating Environment system. You want the Data Warehouse Center components to communicate with each other starting at port 11010, and you want to have a maximum of 10 clients running concurrently on that system. You should set DWC_MIN to 11010 and DWC_MAX to 11030 (10 concurrent clients = 20 ports).

   The Data Warehouse Center client will now communicate only across ports within the DWC_MIN and DWC_MAX range. If a new Data Warehouse Center client tries to log on and finds no available ports in that range, an error message will be displayed. The client must wait until ports within that range become available, or the system administrator will need to

increase the number of available ports by increasing the value of the DWC_MAX to accommodate additional clients.

2. On server systems, open the `.../etc/services` file and add the following lines:

Service Name	Service Port
DWC_MIN	nnnnn
DWC_MAX	nnnnn + x

   where:

   - nnnnn is the lowest port number that you want the Data Warehouse Center server to use when it communicates with a Data Warehouse Center client or agent across the firewall
   - nnnnn + x is the highest port number that you want the Data Warehouse Center server to use when it communicates with a Data Warehouse Center client or agent across the firewall

   You will need to make sure that the range between the value of DWC_MIN and DWC_MAX is sufficient to support two ports for each client that could run concurrently on that system, one port for each Data Warehouse Center agent instance with which the server will communicate concurrently, and one port for communications with the Data Warehouse Center logger.

   For example, say that you installed the Data Warehouse Center server on a Windows NT system and this server will communicate with ten Data Warehouse Center clients on an AIX system and five Data Warehouse Center clients on Windows NT systems. The server will communicate with the local Data Warehouse Center logger and an AIX, Windows NT, and z/OS agent (of which a maximum of five instances can be run concurrently). If you want the Data Warehouse Center components to communicate with each other starting at port 11055, set DWC_MIN to 11055 and set DWC_MAX to 11101 (15 concurrent clients = 30 ports) , 3 agents with 5 concurrent instances each = 15 ports, and 1 logger = 1 port for a total of 46 ports).

   The Data Warehouse Center server will now communicate only across ports within the DWC_MIN and DWC_MAX range. If the Data Warehouse Center server tries to respond to a new Data Warehouse Center client logon attempt or to start a new agent instance, but no ports are available within the defined range, an error message will be displayed. The client must wait until ports within the defined range become available, or the system administrator must specify a larger number of available ports by increasing the value of the DWC_MAX to accommodate additional clients or agent instances.

3. On warehouse agent site systems, open the .../etc/services file and add the following lines:

Service Name	Service Port
DWC_MIN	nnnnn
DWC_MAX	nnnnn + x

   where:

   - nnnnn is the lowest port number that you want the warehouse agent to use when it communicates with a Data Warehouse Center server across the firewall
   - nnnnn + x is the highest port number that you want the warehouse agent to use when it communicates with a Data Warehouse Center server across the firewall

   You will need to make sure that the range between the value of DWC_MIN and DWC_MAX is sufficient to support one port for each warehouse agent instance that could run concurrently on that system.

   For example, say that you install the Data Warehouse Center server on a Windows system. You want to have the Data Warehouse Center server communicate with the warehouse agent starting at port 11025, and you want to have a maximum of ten agents running concurrently (ten Data Warehouse Center steps that could be run concurrently by this agent will cause ten agent instances to execute concurrently). You should set DWC_MIN to 11025 and DWC_MAX to 11035 (10 concurrent steps = 10 concurrent agent instances = 10 ports).

Related concepts:

- "Communications between Data Warehouse Center clients and servers" on page 101
- "Communications between Data Warehouse Center servers and warehouse agents" on page 102

Communications between Data Warehouse Center clients and servers

When the Data Warehouse Center client tries to log on to the Data Warehouse Center server, it first finds an available port on its system. Then the client sends a message to the Data Warehouse Center server using port 11000 to request a logon. The message specifies the available port on the client system. The client uses this port to listen for a response from the server. When the server receives the message to log on, it validates the client system that made the request and then searches the message to determine the port on which it is to respond. Next, the server opens communication with the client (on the

port that the client specified). The server also specifies to the client the port through which the client should communicate with the server.

Each Data Warehouse Center client (Data Warehouse Center Admin Client and Work in Progress client) that will connect to a Data Warehouse Center server must open a communications pipe with a server as described in this section. In this way, a server can process many client requests concurrently. The server must have a separate communications pipe (open TCPIP port) with each concurrently connected client.

Related concepts:

- "Communications between Data Warehouse Center servers and warehouse agents" on page 102

Related reference:

- "Warehouse agent environment structures" on page 97

Communications between Data Warehouse Center servers and warehouse agents

When the Data Warehouse Center server is asked to complete a task that requires the use of a warehouse agent (for example, to run a step with a schedule or to read metadata from the database catalogs), the server finds an available port on its system and then sends a message to the warehouse agent daemon (listening on port 11001) at the agent site. The message contains the number of the port to which the the agent can respond. The warehouse agent daemon receives this message and does some basic validation. After the message is validated, the warehouse agent daemon starts a warehouse agent instance to process the request from the server. The warehouse agent starts and accepts the message sent by the Data Warehouse Center server, finds an available port on its system, and responds to the server, using the port specified in the message from the server. During the response, the warehouse agent also indicates the port on the warehouse agent system that it will use to receive additional requests from the server.

Every time the Data Warehouse Center server needs an agent do a specific task, it must perform a handshake on the ports over which it communicates with the agent. Because the server can handle multiple schedules and client requests at one time, many communication pipes between an agent and a server might be open at the same time.

Related concepts:

- "Communications between Data Warehouse Center clients and servers" on page 101

Related reference:

- "Warehouse agent environment structures" on page 97

DB2 Universal Database technical information

Overview of DB2 Universal Database technical information

DB2 Universal Database technical information can be obtained in the following formats:

- Books (PDF and hard-copy formats)
- A topic tree (HTML format)
- Help for DB2 tools (HTML format)
- Sample programs (HTML format)
- Command line help
- Tutorials

This section is an overview of the technical information that is provided and how you can access it.

FixPaks for DB2 documentation

IBM may periodically make documentation FixPaks available. Documentation FixPaks allow you to update the information that you installed from the *DB2 HTML Documentation CD* as new information becomes available.

Note: If you do install documentation FixPaks, your HTML documentation will contain more recent information than either the DB2 printed or online PDF manuals.

Categories of DB2 technical information

The DB2 technical information is categorized by the following headings:

- Core DB2 information
- Administration information
- Application development information
- Business intelligence information
- DB2 Connect information
- Getting started information
- Tutorial information
- Optional component information
- Release notes

The following tables describe, for each book in the DB2 library, the information needed to order the hard copy, print or view the PDF, or locate the HTML directory for that book. A full description of each of the books in

the DB2 library is available from the IBM Publications Center at www.ibm.com/shop/publications/order

The installation directory for the HTML documentation CD differs for each category of information:

htmlcdpath/doc/htmlcd/*%L*/*category*

where:

- *htmlcdpath* is the directory where the HTML CD is installed.
- *%L* is the language identifier. For example, en_US.
- *category* is the category identifier. For example, core for the core DB2 information.

In the PDF file name column in the following tables, the character in the sixth position of the file name indicates the language version of a book. For example, the file name db2d1e80 identifies the English version of the *Administration Guide: Planning* and the file name db2d1g80 identifies the German version of the same book. The following letters are used in the sixth position of the file name to indicate the language version:

Language	Identifier
Arabic	w
Brazilian Portuguese	b
Bulgarian	u
Croatian	9
Czech	x
Danish	d
Dutch	q
English	e
Finnish	y
French	f
German	g
Greek	a
Hungarian	h
Italian	i
Japanese	j
Korean	k
Norwegian	n
Polish	p
Portuguese	v
Romanian	8
Russian	r
Simp. Chinese	c
Slovakian	7
Slovenian	l

Spanish	z
Swedish	s
Trad. Chinese	t
Turkish	m

No form number indicates that the book is only available online and does not have a printed version.

Core DB2 information

The information in this category covers DB2 topics that are fundamental to all DB2 users. You will find the information in this category useful whether you are a programmer, a database administrator, or you work with DB2 Connect, DB2 Warehouse Manager, or other DB2 products.

The installation directory for this category is `doc/htmlcd/%L/core`.

Table 12. Core DB2 information

Name	Form Number	PDF File Name
IBM DB2 Universal Database Command Reference	SC09-4828	db2n0x80
IBM DB2 Universal Database Glossary	No form number	db2t0x80
IBM DB2 Universal Database Master Index	SC09-4839	db2w0x80
IBM DB2 Universal Database Message Reference, Volume 1	GC09-4840	db2m1x80
IBM DB2 Universal Database Message Reference, Volume 2	GC09-4841	db2m2x80
IBM DB2 Universal Database What's New	SC09-4848	db2q0x80

Administration information

The information in this category covers those topics required to effectively design, implement, and maintain DB2 databases, data warehouses, and federated systems.

The installation directory for this category is doc/htmlcd/%L/admin.

Table 13. Administration information

Name	Form number	PDF file name
IBM DB2 Universal Database Administration Guide: Planning	SC09-4822	db2d1x80
IBM DB2 Universal Database Administration Guide: Implementation	SC09-4820	db2d2x80
IBM DB2 Universal Database Administration Guide: Performance	SC09-4821	db2d3x80
IBM DB2 Universal Database Administrative API Reference	SC09-4824	db2b0x80
IBM DB2 Universal Database Data Movement Utilities Guide and Reference	SC09-4830	db2dmx80
IBM DB2 Universal Database Data Recovery and High Availability Guide and Reference	SC09-4831	db2hax80
IBM DB2 Universal Database Data Warehouse Center Administration Guide	SC27-1123	db2ddx80
IBM DB2 Universal Database Federated Systems Guide	GC27-1224	db2fpx80
IBM DB2 Universal Database Guide to GUI Tools for Administration and Development	SC09-4851	db2atx80
IBM DB2 Universal Database Replication Guide and Reference	SC27-1121	db2e0x80
IBM DB2 Installing and Administering a Satellite Environment	GC09-4823	db2dsx80
IBM DB2 Universal Database SQL Reference, Volume 1	SC09-4844	db2s1x80
IBM DB2 Universal Database SQL Reference, Volume 2	SC09-4845	db2s2x80
IBM DB2 Universal Database System Monitor Guide and Reference	SC09-4847	db2f0x80

Application development information

The information in this category is of special interest to application developers or programmers working with DB2. You will find information about supported languages and compilers, as well as the documentation required to access DB2 using the various supported programming interfaces, such as embedded SQL, ODBC, JDBC, SQLj, and CLI. If you view this information online in HTML you can also access a set of DB2 sample programs in HTML.

The installation directory for this category is `doc/htmlcd/%L/ad`.

Table 14. Application development information

Name	Form number	PDF file name
IBM DB2 Universal Database Application Development Guide: Building and Running Applications	SC09-4825	db2axx80
IBM DB2 Universal Database Application Development Guide: Programming Client Applications	SC09-4826	db2a1x80
IBM DB2 Universal Database Application Development Guide: Programming Server Applications	SC09-4827	db2a2x80
IBM DB2 Universal Database Call Level Interface Guide and Reference, Volume 1	SC09-4849	db2l1x80
IBM DB2 Universal Database Call Level Interface Guide and Reference, Volume 2	SC09-4850	db2l2x80
IBM DB2 Universal Database Data Warehouse Center Application Integration Guide	SC27-1124	db2adx80
IBM DB2 XML Extender Administration and Programming	SC27-1234	db2sxx80

Business intelligence information

The information in this category describes how to use components that enhance the data warehousing and analytical capabilities of DB2 Universal Database.

The installation directory for this category is doc/htmlcd/%L/wareh.

Table 15. Business intelligence information

Name	Form number	PDF file name
IBM DB2 Warehouse Manager Information Catalog Center Administration Guide	SC27-1125	db2dix80
IBM DB2 Warehouse Manager Installation Guide	GC27-1122	db2idx80

DB2 Connect information

The information in this category describes how to access host or iSeries data using DB2 Connect Enterprise Edition or DB2 Connect Personal Edition.

The installation directory for this category is doc/htmlcd/%L/conn.

Table 16. DB2 Connect information

Name	Form number	PDF file name
APPC, CPI-C, and SNA Sense Codes	No form number	db2apx80
IBM Connectivity Supplement	No form number	db2h1x80
IBM DB2 Connect Quick Beginnings for DB2 Connect Enterprise Edition	GC09-4833	db2c6x80
IBM DB2 Connect Quick Beginnings for DB2 Connect Personal Edition	GC09-4834	db2c1x80
IBM DB2 Connect User's Guide	SC09-4835	db2c0x80

Getting started information

The information in this category is useful when you are installing and configuring servers, clients, and other DB2 products.

The installation directory for this category is doc/htmlcd/%L/start.

Table 17. Getting started information

Name	Form number	PDF file name
IBM DB2 Universal Database Quick Beginnings for DB2 Clients	GC09-4832	db2itx80

Table 17. Getting started information (continued)

Name	Form number	PDF file name
IBM DB2 Universal Database Quick Beginnings for DB2 Servers	GC09-4836	db2isx80
IBM DB2 Universal Database Quick Beginnings for DB2 Personal Edition	GC09-4838	db2i1x80
IBM DB2 Universal Database Installation and Configuration Supplement	GC09-4837	db2iyx80
IBM DB2 Universal Database Quick Beginnings for DB2 Data Links Manager	GC09-4829	db2z6x80

Tutorial information

Tutorial information introduces DB2 features and teaches how to perform various tasks.

The installation directory for this category is `doc/htmlcd/%L/tutr`.

Table 18. Tutorial information

Name	Form number	PDF file name
Business Intelligence Tutorial: Introduction to the Data Warehouse	No form number	db2tux80
Business Intelligence Tutorial: Extended Lessons in Data Warehousing	No form number	db2tax80
Development Center Tutorial for Video Online using Microsoft Visual Basic	No form number	db2tdx80
Information Catalog Center Tutorial	No form number	db2aix80
Video Central for e-business Tutorial	No form number	db2twx80
Visual Explain Tutorial	No form number	db2tvx80

Optional component information

The information in this category describes how to work with optional DB2 components.

The installation directory for this category is doc/htmlcd/%L/opt.

Table 19. Optional component information

Name	Form number	PDF file name
IBM DB2 Life Sciences Data Connect Planning, Installation, and Configuration Guide	GC27-1235	db2lsx80
IBM DB2 Spatial Extender User's Guide and Reference	SC27-1226	db2sbx80
IBM DB2 Universal Database Data Links Manager Administration Guide and Reference	SC27-1221	db2z0x80
IBM DB2 Universal Database Net Search Extender Administration and Programming Guide **Note:** HTML for this document is not installed from the HTML documentation CD.	SH12-6740	N / A

Release notes

The release notes provide additional information specific to your product's release and FixPak level. They also provides summaries of the documentation updates incorporated in each release and FixPak.

Table 20. Release notes

Name	Form number	PDF file name	HTML directory
DB2 Release Notes	See note.	See note.	See note.
DB2 Installation Notes	Available on product CD-ROM only.	Available on product CD-ROM only.	

Note: The HTML version of the release notes is available from the Information Center and on the product CD-ROMs. To view the ASCII file on UNIX-based platforms, see the Release.Notes file. This file is located in the DB2DIR/Readme/%L directory, where %L represents the locale name and DB2DIR represents:

- /usr/opt/db2_08_01 on AIX
- /opt/IBM/db2/V8.1 on all other UNIX operating systems

Related tasks:

Printing DB2 books from PDF files

You can print DB2 books from the PDF files on the *DB2 PDF Documentation* CD. Using Adobe Acrobat Reader, you can print either the entire book or a specific range of pages.

Prerequisites:

Ensure that you have Adobe Acrobat Reader. It is available from the Adobe Web site at www.adobe.com

Procedure:

To print a DB2 book from a PDF file:

1. Insert the *DB2 PDF Documentation* CD. On UNIX operating systems, mount the DB2 PDF Documentation CD. Refer to your *Quick Beginnings* book for details on how to mount a CD on UNIX operating systems.
2. Start Adobe Acrobat Reader.
3. Open the PDF file from one of the following locations:
   - On Windows operating systems:

     x:\doc*language* directory, where *x* represents the CD-ROM drive letter and *language* represents the two-character territory code that represents your language (for example, EN for English).
   - On UNIX operating systems:

     /cdrom/doc/*%L* directory on the CD-ROM, where */cdrom* represents the mount point of the CD-ROM and *%L* represents the name of the desired locale.

Related tasks:

- "Ordering printed DB2 books" on page 114
- "Finding product information by accessing the DB2 Information Center from the administration tools" on page 118
- "Viewing technical documentation online directly from the DB2 HTML Documentation CD" on page 120

Related reference:

- "Overview of DB2 Universal Database technical information" on page 105

Ordering printed DB2 books

Procedure:

To order printed books:

- Contact your IBM authorized dealer or marketing representative. To find a local IBM representative, check the IBM Worldwide Directory of Contacts at www.ibm.com/planetwide
- Phone 1-800-879-2755 in the United States or 1-800-IBM-4YOU in Canada.
- Visit the IBM Publications Center at www.ibm.com/shop/publications/order

You can also obtain printed DB2 manuals by ordering Doc Packs for your DB2 product from your IBM Reseller. The Doc Packs are subsets of the manuals in the DB2 library selected to help you to get started using the DB2 product that you purchased. The manuals in the Doc Packs are the same as those that are available in PDF format on the *DB2 PDF Documentation CD* and contain the same content as the documentation that is available on the *DB2 HTML Documentation CD*.

Related tasks:

- "Printing DB2 books from PDF files" on page 113
- "Finding topics by accessing the DB2 Information Center from a browser" on page 116
- "Viewing technical documentation online directly from the DB2 HTML Documentation CD" on page 120

Related reference:

- "Overview of DB2 Universal Database technical information" on page 105

Accessing online help

The online help that comes with all DB2 components is available in three types:

- Window and notebook help
- Command line help
- SQL statement help

Window and notebook help explain the tasks that you can perform in a window or notebook and describe the controls. This help has two types:

- Help accessible from the **Help** button
- Infopops

The **Help** button gives you access to overview and prerequisite information. The infopops describe the controls in the window or notebook. Window and notebook help are available from DB2 centers and components that have user interfaces.

Command line help includes Command help and Message help. Command help explains the syntax of commands in the command line processor. Message help describes the cause of an error message and describes any action you should take in response to the error.

SQL statement help includes SQL help and SQLSTATE help. DB2 returns an SQLSTATE value for conditions that could be the result of an SQL statement. SQLSTATE help explains the syntax of SQL statements (SQL states and class codes).

Note: SQL help is not available for UNIX operating systems.

Procedure:

To access online help:

- For window and notebook help, click **Help** or click that control, then click **F1**. If the **Automatically display infopops** check box on the **General** page of the **Tool Settings** notebook is selected, you can also see the infopop for a particular control by holding the mouse cursor over the control.
- For command line help, open the command line processor and enter:
  - For Command help:

    ```
    ? command
    ```

    where *command* represents a keyword or the entire command.

    For example, `? catalog` displays help for all the CATALOG commands, while `? catalog database` displays help for the CATALOG DATABASE command.
- For Message help:

  ```
  ? XXXnnnnn
  ```

  where *XXXnnnnn* represents a valid message identifier.

  For example, `? SQL30081` displays help about the SQL30081 message.

- For SQL statement help, open the command line processor and enter:

  ? *sqlstate* or ? *class code*

  where *sqlstate* represents a valid five-digit SQL state and *class code* represents the first two digits of the SQL state.

  For example, ? 08003 displays help for the 08003 SQL state, while ? 08 displays help for the 08 class code.

Related tasks:

- "Finding topics by accessing the DB2 Information Center from a browser" on page 116
- "Viewing technical documentation online directly from the DB2 HTML Documentation CD" on page 120

Finding topics by accessing the DB2 Information Center from a browser

The DB2 Information Center accessed from a browser enables you to access the information you need to take full advantage of DB2 Universal Database and DB2 Connect. The DB2 Information Center also documents major DB2 features and components including replication, data warehousing, metadata, and DB2 extenders.

The DB2 Information Center accessed from a browser is composed of the following major elements:

Navigation tree
: The navigation tree is located in the left frame of the browser window. The tree expands and collapses to show and hide topics, the glossary, and the master index in the DB2 Information Center.

Navigation toolbar
: The navigation toolbar is located in the top right frame of the browser window. The navigation toolbar contains buttons that enable you to search the DB2 Information Center, hide the navigation tree, and find the currently displayed topic in the navigation tree.

Content frame
: The content frame is located in the bottom right frame of the browser window. The content frame displays topics from the DB2 Information Center when you click on a link in the navigation tree, click on a search result, or follow a link from another topic or from the master index.

Prerequisites:

To access the DB2 Information Center from a browser, you must use one of the following browsers:

- Microsoft Explorer, version 5 or later
- Netscape Navigator, version 6.1 or later

Restrictions:

The DB2 Information Center contains only those sets of topics that you chose to install from the *DB2 HTML Documentation CD*. If your Web browser returns a `File not found` error when you try to follow a link to a topic, you must install one or more additional sets of topics from the *DB2 HTML Documentation CD*.

Procedure:

To find a topic by searching with keywords:

1. In the navigation toolbar, click **Search**.
2. In the top text entry field of the Search window, enter one or more terms related to your area of interest and click **Search**. A list of topics ranked by accuracy displays in the **Results** field. The numerical ranking beside the hit provides an indication of the strength of the match (bigger numbers indicate stronger matches).

   Entering more terms increases the precision of your query while reducing the number of topics returned from your query.
3. In the **Results** field, click the title of the topic you want to read. The topic displays in the content frame.

To find a topic in the navigation tree:

1. In the navigation tree, click the book icon of the category of topics related to your area of interest. A list of subcategories displays underneath the icon.
2. Continue to click the book icons until you find the category containing the topics in which you are interested. Categories that link to topics display the category title as an underscored link when you move the cursor over the category title. The navigation tree identifies topics with a page icon.
3. Click the topic link. The topic displays in the content frame.

To find a topic or term in the master index:

1. In the navigation tree, click the “Index” category. The category expands to display a list of links arranged in alphabetical order in the navigation tree.
2. In the navigation tree, click the link corresponding to the first character of the term relating to the topic in which you are interested. A list of terms

with that initial character displays in the content frame. Terms that have multiple index entries are identified by a book icon.

3. Click the book icon corresponding to the term in which you are interested. A list of subterms and topics displays below the term you clicked. Topics are identified by page icons with an underscored title.
4. Click on the title of the topic that meets your needs. The topic displays in the content frame.

Related concepts:

- "Accessibility" on page 125
- "DB2 Information Center accessed from a browser" on page 128

Related tasks:

- "Finding product information by accessing the DB2 Information Center from the administration tools" on page 118
- "Updating the HTML documentation installed on your machine" on page 120
- "Troubleshooting DB2 documentation search with Netscape 4.x" on page 123
- "Searching the DB2 documentation" on page 124

Related reference:

- "Overview of DB2 Universal Database technical information" on page 105

Finding product information by accessing the DB2 Information Center from the administration tools

The DB2 Information Center provides quick access to DB2 product information and is available on all operating systems for which the DB2 administration tools are available.

The DB2 Information Center accessed from the tools provides six types of information.

Tasks Key tasks you can perform using DB2.

Concepts
Key concepts for DB2.

Reference
DB2 reference information, such as keywords, commands, and APIs.

Troubleshooting
Error messages and information to help you with common DB2 problems.

Samples
Links to HTML listings of the sample programs provided with DB2.

Tutorials
Instructional aid designed to help you learn a DB2 feature.

Prerequisites:

Some links in the DB2 Information Center point to Web sites on the Internet. To display the content for these links, you will first have to connect to the Internet.

Procedure:

To find product information by accessing the DB2 Information Center from the tools:

1. Start the DB2 Information Center in one of the following ways:
   - From the graphical administration tools, click on the **Information Center** icon in the toolbar. You can also select it from the **Help** menu.
   - At the command line, enter **db2ic**.
2. Click the tab of the information type related to the information you are attempting to find.
3. Navigate through the tree and click on the topic in which you are interested. The Information Center will then launch a Web browser to display the information.
4. To find information without browsing the lists, click the **Search** icon to the right of the list.

   Once the Information Center has launched a browser to display the information, you can perform a full-text search by clicking the **Search** icon in the navigation toolbar.

Related concepts:

- "Accessibility" on page 125
- "DB2 Information Center accessed from a browser" on page 128

Related tasks:

- "Finding topics by accessing the DB2 Information Center from a browser" on page 116
- "Searching the DB2 documentation" on page 124

Viewing technical documentation online directly from the DB2 HTML Documentation CD

All of the HTML topics that you can install from the *DB2 HTML Documentation CD* can also be read directly from the CD. Therefore, you can view the documentation without having to install it.

Restrictions:

As the Tools help is installed from the DB2 product CD and not from the *DB2 HTML Documentation CD*, you must install the DB2 product to view the help.

Procedure:

1. Insert the *DB2 HTML Documentation* CD. On UNIX operating systems, mount the *DB2 HTML Documentation CD*. Refer to your *Quick Beginnings* book for details on how to mount a CD on UNIX operating systems.
2. Start your HTML browser and open the appropriate file:
   - For Windows operating systems:

     ```
     e:\program files\IBM\SQLLIB\doc\htmlcd\%L\index.htm
     ```

     where *e* represents the CD-ROM drive, and %L is the locale of the documentation that you wish to use, for example, **en_US** for English.
   - For UNIX operating systems:

     ```
     /cdrom/program files/IBM/SQLLIB/doc/htmlcd/%L/index.htm
     ```

     where */cdrom/* represents where the CD is mounted, and %L is the locale of the documentation that you wish to use, for example, **en_US** for English.

Related tasks:

- "Finding topics by accessing the DB2 Information Center from a browser" on page 116
- "Copying files from the DB2 HTML Documentation CD to a Web server" on page 122

Related reference:

- "Overview of DB2 Universal Database technical information" on page 105

Updating the HTML documentation installed on your machine

It is now possible to update the HTML installed from the *DB2 HTML Documentation CD* when updates are made available from IBM. This can be done in one of two ways:

- Using the Information Center (if you have the DB2 administration GUI tools installed).
- By downloading and applying a DB2 HTML documentation FixPak .

Note: This will NOT update the DB2 code; it will only update the HTML documentation installed from the *DB2 HTML Documentation CD*.

Procedure:

To use the Information Center to update your local documentation:

1. Start the DB2 Information Center in one of the following ways:
   - From the graphical administration tools, click on the **Information Center** icon in the toolbar. You can also select it from the **Help** menu.
   - At the command line, enter **db2ic**.
2. Ensure your machine has access to the external Internet; the updater will download the latest documentation FixPak from the IBM server if required.
3. Select **Information Center** —> **Update Local Documentation** from the menu to start the update.
4. Supply your proxy information (if required) to connect to the external Internet.

The Information Center will download and apply the latest documentation FixPak, if one is available.

To manually download and apply the documentation FixPak :

1. Ensure your machine is connected to the Internet.
2. Open the DB2 support page in your Web browser at: www.ibm.com/software/data/db2/udb/winos2unix/support.
3. Follow the link for Version 8 and look for the ″Documentation FixPaks″ link.
4. Determine if the version of your local documentation is out of date by comparing the documentation FixPak level to the documentation level you have installed. This current documentation on your machine is at the following level: **DB2 v8.1 GA**.
5. If there is a more recent version of the documentation available then download the FixPak applicable to your operating system. There is one FixPak for all Windows platforms, and one FixPak for all UNIX platforms.
6. Apply the FixPak:
   - For Windows operating systems: The documentation FixPak is a self extracting zip file. Place the downloaded documentation FixPak in an empty directory, and run it. It will create a **setup** command which you can run to install the documentation FixPak.

- For UNIX operating systems: The documentation FixPak is a compressed tar.Z file. Uncompress and untar the file. It will create a directory named `delta_install` with a script called **installdocfix**. Run this script to install the documentation FixPak.

Related tasks:

- "Copying files from the DB2 HTML Documentation CD to a Web server" on page 122

Related reference:

- "Overview of DB2 Universal Database technical information" on page 105

Copying files from the DB2 HTML Documentation CD to a Web server

The entire DB2 information library is delivered to you on the *DB2 HTML Documentation CD* and may be installed on a Web server for easier access. Simply copy to your Web server the documentation for the languages that you want.

Note: You might encounter slow performance if you access the HTML documentation from a Web server through a low-speed connection.

Procedure:

To copy files from the *DB2 HTML Documentation CD* to a Web server, use the appropriate source path:

- For Windows operating systems:
  E:\program files\IBM\SQLLIB\doc\htmlcd*%L**.*

  where *E* represents the CD-ROM drive and *%L* represents the language identifier.
- For UNIX operating systems:
  /*cdrom*/program files/IBM/sqllib/doc/htmlcd/*%L*/*.*

  where *cdrom* represents the mount point for the CD-ROM drive and *%L* represents the language identifier.

Related tasks:

- "Searching the DB2 documentation" on page 124

Related reference:

- "Supported DB2 interface languages, locales, and code pages" in the *Quick Beginnings for DB2 Servers*
- "Overview of DB2 Universal Database technical information" on page 105

Troubleshooting DB2 documentation search with Netscape 4.x

Most search problems are related to the Java support provided by web browsers. This task describes possible workarounds.

Procedure:

A common problem with Netscape 4.x involves a missing or misplaced security class. Try the following workaround, especially if you see the following line in the browser Java console:

```
Cannot find class  java/security/InvalidParameterException
```

- On Windows operating systems:

  From the *DB2 HTML Documentation CD*, copy the supplied *x*:program files\IBM\SQLLIB\doc\htmlcd*locale*\InvalidParameterException.class file to the java\classes\java\security\ directory relative to your Netscape browser installation, where *x* represents the CD-ROM drive letter and *locale* represents the name of the desired locale.

  Note: You may have to create the java\security\ subdirectory structure.

- On UNIX operating systems:

  From the *DB2 HTML Documentation CD*, copy the supplied /*cdrom*/program files/IBM/SQLLIB/doc/htmlcd/*locale*/InvalidParameterException.class file to the java/classes/java/security/ directory relative to your Netscape browser installation, where *cdrom* represents the mount point of the CD-ROM and *locale* represents the name of the desired locale.

  Note: You may have to create the java/security/ subdirectory structure.

If your Netscape browser still fails to display the search input window, try the following:

- Stop all instances of Netscape browsers to ensure that there is no Netscape code running on the machine. Then open a new instance of the Netscape browser and try to start the search again.
- Purge the browser's cache.
- Try a different version of Netscape, or a different browser.

Related tasks:

- "Searching the DB2 documentation" on page 124

Searching the DB2 documentation

You can search the library of DB2 documentation to locate information that you need. A pop-up search window opens when you click the search icon in the navigation toolbar of the DB2 Information Center (accessed from a browser). The search can take a minute to load, depending on the speed of your computer and network.

Prerequisites:

You need Netscape 6.1 or higher, or Microsoft's Internet Explorer 5 or higher. Ensure that your browser's Java support is enabled.

Restrictions:

The following restrictions apply when you use the documentation search:

- Search is not case sensitive.
- Boolean searches are not supported.
- Wildcard and partial searches are not supported. A search on *java** (or *java*) will only look for the literal string *java** (or *java*) and would not, for example, find *javadoc*.

Procedure:

To search the DB2 documentation:

1. In the navigation toolbar, click the **Search** icon.
2. In the top text entry field of the Search window, enter one or more terms (separated by a space) related to your area of interest and click **Search**. A list of topics ranked by accuracy displays in the **Results** field. The numerical ranking beside the hit provides an indication of the strength of the match (bigger numbers indicate stronger matches).

   Entering more terms increases the precision of your query while reducing the number of topics returned from your query.
3. In the **Results** list, click the title of the topic you want to read. The topic displays in the content frame of the DB2 Information Center.

Note: When you perform a search, the first (highest-ranking) result is automatically loaded into your browser frame. To view the contents of other search results, click on the result in the results list.

Related tasks:

- "Troubleshooting DB2 documentation search with Netscape 4.x" on page 123

Online DB2 troubleshooting information

With the release of DB2® UDB Version 8, there will no longer be a *Troubleshooting Guide*. The troubleshooting information once contained in this guide has been integrated into the DB2 publications. By doing this, we are able to deliver the most up-to-date information possible. To find information on the troubleshooting utilities and functions of DB2, access the DB2 Information Center from any of the tools.

Refer to the DB2 Online Support site if you are experiencing problems and want help finding possible causes and solutions. The support site contains a large, constantly updated database of DB2 publications, TechNotes, APAR (product problem) records, FixPaks, and other resources. You can use the support site to search through this knowledge base and find possible solutions to your problems.

Access the Online Support site at www.ibm.com/software/data/db2/udb/winos2unix/support, or by clicking the **Online Support** button in the DB2 Information Center. Frequently changing information, such as the listing of internal DB2 error codes, is now also available from this site.

Related concepts:

- "DB2 Information Center accessed from a browser" on page 128

Related tasks:

- "Finding product information by accessing the DB2 Information Center from the administration tools" on page 118

Accessibility

Accessibility features help users with physical disabilities, such as restricted mobility or limited vision, to use software products successfully. These are the major accessibility features in DB2® Universal Database Version 8:

- DB2 allows you to operate all features using the keyboard instead of the mouse. See "Keyboard Input and Navigation" on page 126.
- DB2 enables you customize the size and color of your fonts. See "Accessible Display" on page 126.
- DB2 allows you to receive either visual or audio alert cues. See "Alternative Alert Cues" on page 126.
- DB2 supports accessibility applications that use the Java™ Accessibility API. See "Compatibility with Assistive Technologies" on page 126.

- DB2 comes with documentation that is provided in an accessible format. See “Accessible Documentation”.

Keyboard Input and Navigation

Keyboard Input

You can operate the DB2 Tools using only the keyboard. You can use keys or key combinations to perform most operations that can also be done using a mouse.

Keyboard Focus

In UNIX-based systems, the position of the keyboard focus is highlighted, indicating which area of the window is active and where your keystrokes will have an effect.

Accessible Display

The DB2 Tools have features that enhance the user interface and improve accessibility for users with low vision. These accessibility enhancements include support for customizable font properties.

Font Settings

The DB2 Tools allow you to select the color, size, and font for the text in menus and dialog windows, using the Tools Settings notebook.

Non-dependence on Color

You do not need to distinguish between colors in order to use any of the functions in this product.

Alternative Alert Cues

You can specify whether you want to receive alerts through audio or visual cues, using the Tools Settings notebook.

Compatibility with Assistive Technologies

The DB2 Tools interface supports the Java Accessibility API enabling use by screen readers and other assistive technologies used by people with disabilities.

Accessible Documentation

Documentation for the DB2 family of products is available in HTML format. This allows you to view documentation according to the display preferences set in your browser. It also allows you to use screen readers and other assistive technologies.

DB2 tutorials

The DB2® tutorials help you learn about various aspects of DB2 Universal Database. The tutorials provide lessons with step-by-step instructions in the areas of developing applications, tuning SQL query performance, working with data warehouses, managing metadata, and developing Web services using DB2.

Before you begin:

Before you can access these tutorials using the links below, you must install the tutorials from the *DB2 HTML Documentation* CD-ROM.

If you do not want to install the tutorials, you can view the HTML versions of the tutorials directly from the *DB2 HTML Documentation CD*. PDF versions of these tutorials are also available on the *DB2 PDF Documentation CD*.

Some tutorial lessons use sample data or code. See each individual tutorial for a description of any prerequisites for its specific tasks.

DB2 Universal Database tutorials:

If you installed the tutorials from the *DB2 HTML Documentation* CD-ROM, you can click on a tutorial title in the following list to view that tutorial.

Business Intelligence Tutorial: Introduction to the Data Warehouse Center
Perform introductory data warehousing tasks using the Data Warehouse Center.

Business Intelligence Tutorial: Extended Lessons in Data Warehousing
Perform advanced data warehousing tasks using the Data Warehouse Center.

Development Center Tutorial for Video Online using Microsoft® Visual Basic
Build various components of an application using the Development Center Add-in for Microsoft Visual Basic.

Information Catalog Center Tutorial
Create and manage an information catalog to locate and use metadata using the Information Catalog Center.

Video Central for e-business Tutorial
Develop and deploy an advanced DB2 Web Services application using WebSphere® products.

Visual Explain Tutorial
Analyze, optimize, and tune SQL statements for better performance using Visual Explain.

DB2 Information Center accessed from a browser

The DB2® Information Center gives you access to all of the information you need to take full advantage of DB2 Universal Database™ and DB2 Connect™ in your business. The DB2 Information Center also documents major DB2 features and components including replication, data warehousing, the Information Catalog Center, Life Sciences Data Connect, and DB2 extenders.

The DB2 Information Center accessed from a browser has the following features if you view it in Netscape Navigator 6.1 or higher or Microsoft Internet Explorer 5. Some features require you to enable support for Java or JavaScript:

Regularly updated documentation
Keep your topics up-to-date by downloading updated HTML.

Search
Search all of the topics installed on your workstation by clicking **Search** in the navigation toolbar.

Integrated navigation tree
Locate any topic in the DB2 library from a single navigation tree. The navigation tree is organized by information type as follows:

- Tasks provide step-by-step instructions on how to complete a goal.
- Concepts provide an overview of a subject.
- Reference topics provide detailed information about a subject, including statement and command syntax, message help, requirements.

Master index
Access the information installed from the *DB2 HTML Documentation CD* from the master index. The index is organized in alphabetical order by index term.

Master glossary
The master glossary defines terms used in the DB2 Information Center. The glossary is organized in alphabetical order by glossary term.

Related tasks:

- "Finding topics by accessing the DB2 Information Center from a browser" on page 116
- "Finding product information by accessing the DB2 Information Center from the administration tools" on page 118
- "Updating the HTML documentation installed on your machine" on page 120

Notices

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 Canada Limited
Office of the Lab Director
8200 Warden Avenue
Markham, Ontario
L6G 1C7
CANADA

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, and have been used in at least one of the documents in the DB2 UDB documentation library.

ACF/VTAM
AISPO
AIX
AIXwindows
AnyNet
APPN
AS/400
BookManager
C Set++
C/370
CICS
Database 2
DataHub
DataJoiner
DataPropagator
DataRefresher
DB2
DB2 Connect
DB2 Extenders
DB2 OLAP Server
DB2 Universal Database
Distributed Relational Database Architecture
DRDA
eServer
Extended Services
FFST
First Failure Support Technology
IBM
IMS
IMS/ESA
iSeries
LAN Distance
MVS
MVS/ESA
MVS/XA
Net.Data
NetView
OS/390
OS/400
PowerPC
pSeries
QBIC
QMF
RACF
RISC System/6000
RS/6000
S/370
SP
SQL/400
SQL/DS
System/370
System/390
SystemView
Tivoli
VisualAge
VM/ESA
VSE/ESA
VTAM
WebExplorer
WebSphere
WIN-OS/2
z/OS
zSeries

The following terms are trademarks or registered trademarks of other companies and have been used in at least one of the documents in the DB2 UDB documentation library:

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

Intel and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. 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.

Index

J

L

M

N

O

P

R

S

Z

Contacting IBM

In the United States, call one of the following numbers to contact IBM:

- 1-800-237-5511 for customer service
- 1-888-426-4343 to learn about available service options
- 1-800-IBM-4YOU (426-4968) for DB2 marketing and sales

In Canada, call one of the following numbers to contact IBM:

- 1-800-IBM-SERV (1-800-426-7378) for customer service
- 1-800-465-9600 to learn about available service options
- 1-800-IBM-4YOU (1-800-426-4968) for DB2 marketing and sales

To locate an IBM office in your country or region, check IBM's Directory of Worldwide Contacts on the web at www.ibm.com/planetwide

Product information

Information regarding DB2 Universal Database products is available by telephone or by the World Wide Web at www.ibm.com/software/data/db2/udb

This site contains the latest information on the technical library, ordering books, client downloads, newsgroups, FixPaks, news, and links to web resources.

If you live in the U.S.A., then you can call one of the following numbers:

- 1-800-IBM-CALL (1-800-426-2255) to order products or to obtain general information.
- 1-800-879-2755 to order publications.

For information on how to contact IBM outside of the United States, go to the IBM Worldwide page at www.ibm.com/planetwide

Part Number: CT16GNA

Printed in U.S.A.

(1P) P/N: CT16GNA

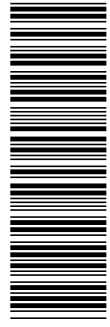

GC27-1122-00

Spine information:

IBM® DB2® Warehouse Manager

Warehouse Manager Installation Guide

Version 8