IBM® DB2 Universal Database™

Business Intelligence Tutorial: Introduction to the Data Warehouse Center

Version 8

IBM® DB2 Universal Database™

Business Intelligence Tutorial: Introduction to the Data Warehouse Center

Version 8

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 the tutorial

This tutorial provides an end-to-end guide for typical business intelligence tasks in the Data Warehouse Center. It is intended for database administrators who have never used the Data Warehouse Center before.

In this tutorial, you will learn how to use the DB2® Control Center and the Data Warehouse Center to create a warehouse database, move and transform source data, and write the data to the warehouse target database. This tutorial is designed for use with Windows NT®, Windows® 2000, Windows XP, Windows ME, and Windows 98.

The tutorial is available in HTML or PDF format. You can view the HTML version of the tutorial from the Data Warehouse Center or the Information Center. The PDF file is available on the DB2 Publications CD-ROM. The *Business Intelligence Tutorial: Extended Lessons in Data Warehousing* is available at http://www.ibm.com/software/data/bi/downloads.html.

This tutorial takes approximately 2 hours to complete.

What is data warehousing?

The systems that contain *operational data*, the data that runs the daily transactions of your business, contain information that is useful to business analysts. For example, analysts can use information about which products were sold in which regions at a specific time of year to look for anomalies or to project future sales. However, there are several problems if analysts access the operational data directly:

- They might not have the expertise to query the operational database. For example, querying IMS™ databases requires an application program that uses a specialized type of data manipulation language. In general, the programmers who have the expertise to query the operational database have a full-time job in maintaining the database and its applications.
- Performance is critical for many operational databases, such as databases for a bank. The system cannot handle users making ad hoc queries.
- The operational data generally is not in the best format for use by business analysts. For example, sales data that is summarized by product, region, and season is much more useful to analysts than raw data.

Data warehousing solves these problems. In *data warehousing*, you create stores of *informational data*, data that is extracted from the operational data and then transformed for decision making. For example, a data warehousing tool might

copy all of the sales data from the operational database, perform calculations to summarize the data, and write the summarized data to a database that is separate from the operational data. Users can query the separate database (the *warehouse*) without impacting the operational databases.

What is covered in this tutorial?

DB2 Universal Database™ offers the Data Warehouse Center, a DB2 component that automates warehouse processing. You can use the Data Warehouse Center to define what data to include in the warehouse. Then, you can use the Data Warehouse Center to automatically schedule when the data in the warehouse is refreshed.

This tutorial covers the most common tasks that are required to set up a warehouse.

In this tutorial, you will:

- Define a *subject area* that identifies and groups the processes that you will create for the tutorial.
- Explore the source data (which is the operational data) and define warehouse sources. *Warehouse sources* identify the source data that you want to use in your warehouse.
- Create a database to use as the warehouse and define *warehouse targets*, which identify the target data to include in your warehouse.
- Specify how to move and transform the source data into its format for the warehouse database. You will define a *process*, which contains the series of movement and transformation steps required to produce a target table in the warehouse from one or more source tables, views, or files. You will then define the process in *steps*, each of which defines one operation in the movement and transformation process. Then you will test the steps that you defined and schedule them to run automatically.

The following map shows the content of the tutorial and the order of lessons.

Conventions that are used in this tutorial

This tutorial uses typographical conventions in the text to help you distinguish between the names of controls and text that you type. For example:

- Menu items are in boldface font:

  Click **Menu —> Menu choice**.
- The names of fields, check boxes, and other controls are also in boldface font:

  Type text in the **Field name** field.

- Text that you type is in monospaced font on a new line:

  ```
  This is the text that you type.
  ```

Business problem

You are a database administrator for a company that is called TBC: The Beverage Company. The company manufactures beverages for sale to other businesses. The financial department wants to track, analyze, and forecast the sales revenue across geographic regions on a periodic basis for all products sold. You have already set up standard queries of the sales data. However, these queries add to the load on your operational database. Also, users sometimes ask for additional ad hoc queries of the data, based on the results of the standard queries.

Your company has decided to create a data warehouse for the sales data. A *data warehouse* is a database that contains data that has been cleansed and transformed into an informational format. Your task is to create this data warehouse.

Before you begin

This section explains the tasks you must complete before you can begin the tutorial.

Prerequisites

Before you begin this tutorial, you must install the products that are covered in the tutorial lessons:

- DB2 server. The DB2 server is included in the typical installation for DB2 Universal Database.
- DB2 Control Center and the Data Warehouse Center administrative interface. You can install the Data Warehouse Center administrative interface on the following operating systems: Windows NT, Windows 2000, Windows XP, Windows ME, and Windows 98, and AIX®.
- Warehouse server. The warehouse server is installed when you select Data warehousing when you install DB2 Enterprise Server Edition.

  For more information about installing DB2 Universal Database and the warehouse server, see *DB2 Universal Database Quick Beginnings*.
- You need sample data to use with the tutorial. The tutorial uses the DB2 Data Warehousing sample data.

  The Data Warehousing sample data is installed on Windows NT, Windows 2000, Windows XP, and AIX, when you install the Data Warehouse Center administrative interface. It must either be installed on the same workstation as the warehouse server, or the remote node for the sample databases must be cataloged on the warehouse server workstation. This tutorial focuses on

completing warehousing tasks using the Data Warehouse Center administrative client on a Windows system.

Creating the sample databases

This tutorial contains several references to sample data in the default directory structure `x:\Program Files\ibm\sqllib`, where x is the drive under which you installed DB2. You could have installed DB2 in a different directory structure. For example, `x:\ sqllib` directory.

You must create the sample databases after you install the files for the sample.

To create the databases:

1. Click **Start —> Programs —> IBM DB2 —> Set-up Tools —> First Steps**.

   The IBM DB2 First Steps window opens.
2. Click **Create Sample Databases**. If this option is not available, the sample databases have already been created.

   The First Steps - Create Sample Databases window opens.
3. Select **Data warehousing sample**. If this option is not available, the sample databases have already been created. The Data Warehouse Center User Name and Password window opens.
4. In the **User name** field, type the DB2 user ID that you will use to access the sample.
5. In the **Password** field, type the password that corresponds to the user ID. Make note of the user ID and password that you use to create the samples. You will need them in a later lesson.
6. Click **OK**. The Data Warehouse Center User Name and Password window closes.

   DB2 starts to create the sample database. A progress window opens. It can take a while for the databases to be created.
7. When the databases have been created, click **OK**.
8. Click **Exit First Steps** to close the IBM DB2 First Steps window.

The following databases are created when you create the data warehousing sample:

DWCTBC
: Contains the operational source tables that are required for the tutorial.

TBC_MD
: Contains metadata for the Data Warehouse Center objects in the sample.

Connecting to the sample databases

Before you begin the tutorial, verify that you can connect to the sample databases:

1. Click **Start —> Programs —> IBM DB2 —> General Administration Tools —> Control Center** to start the DB2 Control Center.
2. Expand the tree until you see one of the sample databases: DWCTBC or TBC_MD.
3. Right-click the name of the database and click **Connect**.

   The Connect window opens.
4. In the **User ID** field, type the user ID that you used to create the sample.
5. In the **Password** field, type the password that you used to create the sample.
6. Click **OK**.

   The DB2 Control Center connects to the database. If the DB2 Control Center is not able to establish a connection, you will see an error message. If you receive an error message, see the *DB2 Universal Database Message Reference*.

Chapter 1. Defining warehouse security

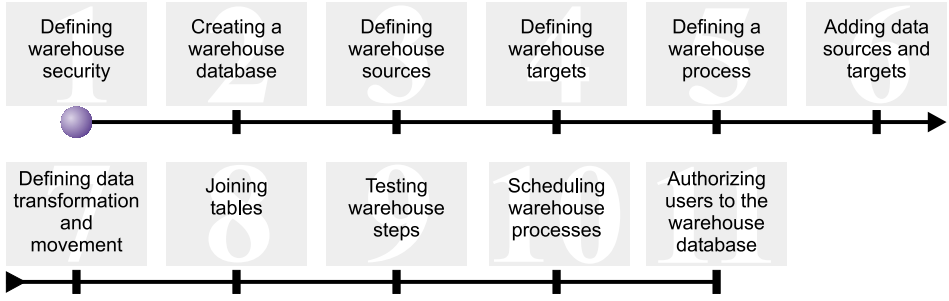

Before you begin defining the TBC data warehouse, you must set up security for the warehouse. In this lesson, you will learn how to perform the following tasks:

- Specify the warehouse control database
- Start the Data Warehouse Center
- Define a warehouse user
- Define the warehouse group

This lesson takes approximately 15 minutes to complete.

After you set up security for the warehouse, you can grant access to the warehouse objects that you create in later lessons.

How security works in the Data Warehouse Center

The first level of security is the logon user ID that is in use when you open the Data Warehouse Center. Although you log on to the DB2 Control Center, the Data Warehouse Center verifies that you are authorized to open the Data Warehouse Center administrative interface by comparing your user ID to entries in the warehouse control database. The *warehouse control database* contains the control tables that are required to store Data Warehouse Center metadata. You initialize the control tables for this database when you install the warehouse server as part of DB2 Universal Database or use the Data Warehouse Center Control Database Management window. During initialization, you specify the ODBC name of the warehouse control database, a valid DB2 user ID, and a password. The Data Warehouse Center authorizes this user ID and password to update the warehouse control database. In the Data Warehouse Center, this user ID is defined as the *default warehouse user*.

The default warehouse user requires a different type of database and operating system authorization for each operating system that the warehouse control database supports.

Access to Data Warehouse center objects and functions

The default warehouse user is authorized to access all Data Warehouse Center objects and perform all Data Warehouse Center functions. However, you probably want to restrict access to certain objects within the Data Warehouse Center and the tasks that users can perform on the objects. For example, warehouse sources and warehouse targets contain the user IDs and passwords for their corresponding databases. You might want to restrict access to the warehouse sources and warehouse targets that contain sensitive data, such as personnel data.

To provide this level of security, the Data Warehouse Center provides a security system that is separate from the database and operating system security. To implement Data Warehouse Center security, you define warehouse users and warehouse groups. A *warehouse group* is a named grouping of warehouse users and their authorization to perform functions. Warehouse users and warehouse groups do not have to match the database users and database groups that are defined for the warehouse control database.

For example, you might define a warehouse user that corresponds to someone who uses the Data Warehouse Center. You might then define a warehouse group that is authorized to access certain warehouse sources, and add the new user to the new warehouse group. The new user is authorized to access the warehouse sources that are included in the group.

You can give users different levels of authorization. You can include any of the different types of authorization in a warehouse group. You can also include a warehouse user in more than one warehouse group. The combination of the groups to which a user belongs is the user's overall authorization.

Specifying the warehouse control database

When you install the Data Warehouse Center, if you specify to prepare the metadata during the installation, the installation process registers the default warehouse control database as the active warehouse control database. However, you must use the TBC_MD database in the sample as the warehouse control database so that you can use the sample metadata. To make TBC_MD the active database, you must reinitialize it.

To specify the warehouse control database, TBC_MD:

1. Click **Start —> Programs —> IBM DB2 —> Set-up Tools —> Warehouse Control Database Management.**

The Data Warehouse Center - Control Database Management window opens.

2. In the **New control database** field, type:

   `TBC_MD`

3. In the **Schema** field, use the default schema of IWH.
4. In the **User ID** field, type the user ID that is required to access the database.
5. In the **Password** field, type the name of the password for the user ID.
6. In the **Verify password** field, type the password again.
7. Click **OK**.

   The window remains open. The Messages field displays messages that indicate the status of the creation and migration process.

8. After the process is complete, click **Close** to close the window. TBC_MD is now the active warehouse control database.

Starting the Data Warehouse Center

In this exercise, you will start the Data Warehouse Center and log on as the default warehouse user. When you log on, you will use the TBC_MD warehouse control database. The default warehouse user for TBC_MD is the user ID that you specified when you created the data warehousing sample databases.

TBC_MD must be a local database or a cataloged remote database on the workstation that contains the warehouse server. It must also be a local or cataloged remote database on the workstation that contains the Data Warehouse Center administrative client.

To start the Data Warehouse Center:

1. Click **Start —> Programs —> IBM DB2 —> Business Intelligence Tools —> Data Warehouse Center**.

   The Data Warehouse Center Logon window opens.

2. Click **Advanced**.

   The Advanced window opens.

3. In the **Control database** field, type the name of the warehouse control database that is included in the sample:

   `TBC_MD`

4. In the **Server host name** field, type the TCP/IP hostname for the workstation where the warehouse server is installed.

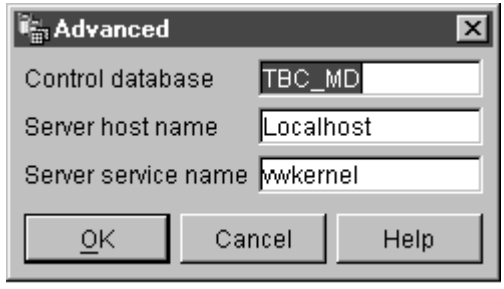

5. Click **OK**.

   The Advanced window closes.

   The next time that you log on, the Data Warehouse Center will use the settings that you specified in the Advanced Logon window.
6. In the **User ID** field of the Data Warehouse Center Logon window, type the default warehouse user ID.
7. In the **Password** field, type the password for the user ID.

8. Click **OK**.

   The Data Warehouse Center Logon window closes.
9. Close the Data Warehouse Center Launchpad window if it is open.

Defining a warehouse user

In this exercise, you will define a new user to the Data Warehouse Center.

The Data Warehouse Center controls access with user IDs. When a user logs on, the user ID is compared to the warehouse users that are defined in the Data Warehouse Center to determine whether the user is authorized to access the Data Warehouse Center. You can authorize additional users to access the Data Warehouse Center by defining new warehouse users.

The user ID for the new user does not require authorization to the operating system or the warehouse control database. The user ID exists only within the Data Warehouse Center.

To define a warehouse user:

1. In the left pane of the main Data Warehouse Center window, expand the **Administration** tree.

2. Expand the **Warehouse Users and Groups** tree.
3. Right-click the **Warehouse Users** folder, and click **Define**.

   The Define Warehouse User notebook opens.
4. In the **Name** field, type the name of the user:

   ```
   Tutorial User
   ```

   The name identifies the user ID within the Data Warehouse Center. This name can be up to 80 characters, including spaces.
5. In the **Administrator** field, type your name as the contact for this user.
6. In the **Description** field, type a short description of the user:

   ```
   This is a user that I created for the tutorial.
   ```
7. In the **User ID** field, type the new user ID:

   ```
   tutuser
   ```

   The user ID must be no longer than 60 characters and cannot contain spaces, dashes, or special characters (such as @, #, $, %, >, +, =). It can contain the underscore character.

   Specifying a unique user ID:

   To determine if a user ID and password is unique:

   a. From the main Data Warehouse Center window, expand the **Administration** tree.

   b. Click on the **Warehouse Users** folder. All of the user IDs for the data warehouse appear in the right panel. Any ID that does not appear in the right panel is a unique ID.
8. In the **Password** field, type the password:

   ```
   password
   ```

   Passwords must be a minimum of six characters and cannot contain spaces, dashes, or special characters.
9. In the **Verify password** field, type your password again.
10. Click **OK** to save the warehouse user and close the notebook.

Defining the warehouse group

In this exercise, you will define a warehouse group to authorize the tutorial user, tutuser, that you just created to perform tasks.

To define the warehouse group:

1. In the main Data Warehouse Center window, right-click the **Warehouse Groups** folder, and click **Define**.

The Define Warehouse Group notebook opens.

2. In the **Name** field, type the name for the new group:

   ```
   Tutorial Warehouse Group
   ```

3. In the **Administrator** field, type your name as the contact for this new group.
4. In the **Description** field, type a short description of the new group:

   ```
   This is the warehouse group for the tutorial.
   ```

5. Click >> to move all of the privileges from the **Available privileges** to the **Selected privileges** list. This selects all of the privileges for your group.

   Your group now has the following privileges:

   Administration
   Users in the warehouse group can define and change warehouse users and warehouse groups, change Data Warehouse Center properties, import metadata, and define which warehouse groups have access to objects when they are created.

   Operations
   Users in the warehouse group can monitor the status of scheduled processing.

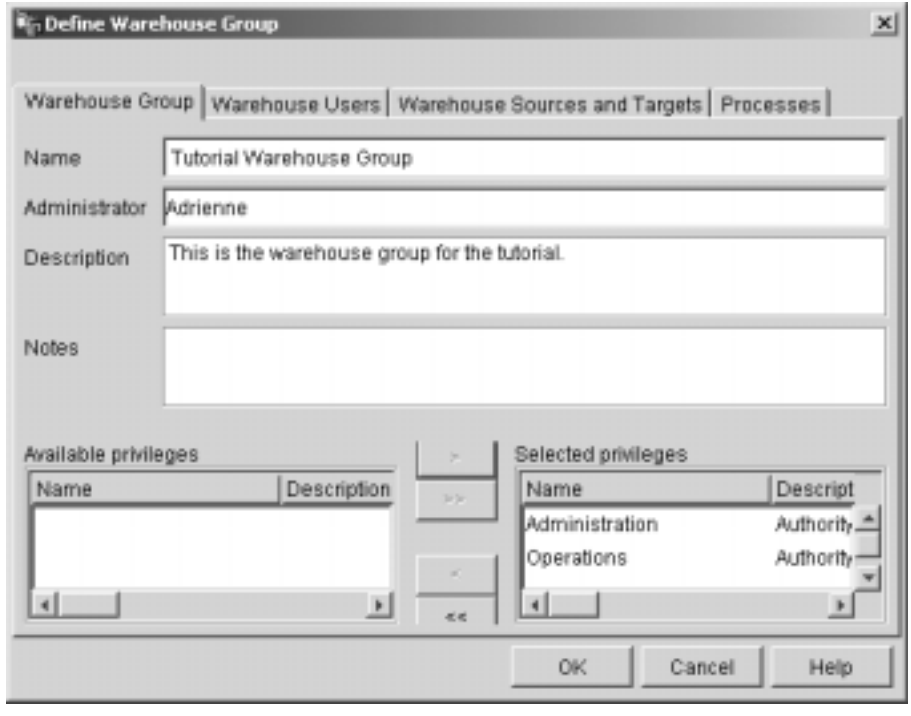

6. Click the **Warehouse Users** tab.

7. From the **Available warehouse users** list, select the **Tutorial User**.
8. Click >.

   The Tutorial User moves to the **Selected warehouse users** list.

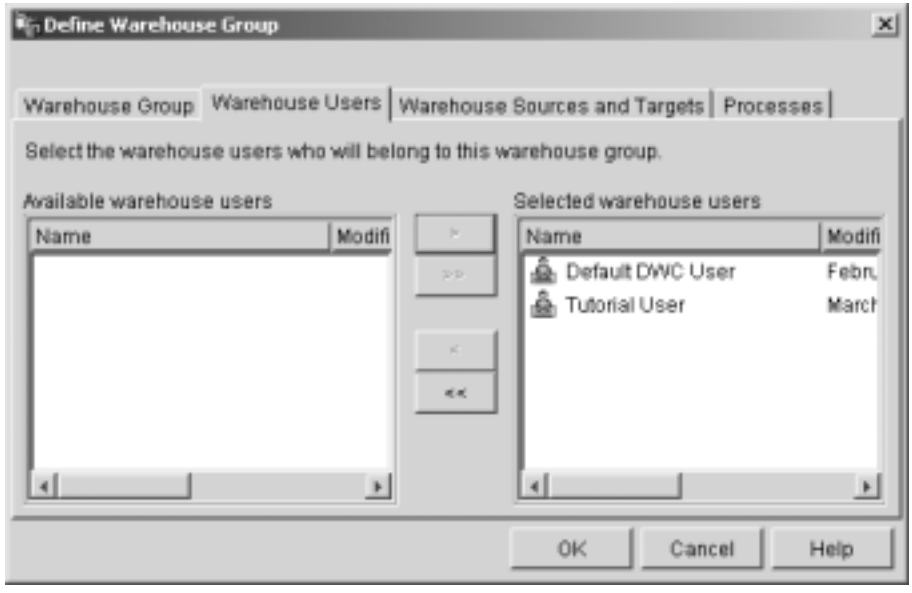

   The user is now part of the warehouse group.

   Skip the Warehouse Sources and Targets page and the Processes page. You will create these objects in subsequent lessons. You will authorize the warehouse group to access objects as you create them.
9. Click **OK** to save the warehouse user group and close the notebook.

What you just did

In this lesson, you:

- Started the Data Warehouse Center interface and logged on as the default user.
- Specified a warehouse control database.
- Defined a warehouse user.
- Defined a warehouse group.

In subsequent lessons, you will authorize the warehouse group to access the objects that you define.

Chapter 2. Creating a warehouse database

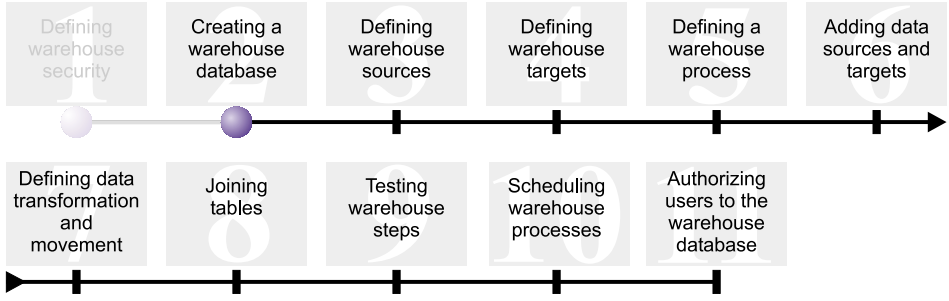

In this lesson, you will create the database for your warehouse. As part of DB2 First Steps, DB2 created the DWCTBC database, which contains the source data for this tutorial. The database that you create in this lesson is where the source data that is transformed for the warehouse will be stored.

This lesson takes approximately 10 minutes to complete.

Creating a database

This exercise shows you how to create the TUTWHS database for your warehouse. TUTWHS is the warehouse database.

To create the database:

1. Click **Start —> Programs —> IBM DB2 —> Command Line Tools —> Command Line Processor**. The DB2 CLP window opens.
2. Enter the following command at the prompt:

   ```
   create db TUTWHS
   ```

   When the database has been created, you will receive a message saying that the CREATE DATABASE command completed successfully.

Registering a database with ODBC

You can use the Configuration Assistant, the DB2 Command Line Processor, or the ODBC Data Source Administrator to register a database with ODBC. In this exercise, you will use the Configuration Assistant to register the TUTWHS database with ODBC.

To register the TUTWHS database with ODBC:

1. Start the Configuration Assistant by clicking **Start —> Programs —> IBM DB2 —> Set-up Tools —> Configuration Assistant**.

   The Configuration Assistant window opens.
2. Click **View —> Advanced View**.

   The Configuration Assistant changes to the Advanced Configuration Assistant view.
3. Click the **Data Sources** tab.
4. Click **Selected —> Add Data Source**.

   The Add Data Source window opens.
5. Select **System data source**. *System data source* means that the data is available to all users on the system.
6. Select TUTWHS in the **Database alias** list.
7. Type TUTWHS in the **Data source name** field. Use the default values for the rest of this window.
8. Click **OK**.

   The TUTWHS database is registered with ODBC and the window closes.
9. Close the Configuration Assistant.

Verifying that the DWCTBC and TBC_MD databases are registered with ODBC

To verify that the DWCTBC and TBC_MD databases are registered with ODBC:

1. Start the Configuration Assistant by clicking **Start —> Programs —> IBM DB2 —> Set-up Tools —> Configuration Assistant**.

   The Configuration Assistant window opens.
2. Verify that DWCTBC and TBC_MD appear in the Alias column.
3. If DWCTBC and TBC_MD appear in the Alias column, verify that the value in the ODBC Registered column is Yes. You might have to scroll to the right to see the ODBC Registered column.
4. If they are not registered, follow the procedure in “Registering a database with ODBC” on page 9 to register the databases with ODBC.

Connecting to the target database

Before you use the database that you defined, you must verify that you can connect to the database.

To connect to the database:

1. From the DB2 Control Center, expand the tree until you see the Databases folder.
2. Right-click the **Databases** folder, and click **Refresh**.

3. Right-click the TUTWHS of the database and click **Connect**.

   The Connect window opens.

4. Type the user ID and password that you used to log on to the DB2 Control Center.
5. Click **OK**.

   The DB2 Control Center connects to the database. If the connection is not successful, you will receive an error message.

What you just did

In this lesson, you:

- Created the TUTWHS database to contain the data for the warehouse.
- Verified that the DWCTBC and TBC_MD databases are registered with ODBC.
- Verified that you can connect to the database.

In the next lesson, you will view the source data that you will later transform and store in the database that you just created.

Chapter 3. Defining warehouse sources

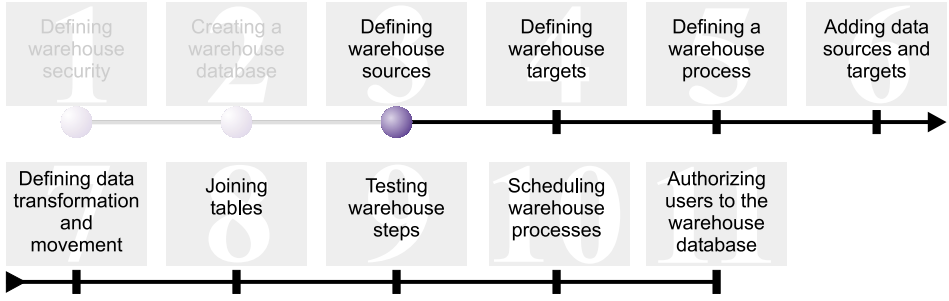

When you design a warehouse, you gather information about the operational data to use as input to the warehouse and the requirements for the warehouse data. The database administrator who is responsible for the operational data is a good source for information about the operational data. The business users who will be making business decisions based on the data in the warehouse are a good source for information about the requirements of the warehouse.

In this lesson you will:

- Specify a user ID and password for the sample sources
- View table and file data
- Define two warehouse sources:

  Tutorial Relational Source
  Corresponds to the GEOGRAPHIES source table in the DWCTBC database.

  Tutorial File Source
  Corresponds to the demographics file, which you will load into the warehouse database in a later lesson.

Warehouse sources are logical definitions of the tables and files. The warehouse sources that you define in this lesson will provide data to the Market dimension table mentioned in the “Business problem” on page vii. The Data Warehouse Center uses the specifications in the warehouse sources to access and select the data.

If you are using source databases that are remote to the warehouse server, you must register the databases on the workstation or AIX system that contains the warehouse server.

This lesson takes approximately 20 minutes to complete.

Specifying a user ID and password for the TBC sample sources

The sample warehouse sources do not have a user ID and password associated with them. You must add a user ID and password before you can work with these sources.

To specify a user ID and password for the TBC sample sources:

1. From the Data Warehouse Center window, expand the **Warehouse Sources** tree.
2. Right-click on **TBC Sample Sources**, and click **Properties**.

   The Properties - TBC Sample Sources window opens.
3. Click the **Database** tab.
4. In the **User ID** field, type the user ID that you used when you created the warehouse database, TUTWHS, in the previous lesson.
5. In the **Password** field, type the password for the user ID.
6. In the **Verify password** field, type the password again.
7. Click **OK**.

Viewing data

When you define warehouse sources, you must first decide which sources to use. You can view the data in the tables and files that you might want to use before you make your final decision. Viewing data in target tables and files after the source data is moved and transformed can help you to verify that the resulting data is what you expected.

The source data that you use consists of DB2 Universal Database tables and a text file. Some other typical types of source data are non-DB2 relational tables, MVS™ data sets, and Microsoft® Excel spreadsheets.

The source data that you use in this tutorial is preselected for you, but the following exercises will show you how to view table and file data for your own warehouse sources and targets. As you view the data, look for relationships among the data and consider what information might be of the most interest to users.

Viewing table data

In this exercise, you will use the DB2 Control Center to view rows in a table.

To view the data in the table:

1. In the Control Center window, expand the objects in the DWCTBC database until you see the Tables folder.

2. Click the **Tables** folder. In the right pane, you see all the tables for the database.
3. Find the GEOGRAPHIES table. Right-click the **GEOGRAPHIES** table, and click **Sample Contents**. The Sample Contents window opens.

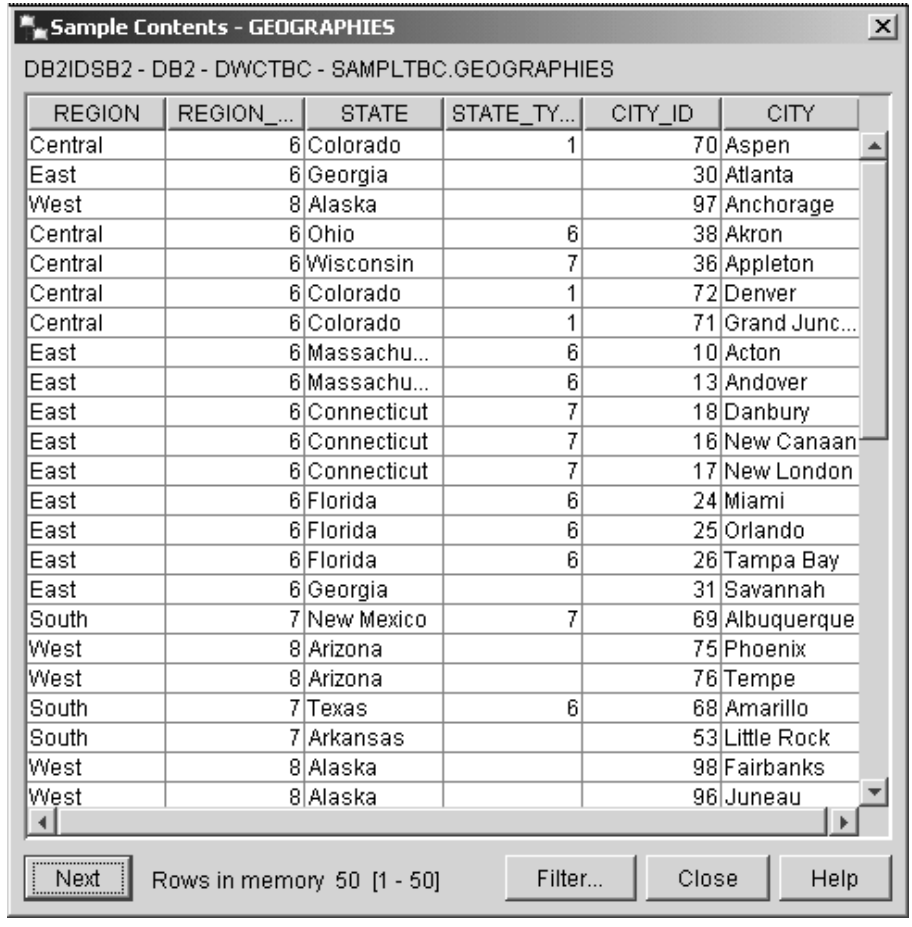

REGION	REGION_...	STATE	STATE_TY...	CITY_ID	CITY
Central	6	Colorado	1	70	Aspen
East	6	Georgia		30	Atlanta
West	8	Alaska		97	Anchorage
Central	6	Ohio	6	38	Akron
Central	6	Wisconsin	7	36	Appleton
Central	6	Colorado	1	72	Denver
Central	6	Colorado	1	71	Grand Junc...
East	6	Massachu...	6	10	Acton
East	6	Massachu...	6	13	Andover
East	6	Connecticut	7	18	Danbury
East	6	Connecticut	7	16	New Canaan
East	6	Connecticut	7	17	New London
East	6	Florida	6	24	Miami
East	6	Florida	6	25	Orlando
East	6	Florida	6	26	Tampa Bay
East	6	Georgia		31	Savannah
South	7	New Mexico	7	69	Albuquerque
West	8	Arizona		75	Phoenix
West	8	Arizona		76	Tempe
South	7	Texas	6	68	Amarillo
South	7	Arkansas		53	Little Rock
West	8	Alaska		98	Fairbanks
West	8	Alaska		96	Juneau

The column names are displayed at the top of the window. You might need to scroll to the right to see all the columns, and scroll down to see all the rows.

4. Click **Close**.

Viewing file data

In this exercise, you will use Microsoft® Notepad to view the contents of the demographics.txt file.

To view the file:

1. Click **Start —> Programs —> Accessories —> Notepad** to open Microsoft Notepad.
2. Click **File —> Open**.
3. Use the Open window to locate the file. For example, it might be located in:

   x:\program files\ibm\sqllib\samples\db2sampl\dwc\demographics.txt

   Where *x* is the drive on which you installed DB2.
4. Select the demographics.txt file, and click **Open** to view its contents.

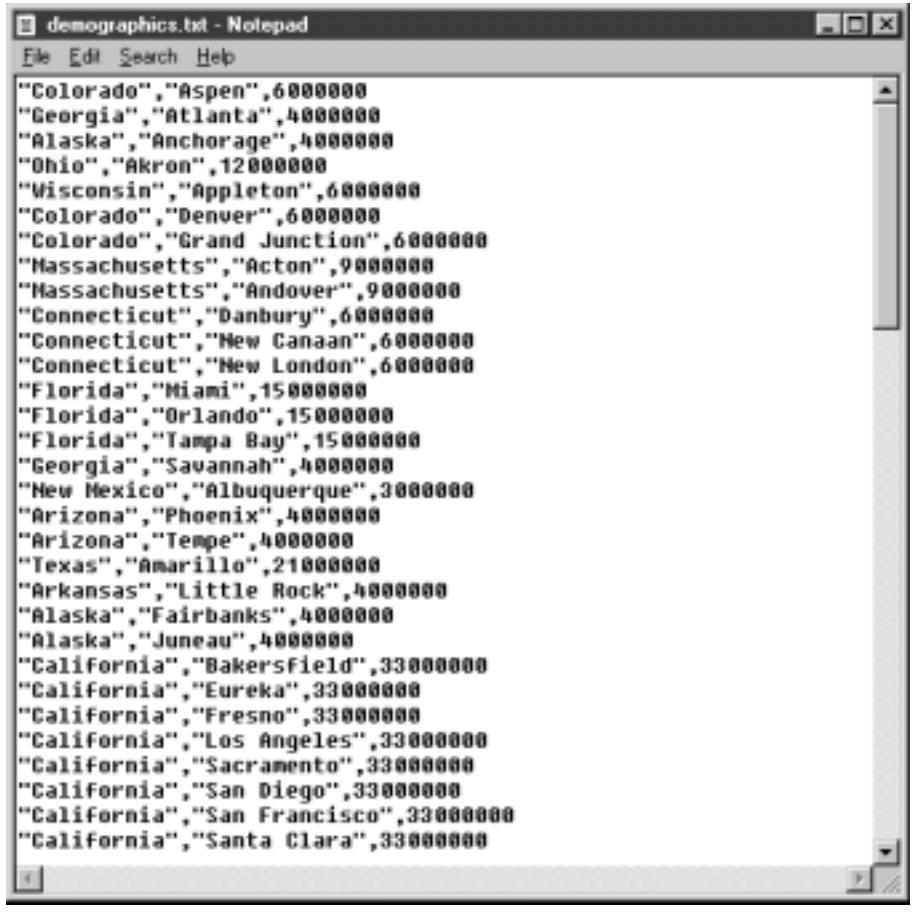

   Note that the file is comma-delimited. You will specify the delimiter for the demographics.txt file in a later lesson.
5. Close Notepad.

Defining a relational warehouse source

To track geographical sales data for TBC, you will define a relational warehouse source called the Tutorial Relational Source. It corresponds to the GEOGRAPHIES relational table that is provided in the DWCTBC database.

You will perform the following tasks when you define the Tutorial Relational Source:

- Open the Define Warehouse Source notebook
- Specify information about the source
- Specify information about the source database
- Import tables into the warehouse source
- Define security for a warehouse source

Opening the Define Warehouse Source notebook

This exercise shows you how to open the Define Warehouse Source notebook for a relational warehouse source called Tutorial Relational Source. This relational warehouse source will be used later in the tutorial.

To open the Define Warehouse Source notebook for the Tutorial Relational Source:

1. From the Data Warehouse Center window, right-click the **Warehouse Sources** folder.
2. Click **Define —> DB2 Family**.

   The Define Warehouse Source notebook opens.

Leave the Define Warehouse Source notebook open for the next task.

Specifying information about a warehouse source

This exercise shows you how to specify general information about the warehouse source that you are defining.

To specify information about a warehouse source:

1. On the Warehouse Source page of the Define Warehouse Source notebook, type the following business name for the warehouse source in the **Name** field:

   ```
   Tutorial Relational Source
   ```

   A *business name* is a descriptive name that users will understand. You will use this name to refer to your warehouse source throughout the Data Warehouse Center.
2. In the **Administrator** field, type your name as the contact for the warehouse source.
3. In the **Description** field, type a short description of the data:

Relational data for the TBC company

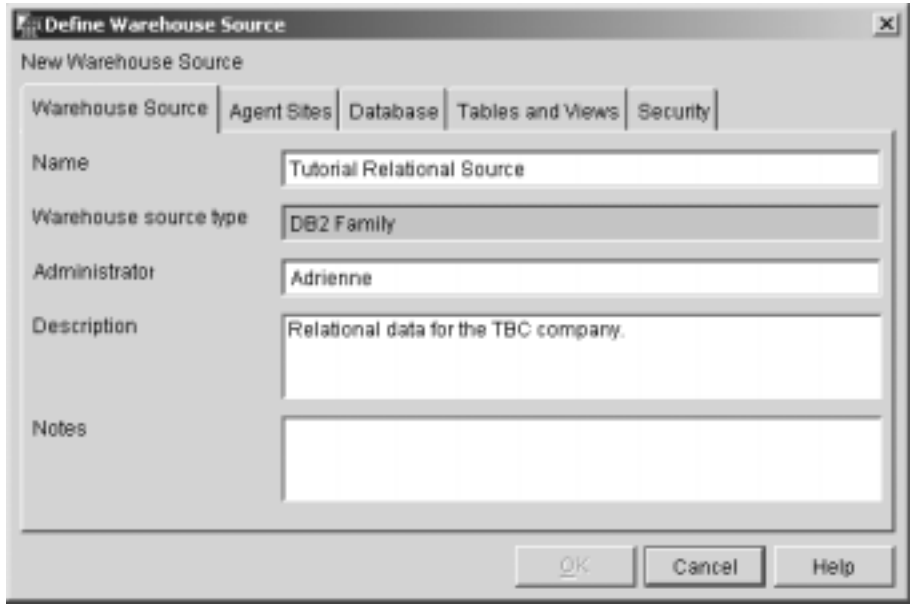

Leave the Define Warehouse Source notebook open for the next task.

Specifying information about the source database

In addition to specifying information about the warehouse source, you must also specify information about the physical source database.

To specify information about the source database:

1. Click the **Database** tab.
2. In the **Database name** field, specify DWCTBC as the name of the physical database.
3. If you are accessing DWCTBC from a workstation that is remote to the warehouse server, in the **System name** field specify the name of the system where DWCTBC resides. If you are accessing DWCTBC locally, skip this field.
4. In the **User ID** field, type a user ID that has access to the database.

   Use the user ID that you specified when you created the sample databases.
5. In the **Password** field, type the password for the user ID.
6. In the **Verify password** field, type the password again.

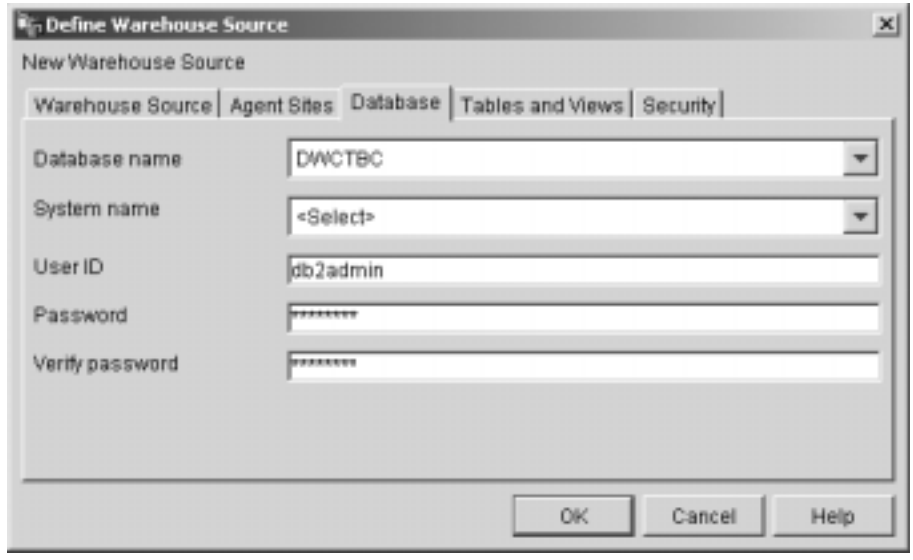

Leave the Define Warehouse Source notebook open for the next task.

Importing tables into your warehouse source

TBC wants to track the geographical sales data that is in the GEOGRAPHIES table. Because the GEOGRAPHIES table exists in a DB2 database, you can import the table definitions from DB2 instead of defining them manually. This exercise will show you how to import tables and views into a warehouse source.

1. In the Define Warehouse Source notebook, click the **Tables and Views** tab.
2. Expand the **Tables** tree.

   The Filter window opens.

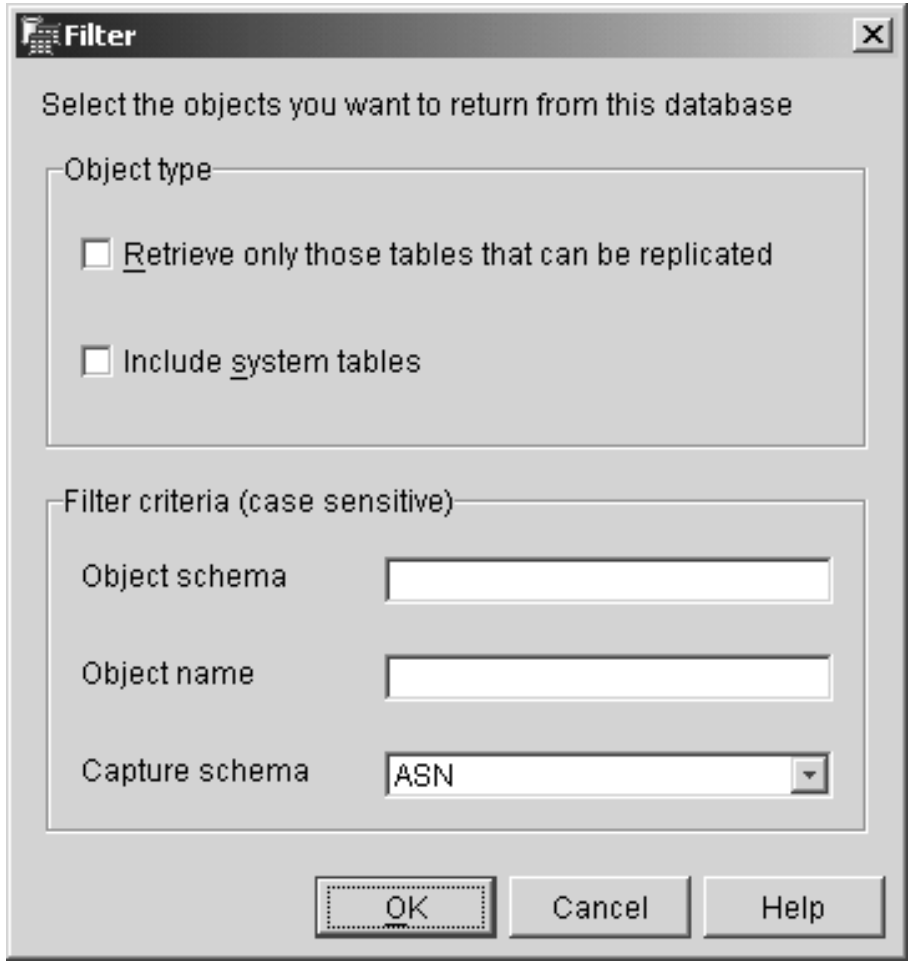

3. Click **OK**.

   After the import process finishes, the Data Warehouse Center lists the imported tables in the **Available tables and views** list.
4. From the **Available tables and views** list, select the **SAMPLTBC.GEOGRAPHIES** table.
5. Click > to move the SAMPLTBC.GEOGRAPHIES table to the **Selected tables and views** list. The table definitions in the Selected tables and views list are imported.

Leave the notebook open for the next task.

Defining security for a warehouse source

This exercise shows you how to define security for a warehouse source.

To define security for a warehouse source:

1. In the Define Warehouse Source notebook, click the **Security** tab.
2. Click the **Tutorial Warehouse Group**.
3. Click >.

   Adding the source to the Selected warehouse groups list authorizes the users in the group (in this case, you) to define tables and views for the source.
4. Click **OK** to save your changes.

   The Define Warehouse Sources notebook closes automatically when the changes are saved.

Defining a file source

In addition to the geographical information, TBC also wants to track demographic sales information. The company's demographic information is in the form of a text file called demographics.txt. In the following exercises, you will complete the tasks that are required to define a file warehouse source in the Data Warehouse Center. The file source is called Tutorial File Source, and it corresponds to the demographics.txt file that is provided with the data warehousing sample. For this tutorial, you will define only one file in the warehouse source, but you can define multiple files in a warehouse source.

The tasks that you will complete to define the file source are:

- Open the Define Warehouse Source notebook for a file source
- Specify information about a warehouse file source
- Specify information about the physical source file
- Define security for a warehouse file source

Opening the Define Warehouse Source notebook for a file source

To open the Define Warehouse Source notebook for a file source:

1. Right-click the **Warehouse Sources** folder.
2. Click **Define —> Flat File —> Local files**.

   The Define Warehouse Source notebook opens.

   The source type is Local files because the file that is used in this exercise was installed on your workstation with the tutorial.

Leave the Define Warehouse Source notebook open for the next task.

Specifying information about a warehouse file source

To specify information about a warehouse file source:

1. In the **Name** field, type the business name for the warehouse source:

   ```
   Tutorial file source
   ```

2. In the **Administrator** field, type your name as the contact for the warehouse source.
3. In the **Description** field, type a short description of the data:

   ```
   File data for the TBC company
   ```

Leave the Define Warehouse Source notebook open for the next task.

Specifying information about the physical source file

When you define a file source to the Data Warehouse Center, you must provide information about the file, such as its location. This exercise shows you how to specify information about the source file. You will complete the following tasks when you specify information about the source file:

- Open the Define Warehouse Source File notebook
- Specify information about the source file
- Specify parameters for the source file
- Specify information for the columns that are derived from the fields in the source file

Opening the Define Warehouse Source File notebook

The Define Warehouse Source File notebook is used to define the attributes of a source file that is being defined as a warehouse source. This exercise shows you how to open the Define Warehouse Source File notebook.

To open the Define Warehouse Source File notebook:

1. From the Define Warehouse Source notebook, click the **Files** tab.
2. Right-click in the blank area of the **Files** page, and click **Define**.

   The Define Warehouse Source File notebook opens.

Leave the Define Warehouse Source File notebook open for the next task.

Specifying information about the source file

After you open the Define Warehouse Source File notebook, you can specify information about the source file:

1. In the **File name** field, type the full path and name of the file. For example, if you installed DB2 in the default location, type:

   ```
   x:\Program Files\ibm\sqllib\samples\db2sampl\dwc\demographics.txt
   ```

   Where `x` is the drive where you installed DB2.
2. In the **Description** field, type a short description of the file:

```
Demographics data for sales regions.
```

3. In the **Business name** field, type:

   ```
   Demographics Data
   ```

   This is the business name for this file source object in the Data Warehouse Center. This metadata can also be published to an information catalog.

Leave the Define Warehouse Source File notebook open for the next task.

Specifying parameters for the source file

After you specify information about the source file, you can specify parameters for the file. The parameters for the file determine how the file data maps to the table structure. For example, you can specify that the first row of the source file contains the column headings. When the file is loaded into the table, the data in the first row is used as the column headings. This exercise shows you how to specify parameters for the source file.

To specify parameters for the source file:

1. Click the **Parameters** tab.

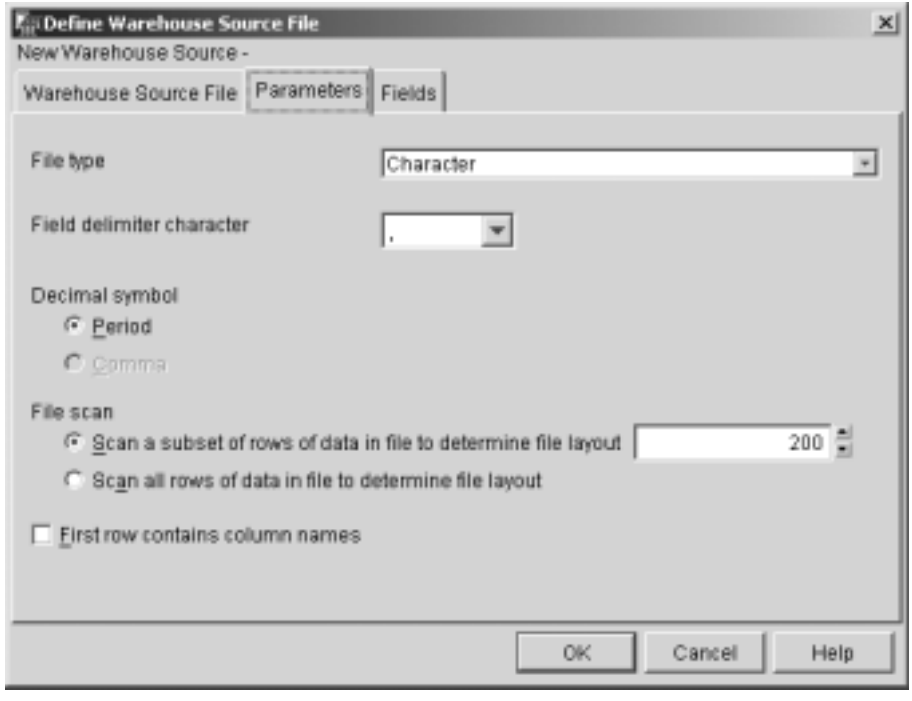

2. Verify that **Character** is selected in the **File type** list.

3. Verify that the comma is selected in the **Field delimiter character** field. The comma is selected as the field delimiter character because the file is a comma-delimited file.
4. Verify that the **First row contains column names** check box is cleared.

   The file does not contain column names.

Leave the Define Warehouse Source File notebook open for the next task.

Specifying information for the columns that are derived from the fields in a source file

The Data Warehouse Center reads the file that you specified on the Warehouse Source File page. It defines columns based on the fields in the file, and displays the column definitions in the **Fields** list. The sample data is displayed in the **File** preview area. You can scroll to see all the sample data.

To define information for the columns that are derived from the fields in a source file:

1. Click the **Fields** tab.
2. Click the default column name, **COL001**, and replace it with the following column name:

   `STATE`
3. Repeat step 2 to rename the rest of the columns. Rename **COL002** as `CITY` and **COL003** as `POPULATION`.

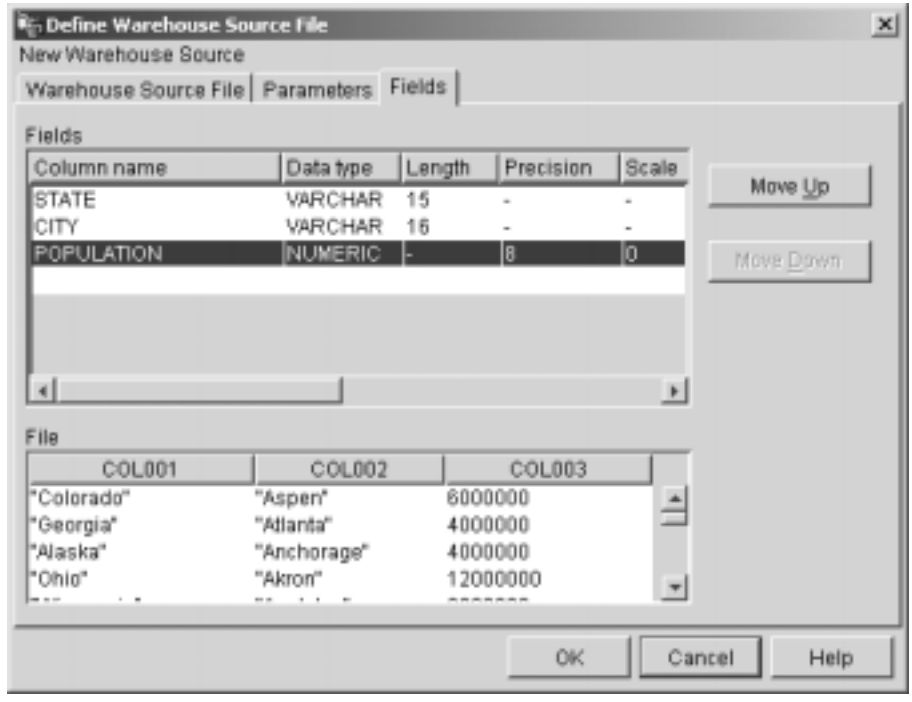

4. Click **OK**.

   The Define Warehouse Source File notebook closes.

Leave the Define Warehouse Source notebook open for the next task.

Defining security for a warehouse file source

When you create a file source, you must define security for the source. This exercise shows you how to define security for a warehouse file source.

To define security for a warehouse file source:

1. In the Define Warehouse Source notebook, click the **Security** tab.
2. Click the **Tutorial Warehouse Group**.
3. Click **>** to move the Tutorial Warehouse Group to the **Selected warehouse groups** list. This gives your user ID the ability to create steps that use this warehouse source.
4. Click **OK** to save your changes and close the Define Warehouse Source notebook.

What you just did

In this lesson, you:

- Viewed the data in a table and a file
- Defined a relational warehouse source
- Defined a warehouse file source

Chapter 4. Defining warehouse targets

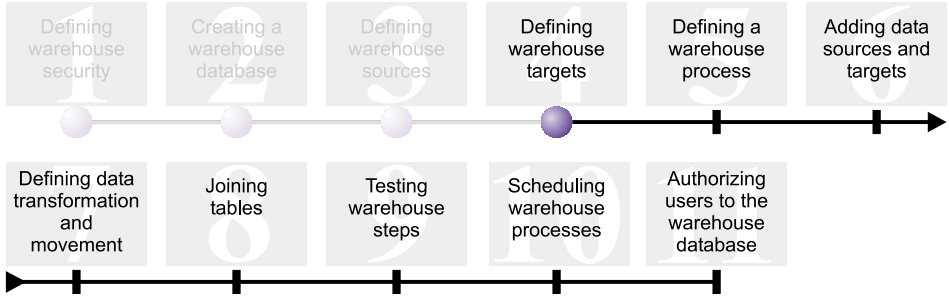

In this lesson, you will define warehouse targets. *Warehouse targets* identify the database and tables that the Data Warehouse Center is to use for your data warehouse. Generally, the target tables that are defined in a warehouse target are also used as the dimension and fact tables in a star schema. However, the warehouse target might also include interim target tables that are used for data transformation. For more information about creating a star schema, see the *Business Intelligence Tutorial: Extended Lessons in Data Warehousing* at http://www.ibm.com/software/data/bi/downloads.html.

In this lesson, you will complete the following tasks:

- Define the Tutorial Targets warehouse target. The Tutorial Targets warehouse target is a logical definition for the warehouse database that you created in Chapter 2, “Creating a warehouse database” on page 9.
- Define the DEMOGRAPHICS_TARGET table manually. This table resides in the warehouse database.

This lesson takes approximately 7 minutes to complete.

Defining a warehouse target

This exercise shows you how to define a warehouse target. To define the Tutorial Targets warehouse target you must complete the following tasks:

- Open the Define Warehouse Target notebook
- Specify information about a warehouse target
- Specify information about the target database
- Define security for the warehouse target

Opening the Define Warehouse Target notebook

To open the Define Warehouse Target notebook:

1. From the Data Warehouse Center window, right-click the **Warehouse Targets** folder.
2. Click **Define —> DB2 Family**.

   The Define Warehouse Target notebook opens.

Leave the Define Warehouse Target notebook open for the next task.

Specifying information about a warehouse target

To specify information about a warehouse target:

1. In the **Name** field, type the business name for the warehouse target:

   ```
   Tutorial Targets
   ```
2. In the **Administrator** field, type your name as the contact for the warehouse target.
3. In the **Description** field, type a short description of the data:

   ```
   Warehouse for The Beverage Company
   ```

Leave the Define Warehouse Target notebook open for the next task.

Specifying information about the target database

When you define a warehouse target, you must specify information about the warehouse database.

To specify information about the warehouse database:

1. In the Define Warehouse Target notebook, click the **Database** tab.
2. In the **Database name** field, click or type the name of the database:

   ```
   TUTWHS
   ```

   You are creating the target database in the default system, so you can skip the **System name** field.
3. In the **User ID** field, type the user ID that you used to create the sample databases.
4. In the **Password** field, type the password for the user ID.
5. In the **Verify password** field, type the password again.

   Use the default values for the rest of the controls on the page.

Leave the Define Warehouse Target notebook open for the next task.

Defining security for a warehouse target

When you define a warehouse target, you must define security for it.

To define security for a warehouse target:

1. In the Define Warehouse Target notebook, click the **Security** tab.
2. Select the **Tutorial Warehouse Group**.
3. Click > to move the Tutorial Warehouse Group to the **Selected warehouse groups** list.

   Adding the target to the warehouse group authorizes the users in the group (in this case, you) to create steps that use this warehouse target.
4. Click **OK** to save your changes and close the Define Warehouse Target notebook.

Defining a target table

In this exercise you will define the DEMOGRAPHICS_TARGET table within the Tutorial Targets warehouse target. You will complete the following tasks:

- Open the Define Warehouse Target Table notebook
- Specify information about a target table
- Add columns to the target table

Opening the Define Warehouse Target Table notebook

To open the Define Warehouse Target Table notebook:

1. Expand the **Warehouse Targets** tree until you see the **Tables** folder under the **Tutorial Targets** warehouse target.
2. Right-click the **Tables** folder, and click **Define**.

   The Define Warehouse Target Table notebook opens.

Leave the Define Warehouse Target Table notebook open for the next task.

Specifying information about a target table

This exercise shows you how to specify information about a target table.

1. In the **Table schema** list, specify the default schema, IWH.
2. In the **Table name** field, type the name of the target table:

   ```
   DEMOGRAPHICS_TARGET
   ```

   Because you are creating the tables in the default table space, you can skip the **Table space** and **Index table space** lists.
3. In the **Description** field, type the description of the table:

   ```
   Demographics data for sales regions
   ```
4. In the **Business name** field, type the business name for the table:

   ```
   Demographics Target
   ```
5. Verify that **Data Warehouse Center created table** is selected.

   The Data Warehouse Center creates this table when the step is promoted to test mode.

You use this option when you want the Data Warehouse Center to create the target table. Otherwise, you can use a target table that is already defined.

6. Verify that **Grant to public** is selected.

   This check box specifies that anyone who has access to the database has access to the table. Use the default values for the rest of this page.

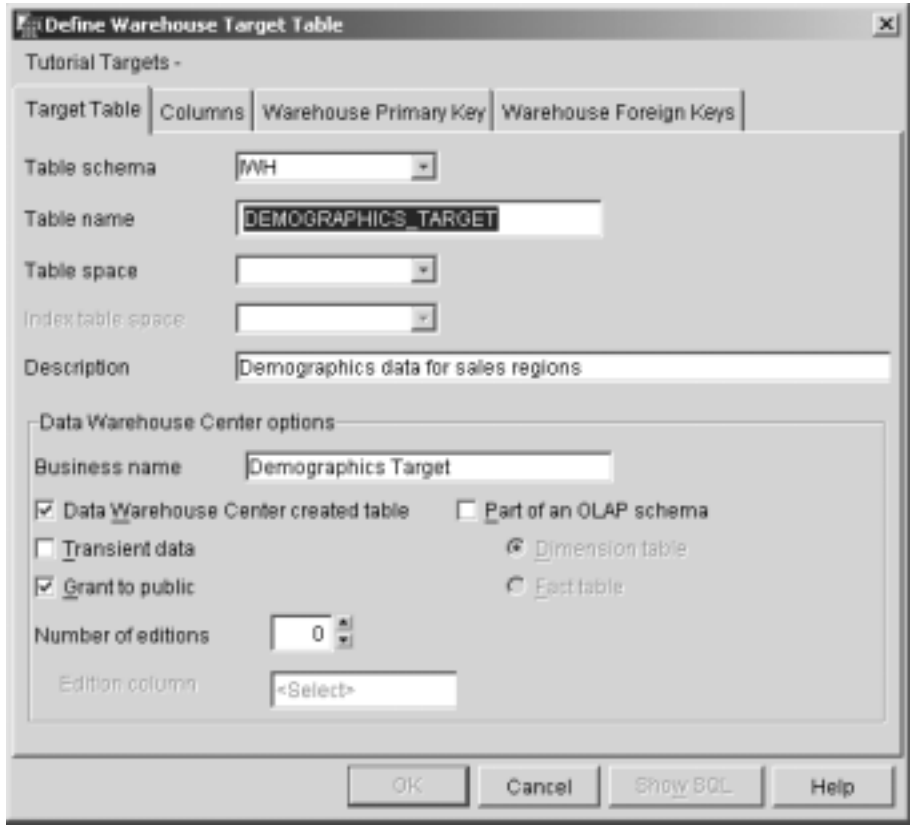

Leave the Define Warehouse Target Table notebook open for the next task.

Adding columns to the target table

To add columns to the target table:

1. Click the **Columns** tab.
2. Right-click the empty space in the list.
3. Click **Add**.

   A row is added to the list, and you can define the **State** column, which is one of the main values for the table.
4. Click the **Column name** column, and type STATE.

5. In the **Data type** column, verify that **CHAR** is specified.
6. Click the **Length** column, and replace the default value with `50`.

   Skip the **Precision** and **Scale** columns, because they apply to decimal data only.
7. Verify that **Nullable** is selected. Scroll to the right if you cannot see the Nullable check box.
8. In the **Business name** field, type `State`.
9. Repeat steps 2 through 8 using the values in the following table to define the rest of the columns in the DEMOGRAPHICS_TARGET table:

Name	Data type	Length	Allow nulls	Business name
CITY	CHAR	50	Clear the **Nullable** check box	City
POPULATION	INTEGER	N/A	Clear the **Nullable** check box	Population

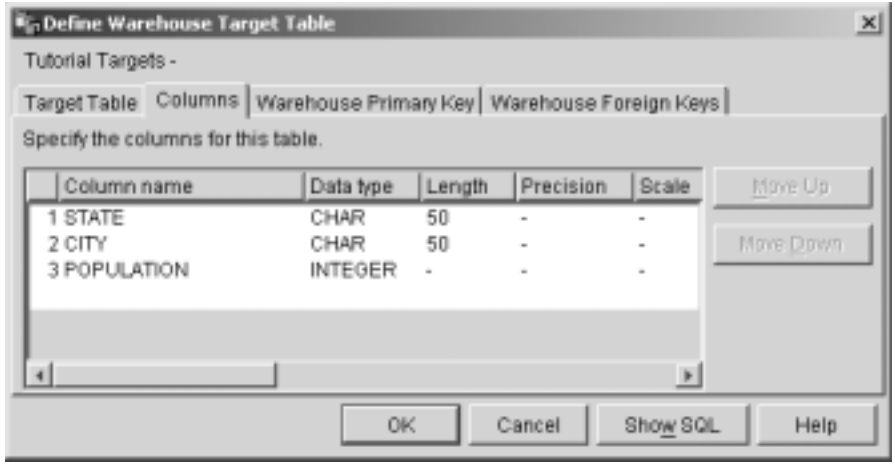

   Skip the rest of the notebook.
10. Click **OK**.

    The Define Warehouse Target Table notebook closes. The DEMOGRAPHICS_TARGET table is displayed under the **Tables** folder.

What you just did

In this lesson, you:

- Defined a warehouse target
- Defined a warehouse target table, DEMOGRAPHICS_TARGET

Chapter 5. Defining a warehouse process

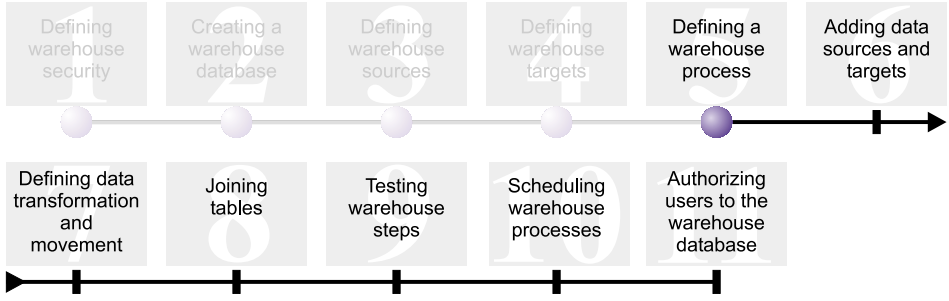

Warehouse processes contain a series of steps that define data transformation and movement. In this lesson you will create the Build Tutorial Market Dimension process that will hold the warehouse sources, target table, and steps that you will use to define data transformation and movement in the next lesson.

To define the Build Tutorial Market Dimension process, you will complete the following tasks:

- Define a subject area
- Open the Define Process notebook
- Specify information about the process
- Define security for the process

This lesson takes approximately 9 minutes to complete.

Defining a subject area

After you create and define your warehouse sources and target, you can define how the data will be moved and transformed. In the Data Warehouse Center, you use subject areas, processes, and steps to organize, move, and transform your source data, and then insert that data into the warehouse database.

In this lesson, you will use the Data Warehouse Center to define the TBC Tutorial subject area. A *subject area* identifies and groups processes that relate to a logical area of the business.

For example, you are building a warehouse of sales and marketing data for TBC, so you define a Sales subject area and a Marketing subject area. You then add the processes that relate to sales to the Sales subject area. Similarly, you add the definitions that relate to the marketing data to the Marketing subject area.

Any user can define a subject area, so you do not need to change the authorizations for the Tutorial Warehouse Group.

Defining the TBC Tutorial subject area

To define the subject area:

1. From the Data Warehouse Center tree, right-click on the **Subject Areas** folder, and click **Define**.

   The Define Subject Area notebook opens.

2. In the **Name** field, type the business name of the subject area for this tutorial:

   ```
   TBC Tutorial
   ```

   The name can be 80 characters including spaces.

3. In the **Administrator** field, type your name as the contact for this new subject.

4. In the **Description** field, type a short description of the subject area:

   ```
   Tutorial subject area
   ```

   You can also use the **Notes** field to provide additional information about the subject area.

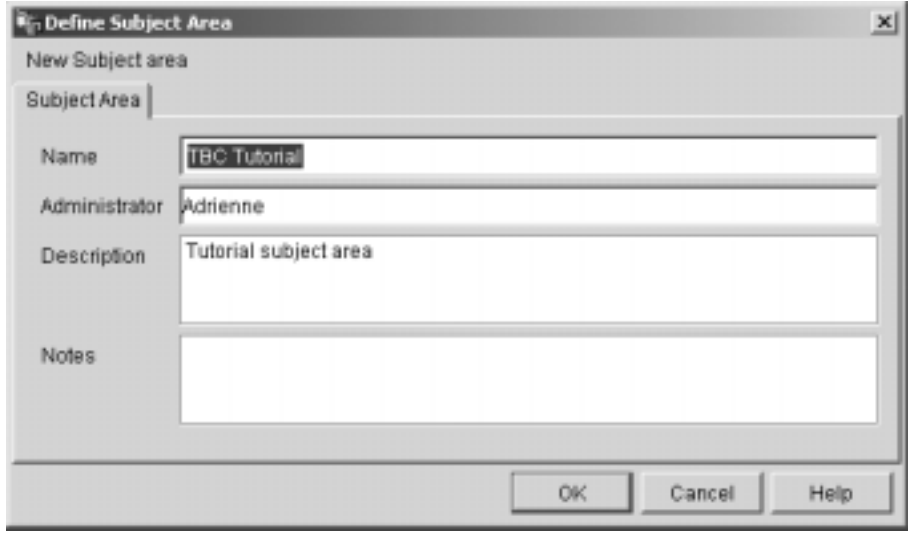

5. Click **OK** to create the subject area in the Data Warehouse Center tree.

Opening the Define Process notebook

To open the Define Process notebook:

1. From the Data Warehouse Center window, expand the **Subject Areas** tree.
2. Expand the **TBC Tutorial** subject area, which you defined in "Defining a subject area" on page 33.
3. Right-click the **Processes** folder, and click **Define**.

   The Define Process notebook opens.

Leave the Define Process notebook open for the next task.

Specifying information about the process

To specify information about the process:

1. From the Define Process notebook, type the name of the process in the **Name** field:

   ```
   Build Tutorial Market Dimension
   ```

   The name can be up to 80 characters long and is case-sensitive. The first character of the name must be alphanumeric.
2. In the **Administrator** field, type your name as the contact for the process definition.
3. In the **Description** field, type the description of the process:

   ```
   Process to create the LOOKUP_MARKET table
   ```

Leave the Define Process notebook open for the next task.

Defining security for the process

To define security for the process:

1. Click the **Security** tab.
2. In the **Available warehouse groups** list, click the **Tutorial Warehouse Group**, which you defined in "Defining the warehouse group" on page 5.
3. Click **>** to move the Tutorial Warehouse Group to the **Selected warehouse groups** list.

   Adding the process to the warehouse group authorizes the users in the group (in this case, you) to open and add objects to the process.
4. Click **OK**.

   The Define Process notebook closes.

What you just did

In this lesson you completed the following tasks to create the Build Tutorial Market Dimension process:

- Defined the TBC Tutorial subject area
- Opened the Define Process notebook
- Specified information about the process
- Defined security for the process

Chapter 6. Adding data sources and targets to a process

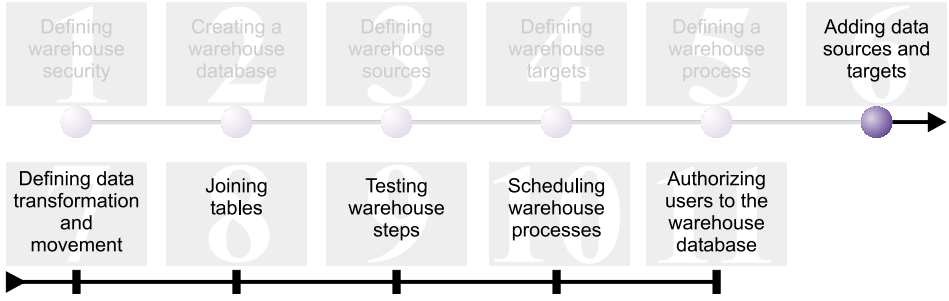

To define the flow of data, you need to add each source that the steps transform and the target tables that result from the transformation. In this lesson, you will add the data sources to the process.

In the Build Tutorial Market Dimension process, you load the demographics.txt file into the target database. You must add to the process the source file and the DEMOGRAPHICS_TARGET table for the step. The demographics.txt source file is part of the Tutorial file source warehouse source, which you defined in Chapter 3, “Defining warehouse sources” on page 13. The DEMOGRAPHICS_TARGET table is part of the Tutorial Targets warehouse target, which you defined in Chapter 4, “Defining warehouse targets” on page 27.

This lesson shows you how to open the process and add objects to it, so that you can graphically define the flow of data.

This lesson takes approximately 7 minutes to complete.

Opening the process

To open the Build Tutorial Market Dimension process:

1. In the Data Warehouse Center window, expand the **Subject Areas** tree.
2. Expand the **TBC Tutorial** tree until you see the **Build Tutorial Market Dimension** process.
3. Right-click the **Build Tutorial Market Dimension** process.
4. Click **Open**.

Adding data sources to a process

The following exercises show you how to add a file source and source table to a process.

Adding the demographics.txt file source to the process

In Chapter 3, "Defining warehouse sources" on page 13, you defined the demographics.txt file as part of the Tutorial file source warehouse source. This exercise shows you how to add the demographics.txt file source to a process.

To add the demographics.txt file to the Build Tutorial Market Dimension process:

1. Click the **Add Data** icon. 
2. Click the canvas (the blank area in the window) at the spot where you want to place the file source. The Add Data window opens.
3. In the **Available source and target tables** list, expand the **Warehouse Sources** tree.

   A list of the available warehouse sources is displayed.
4. Expand the tree for the Tutorial file source warehouse source.
5. Expand the **Files** tree.

   In the tree, select the demographics.txt file:

   x:\program files\ibm\sqllib\samples\db2sampl\dwc\demographics.txt, where *x* is the drive on which you installed the sample.
6. Click **>** to add the demographics file to the **Selected source and target tables** list.
7. Click **OK** to close the window and add the table to the process.

Adding the SAMPLTBC.GEOGRAPHIES table to the process

Now you need to add the source table to the process.

To add the SAMPLTBC.GEOGRAPHIES source table to the process:

1. Click the **Add Data** icon. 
2. Click the canvas at the spot where you want to place the table. The Add Data window opens.
3. Expand the **Warehouse Sources** tree until you see the Tutorial Relational Source warehouse source.
4. Expand the **Tutorial Relational Source** tree until you see the SAMPLTBC.GEOGRAPHIES table.
5. Click the **SAMPLTBC.GEOGRAPHIES** table.
6. Click **>** to add the SAMPLTBC.GEOGRAPHIES table to the **Selected source and target tables** list.

7. Click **OK** to close the window and add the table to the process.

   The SAMPLTBC.GEOGRAPHIES table is displayed in the Process Model window with the demographics.txt file source.
8. Move the demographics.txt icon to the upper left corner of the Process Model window, and place the GEOGRAPHIES icon in the upper right corner of the Process Model window.

Adding the DEMOGRAPHICS_TARGET table to the process

Target tables hold the source data after it is transformed. In this exercise, you will add the DEMOGRAPHICS_TARGET table to the process.

To add the DEMOGRAPHICS_TARGET table to the process:

1. Click the **Add Data** icon. 
2. Click the canvas below the demographics.txt file source icon.

   The Add Data window opens.
3. In the **Available source and target tables** list, expand the **Warehouse Targets** tree.

   A list of the available warehouse targets is displayed.
4. Expand the **Tutorial Targets** warehouse target tree.
5. Expand the **Tables** tree.

   The DEMOGRAPHICS_TARGET table is displayed in the list.
6. Click the DEMOGRAPHICS_TARGET table.
7. Click **>** to add the DEMOGRAPHICS_TARGET table to the **Selected source and target tables** list.
8. Click **OK** to close the window, and add the table to the process.

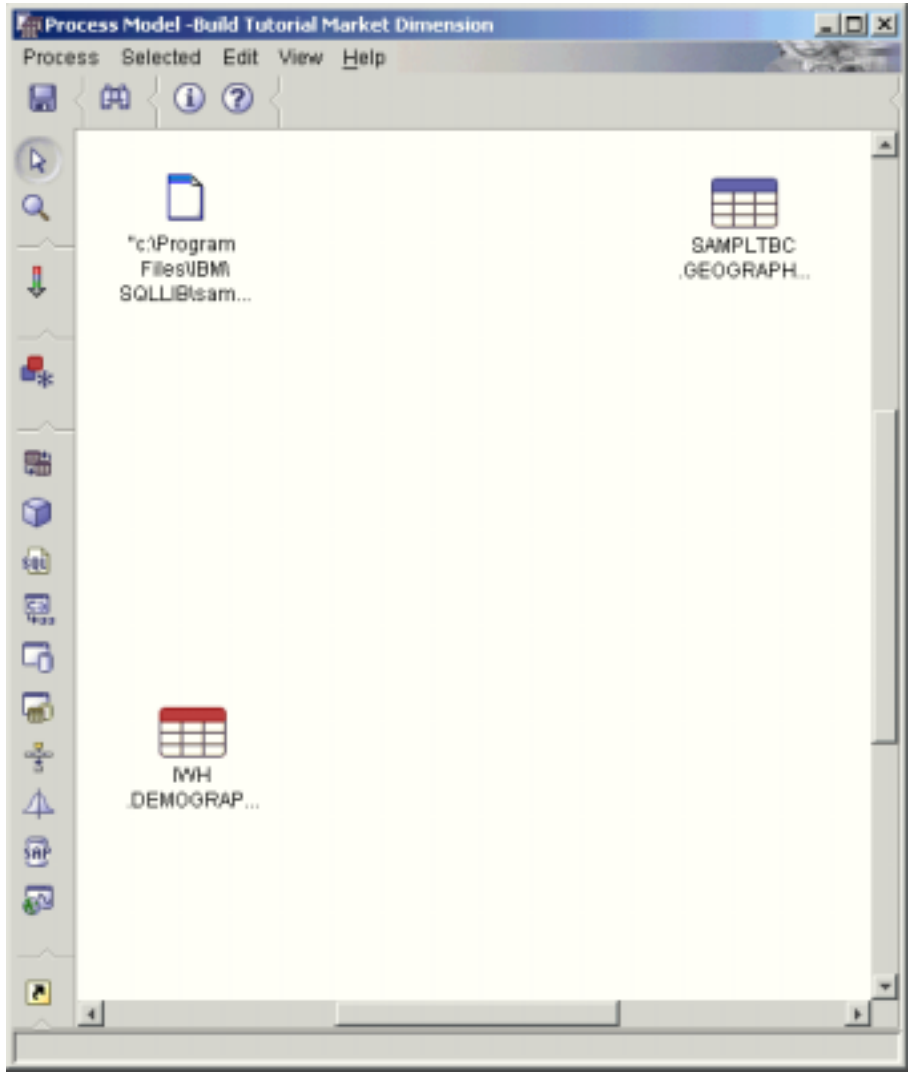

What you just did

In this lesson you completed the following tasks:

- Opened the process
- Added data sources to the process
- Added a target table to the process

Chapter 7. Defining data transformation and movement

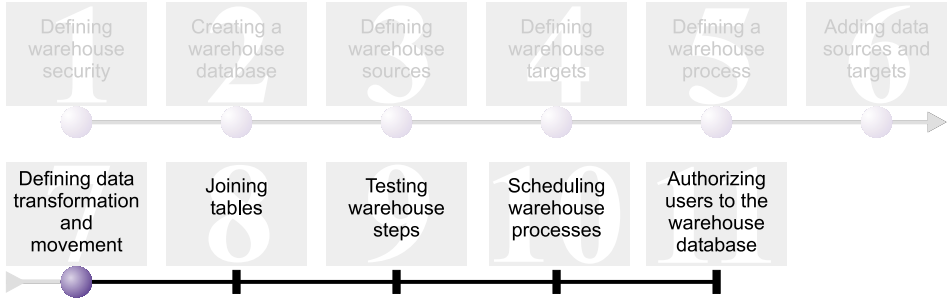

Now that you have warehouse sources, a warehouse target, and a warehouse process, you can move and transform the geographic and demographic data for TBC's data warehouse. In this lesson, you will define how the Data Warehouse Center is to move and transform data into a format for the data warehouse.

You will define steps, sources, and targets in the Build Tutorial Market Dimension process. You will complete the following tasks in this lesson:

- Add a step that loads the demographics.txt file into the warehouse database.
- Add a step that selects data from the GEOGRAPHIES table and inserts it into a target table using an SQL statement.

This lesson takes approximately 15 minutes to complete.

Adding steps to a process

After you define warehouse sources and targets, you can add the steps that define how TBC's demographic and geographic source data is to be transformed into the target data. To do this, you will define the following two steps in this exercise:

Load Demographics Data
A DB2 program step that loads data from the demographics.txt into a table in the TBC Warehouse database, which you created in Chapter 2, "Creating a warehouse database" on page 9.

Select Geographies Data
An SQL step that selects columns from the GEOGRAPHIES source table.

In a later lesson, you will define an SQL Select and Insert step that joins the two target tables that you are populating in this exercise.

Defining the Load Demographics Data step

To define the Load Demographics Data step:

1. From the palette (the toolbar on the left side of the window), click the DB2 Programs icon. 

   Each program option in the menu represents a program group, which is a grouping of similar programs.
2. Click **DB2 UDB —> Load**.
3. Click the spot on the canvas between the demographics.txt file and the DEMOGRAPHICS_TARGET table.

   An icon for the step is added to the window.
4. Right-click on the step, and click **Properties**.

   The Properties notebook for the step opens.

Leave the Properties notebook open for the next task.

Specifying general information about the step

To specify general information about the step:

1. From the DB2 Universal Database page in the Properties notebook for the step, type the following name in the **Name** field:

   ```
   Load Demographics Data
   ```
2. In the **Administrator** field, type your name as the contact for this step.
3. In the **Description** field, type the description of the step:

   ```
   Loads demographics data into the warehouse.
   ```
4. Click **OK**.

   The Properties notebook for the step closes.

Linking the Load Demographics Data step to a warehouse source and target

When you define a DB2 UDB Load step, you must link it to the source data that it will load.

To link the Load Demographics Data step to source data:

1. Click the **Task Flow** icon. 
2. Click the **Data Link** icon. 

   The Data Link icon defines the flow of data from the source file, through transformation by a step, to the target table.

3. Click the middle of the demographics.txt file, and drag the pointer to the Load Demographics Data step.

   The Data Warehouse Center draws a line between the file and the step. The line indicates that the demographics.txt file contains the source data for the step.

4. Click the middle of the Load Demographics Data step, and drag the pointer to the DEMOGRAPHICS_TARGET table.

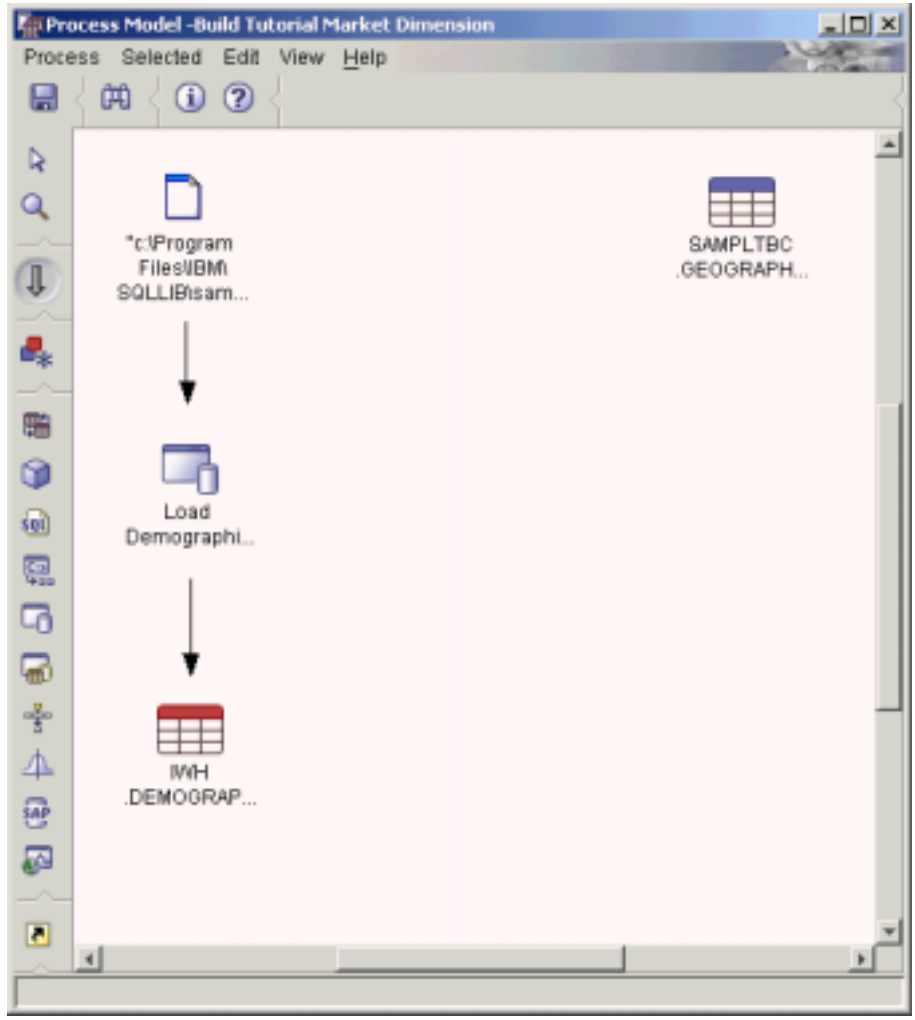

   This line indicates that the DEMOGRAPHICS_TARGET table contains the target data for the step.

Defining parameters for the Load Demographics Data step

To define parameters for the Load Demographics Data step:

1. Right-click on the Load Demographics Data step, and click **Properties**. The Properties notebook for the Load Demographics Data step opens.
2. Click the **Parameters** tab.
3. From the **Load mode** list, select **REPLACE**.

   For this exercise, use the default values for the rest of the page.

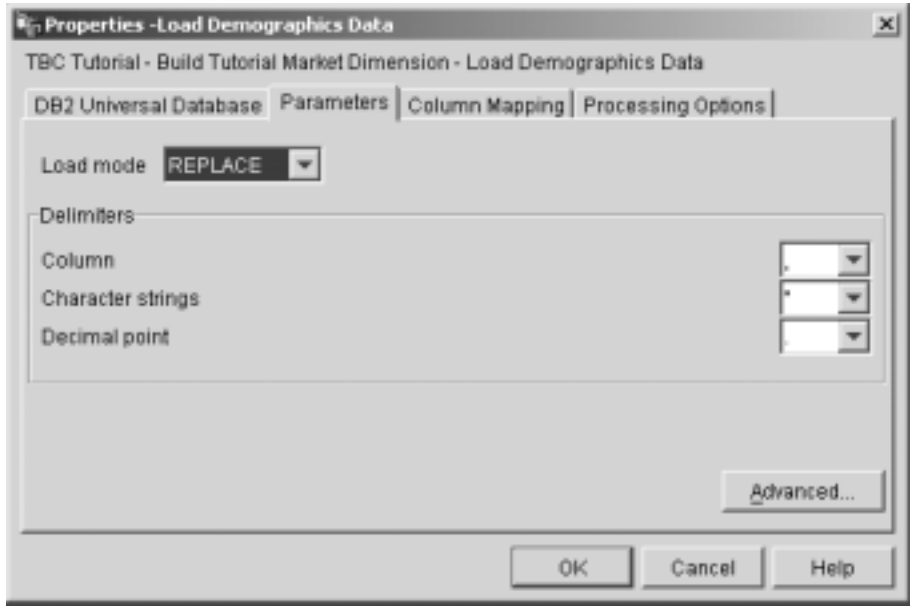

   You can skip the Column Mapping page because the Load program does not use column mapping.
4. Click **OK**.

   The Properties notebook for the step closes.

Defining the Select Geographies Data step

The Select Geographies Data step selects columns from the source table, GEOGRAPHIES, and inserts them into the target table that you will create in a later exercise.

To define the Select Geographies Data step:

1. From the palette, click the **SQL** icon , and click **Select and Insert**.
2. Click a spot on the canvas below the GEOGRAPHIES table to add a step icon to the window.
3. Open the Properties notebook for the step.
4. On the SQL page, type the following name in the **Name** field:

   `Select Geographies Data`

5. In the **Administrator** field, type your name as the name of the contact for the step.
6. In the **Description** field, type the description of the step:

   ```
   Selects geographical data from the warehouse source
   ```

7. Click **OK**.

   The Properties notebook for the step closes.

Linking the Select Geographies Data step to a warehouse source and target

You must specify a source and a target for the Select Geographies Data step to define the data that it will move and where the data will be placed after it is moved and transformed.

To link the Select Geographies Data step to a warehouse source and target:

1. In the Process Model window, click the **Task Flow** icon. 
2. Click the **Data Link** icon. 
3. Click the middle of the GEOGRAPHIES source table, and drag the pointer to the middle of the Select Geographies Data step.

   The Data Warehouse Center draws a line that indicates that the GEOGRAPHIES source table contains the source data for the step.

   You do not need to link a target table to the step because you will use the Data Warehouse Center to create a target table.

Selecting columns from the GEOGRAPHIES source table

To select columns from the GEOGRAPHIES source table:

1. Open the Properties notebook for the Select Geographies Data step.
2. Click the **SQL Statement** tab.
3. Click **Build SQL**.

   The Build SQL wizard opens. The **Tables** page is displayed.
4. Verify that SAMPLTBC.GEOGRAPHIES is listed.

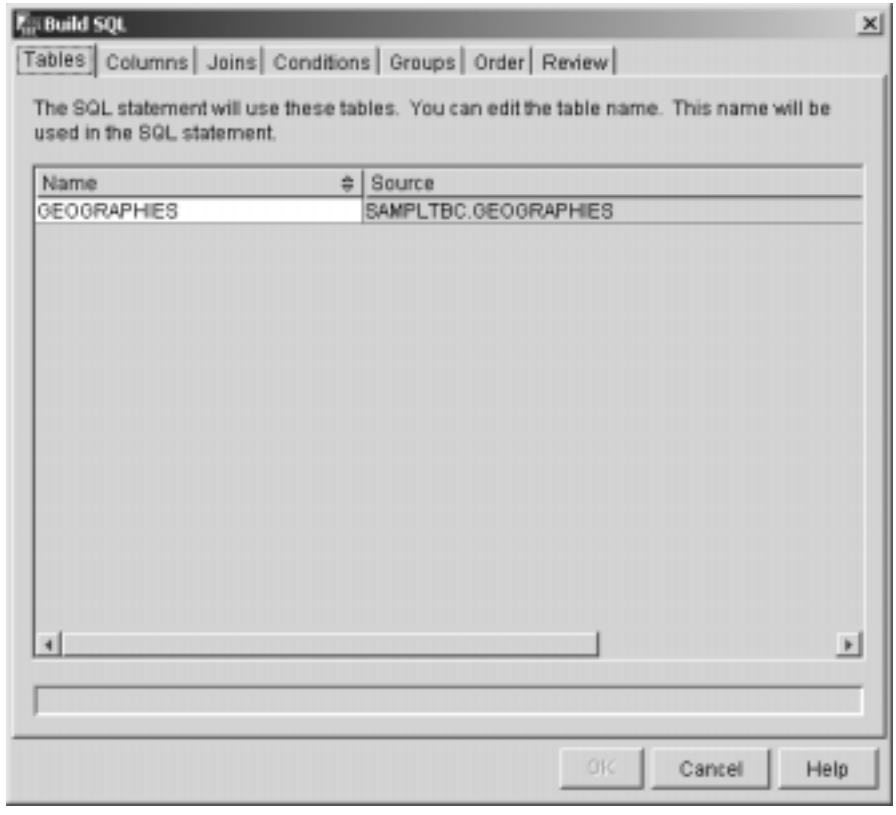

5. Click the **Columns** tab.
6. Click >> to add all the columns from the SAMPLTBC.GEOGRAPHIES table to the **Selected columns** list.

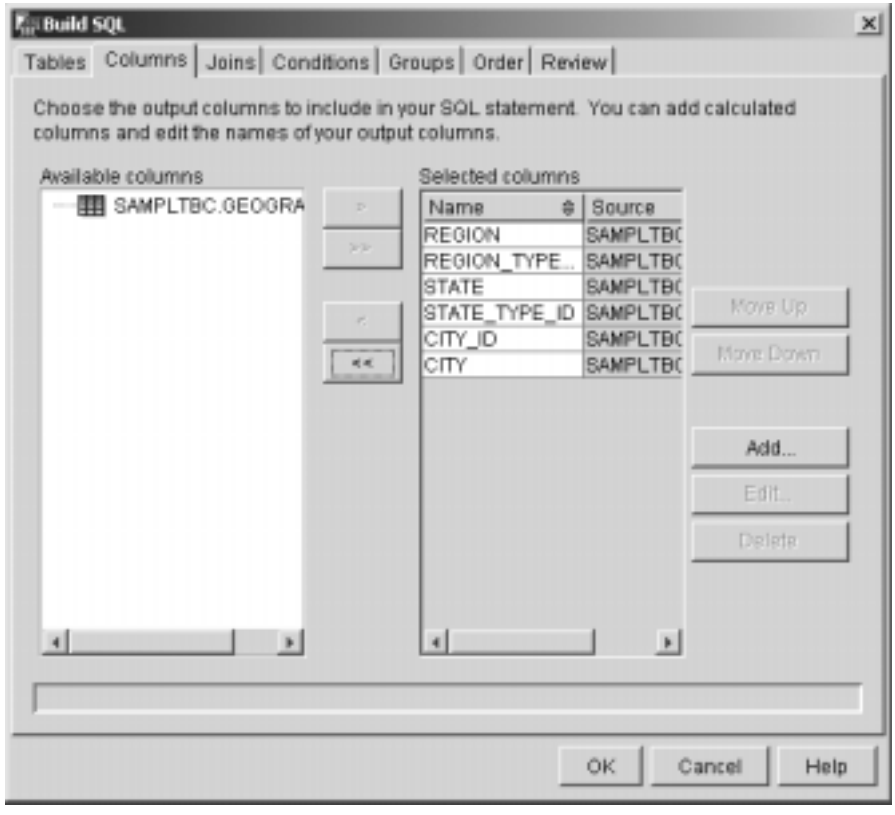

7. Click the **Review** tab to look at the SQL statement that you created.
8. Click **OK**.

The Build SQL wizard closes. The SQL statement that you created is displayed on the SQL Statement page.

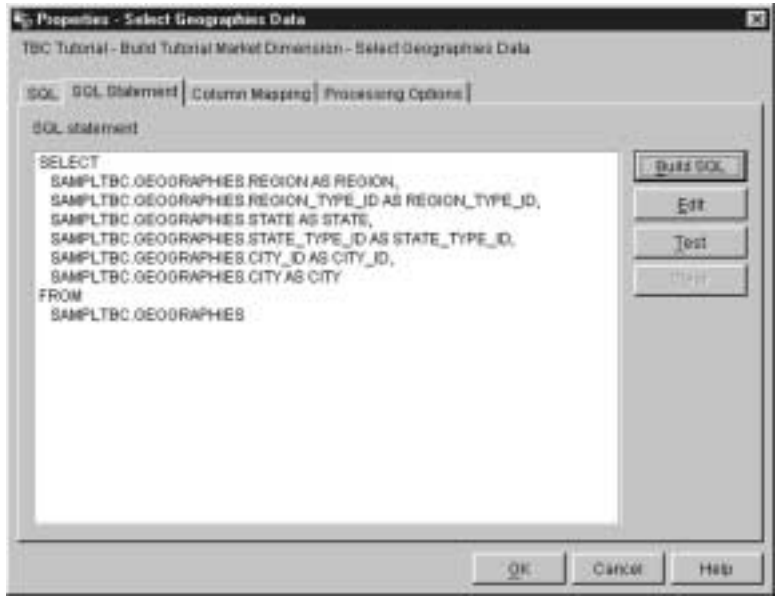

9. Click **Test** to test the SQL that you created.

   The Data Warehouse Center returns sample results of your SELECT statement. Compare your results to the results that you got in the sample shown in "Viewing table data" on page 14.
10. Click **Close** to close the window.

Now that you have specified which columns the SQL step will access in the source table, you are ready to create the target table. Leave the Properties notebook for the step open for the next task.

Creating the GEOGRAPHIES_TARGET table

In this exercise, you will create the GEOGRAPHIES_TARGET table using the Column Mapping page in the Select Geographies Data step.

To create the GEOGRAPHIES_TARGET table:

1. In the Properties notebook for the Select Geographies Data step, click the **Column Mapping** tab.

   The source columns are displayed on the left side of the page, and the target columns list is on the right side of the page. There are no target columns in the list because you have not specified the parameters for creating the target table.
2. Click **Generate Default Table**.

   The Generate Default Table window opens.

3. In the **Warehouse target** list, click **Tutorial Targets**.

   The warehouse target is the database or file system in which the target table is created.
4. In the **Table schema** list, specify the default table schema, IWH.

   Because you are creating the table in the default table space, you can skip the **Table space** list.
5. In the **Table name** field, type:

   `GEOGRAPHIES_TARGET`
6. Click **OK** to close the Generate Default Table window.

   The columns of the GEOGRAPHIES_TARGET table are displayed in the target columns list on the right side of the Column Mapping page.
7. Verify that the source columns map to the correct target columns.

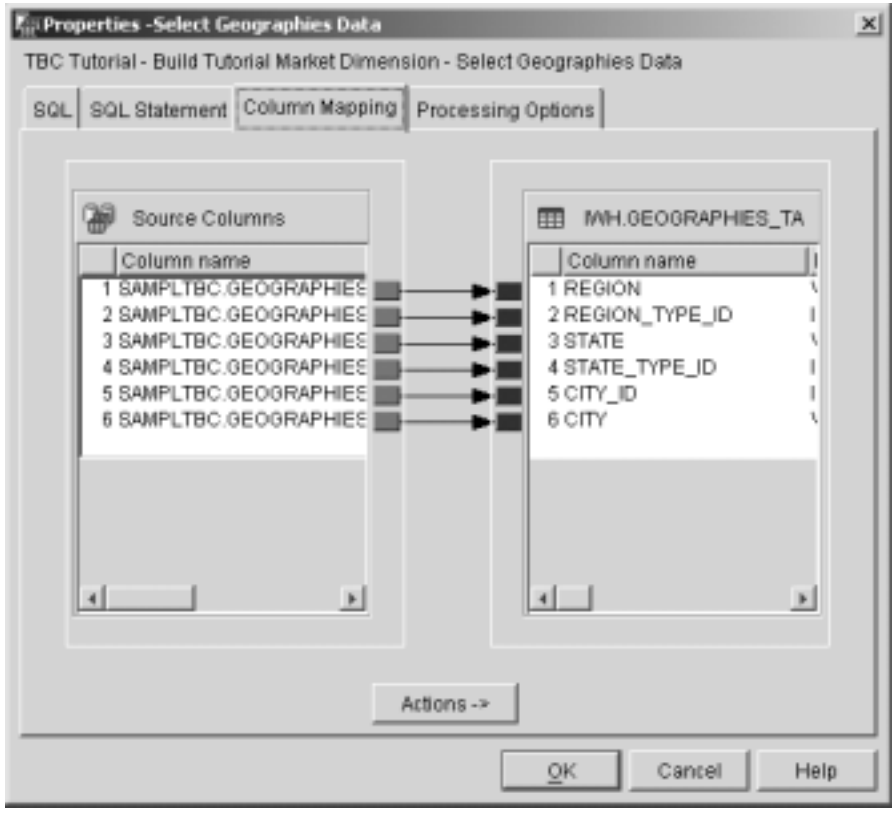

8. Click the **Processing Options** tab.
9. Verify that **Replace** is selected in the **Population type** list.
10. Verify that the **Run on demand** check box is selected.

    For this exercise, you will use the default values on this page.

11. Click **OK**.

    The Properties notebook for the step closes. The Data Warehouse Center creates a target table called GEOGRAPHIES_TARGET and links it to the step.

Specifying properties for the GEOGRAPHIES_TARGET table

In this exercise, you will specify the properties for the GEOGRAPHIES_TARGET table that you created.

To specify the properties of the GEOGRAPHIES_TARGET table:

1. In the Process Model window, right-click the GEOGRAPHIES_TARGET table, and click **Properties**.
2. In the **Business name** field, type a descriptive name for the table:

   ```
   Geographies Target
   ```

   Use the default values for the rest of the notebook.
3. Click **OK**. The Properties notebook for the table closes.

What you just did

In this lesson you completed the following tasks:

- Defined two steps for moving and transforming data:
  - Load Demographics Data
  - Select Geographies Data

  The sources and targets for each of the steps are shown in the following table:

Step	Sources	Target
Load Demographics Data	demographics.txt file	DEMOGRAPHICS_TARGET table
Select Geographies Data	GEOGRAPHIES table	GEOGRAPHIES_TARGET table

- Added the data links for each step.
- Created a default target table, GEOGRAPHIES_TARGET.

Chapter 8. Joining tables

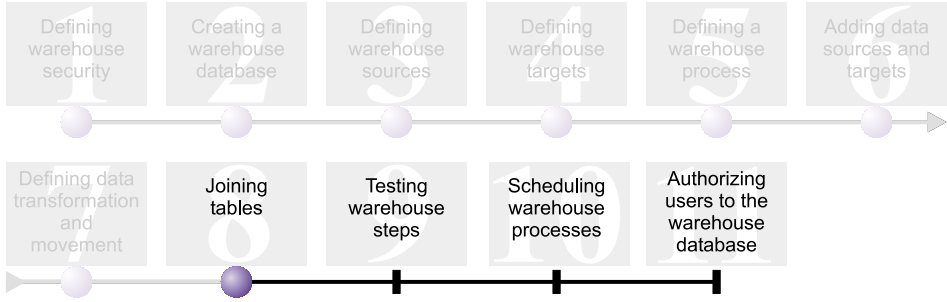

Now that you have loaded and inserted source data into target tables in the Build Tutorial Market Dimension process, you can join the geographic and demographic data into a single table. This lesson will show you how to create a join step that will join the data from the DEMOGRAPHICS_TARGET and GEOGRAPHIES_TARGET tables and write the joined table into the Tutorial Warehouse database.

You will complete the following steps to join the two tables:

- Define the Join Market Data step
- Join the DEMOGRAPHICS_TARGET and GEOGRAPHIES_TARGET tables
- Create the LOOKUP_MARKET table
- Define properties for the LOOKUP_MARKET table

This lesson takes approximately 10 minutes to complete.

Defining the Join Market Data step

You created steps that pull data from your sources and insert it into target tables. However, TBC wants to track data over both demographic and geographic areas, so you need to join data from the two target tables you created into a single table that contains both the demographic and geographic data.

Adding the Join Market Data step to the process

This exercise shows you how to use a Select and Insert SQL step to join data from the two target tables that you created and insert the data into another target table.

To add the Join Market Data step to the canvas and specify general information for it:

1. From the palette, click the SQL icon  , and click **Select and Insert**.
2. Click the spot on the canvas below the two columns of steps and tables.

   An icon for the step is added to the window.
3. Open the properties notebook for the step.
4. On the SQL page, type the following name in the **Name** field:

   `Join Market Data`
5. In the **Administrator** field, type your name as the contact for the step.
6. In the **Description** field, type the description of the step:

   `Joins the GEOGRAPHIES_TARGET table with the DEMOGRAPHICS_TARGET table.`
7. Click **OK**.

   The Properties notebook for the step closes.

Linking the Join Market Data step to source data

This exercise shows you how to link the Join Market Data step to the DEMOGRAPHICS_TARGET and GEOGRAPHIES_TARGET tables. These two tables serve as sources for the Join Market Data step.

To define the data flow for the Join Market Data Step:

1. Click the **Task Flow** icon. 
2. Click the **Data Link** icon. 
3. Click the middle of the GEOGRAPHIES_TARGET table, and drag the pointer to the Join Market Data step.

   The Data Warehouse Center draws a line that indicates that the GEOGRAPHIES_TARGET table contains source data for the step.
4. Repeat step 3 with the DEMOGRAPHICS_TARGET table and the Join Market Data step.

   The Data Warehouse Center draws a line that indicates that the DEMOGRAPHICS_TARGET table contains source data for the step.

   Because you will specify that the Data Warehouse Center is to create the target table, you do not need to link a target table to the step.

Joining the DEMOGRAPHICS_TARGET and GEOGRAPHIES_TARGET tables

In this exercise, you will join the DEMOGRAPHICS_TARGET and GEOGRAPHIES_TARGET tables. This exercise shows you how to complete the following tasks to create the join:

- Open the Build SQL window

- Specify the columns that you want to join
- Join the columns

Opening the Build SQL window

To open the Build SQL window:

1. Right-click the Join Market Data step.
2. Click **Properties**.

   The Properties notebook for the step opens.
3. Click the **SQL Statement** tab.
4. Click **Build SQL** to have the Data Warehouse Center create SQL. (Otherwise, you can create your own SQL.)

   The Build SQL window opens.
5. On the Tables page, verify that the DEMOGRAPHICS_TARGET and GEOGRAPHIES_TARGET tables are listed.

Leave the Build SQL window open for the next task.

Specifying columns for the join

This exercise shows you how to specify columns from each of the tables for the join.

To specify columns for the join:

1. From the Columns page of the Build SQL notebook, select each of the columns in the **Available columns** list, except for the STATE and CITY columns in the DEMOGRAPHICS_TARGET table, and click > to move them to the **Selected columns** list.

   You do not need the STATE and CITY columns in the DEMOGRAPHICS_TARGET table because they are already defined in the GEOGRAPHIES_TARGET table.
2. Click **Add**.

   The Expression Builder - Columns window opens.
3. In the **Expression** field, type the following CASE statement:

```
CASE
WHEN POPULATION < 9000000   THEN 1
WHEN POPULATION < 15000000  THEN 2
ELSE 3
END
```

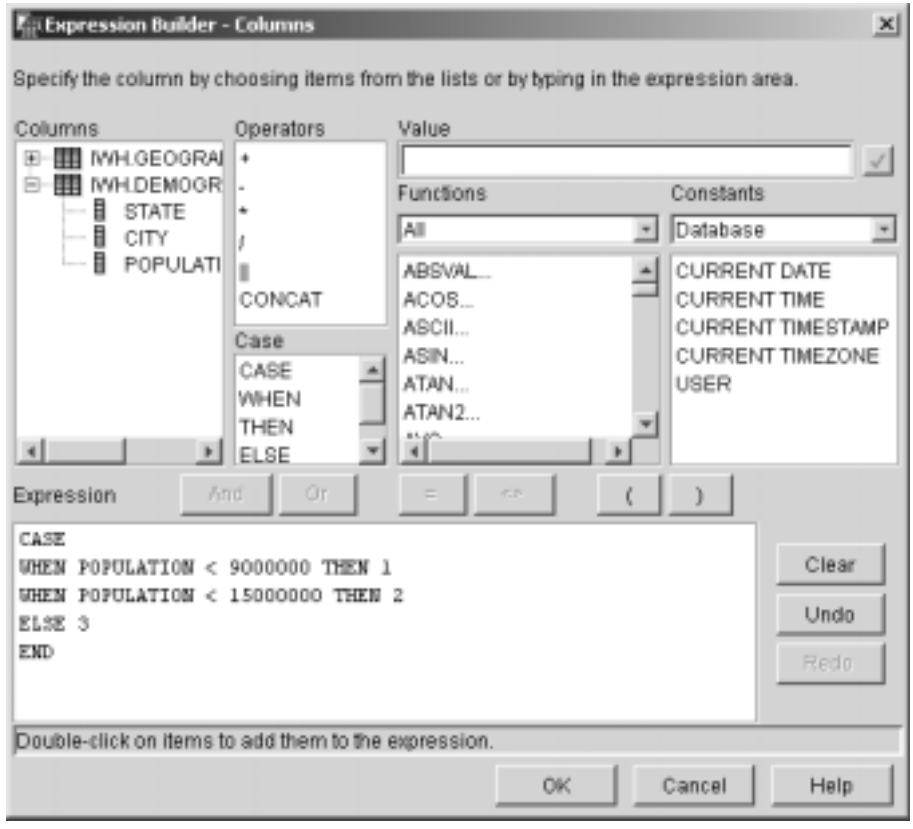

4. Click **OK**.

   A new column is added to the **Selected columns** list.
5. Click the **Name** field of the new column, and type the name of the column:

   `SIZE_ID`
6. Press Enter.
7. Click **Move Up** until the SIZE_ID column is above the POPULATION column.

Joining the columns

This exercise shows you how to join the CITY column in the GEOGRAPHIES_TARGET table to the CITY column in the DEMOGRAPHICS_TARGET table. You will join the tables using the CITY column because it is unique to both tables.

To join the columns:

1. Click the **Joins** tab.
2. In the GEOGRAPHIES_TARGET table, click the CITY column.

3. In the DEMOGRAPHICS_TARGET table, click the CITY column.
4. Click **Join**.

   The line between the CITY columns turns green indicating that the tables are joined on that column.
5. Click the **Review** tab to see the SQL statement that you just created.
6. Click **OK**.

   The Build SQL wizard closes.

Leave the Properties notebook open for the next step.

Creating the LOOKUP_MARKET table

This exercise will show you how to create the LOOKUP_MARKET table from the Column Mapping page of the Join Market Data step Properties notebook. This table will contain the demographic and geographic data that is pulled from the DEMOGRAPHICS_TARGET and GEOGRAPHIES_TARGET tables by the Join Market Data step.

To create the LOOKUP_MARKET table:

1. Click the **Column Mapping** tab.
2. Click **Generate Default Table**.

   The Generate Default Table window opens.
3. In the **Warehouse target** list, click **Tutorial Targets**.
4. In the **Table schema** list, type or click **IWH**. Skip the **Table space** list.
5. In the **Table name** field, type:

   `LOOKUP_MARKET`
6. Click **OK**.

   The Generate Default Table window closes. The target columns are displayed in the target columns list on the right side of the Column Mapping page.
7. Click the **Processing Options** tab.
8. In the **Population type** list, verify that **Replace** is selected.
9. Verify that **Run on demand** is selected.

   Use the default values for the rest of this page.
10. Click **OK**.

    The Properties notebook for the step closes, and the LOOKUP_MARKET table is displayed in the Process Model window.

Defining the properties of the LOOKUP_MARKET table

This exercise will show you how to define the properties for the LOOKUP_MARKET table. The table is defined as a dimension table, so that it can be used as a dimension table in a star schema.

To define the properties of the LOOKUP_MARKET table:

1. Right-click the LOOKUP_MARKET table, and click **Properties**.

   The Properties notebook for the table opens.
2. In the **Description** field, type a description of the table:

   `Market dimension data`
3. Select the **Part of an OLAP schema** check box and the **Dimension table** radio button.

   The LOOKUP_MARKET table is one of the dimension tables that you will include in a star schema in the *Business Intelligence Tutorial: Extended Lessons in Data Warehousing*.
4. Click the **Columns** tab.
5. Clear **Nullable** for the CITY_ID column.
6. Click the **Warehouse Primary Key** tab.
7. Click **CITY_ID** in the **Available columns** list.
8. Click **>** to move the CITY_ID column to the **Primary key columns** list.
9. In the **Constraint name** field, type:

   `Whse Market PK`
10. Click **OK**. The Properties notebook for the table closes.
11. Click the Save icon  in the toolbar to save the process.

What you just did

In this lesson, you completed the following tasks:

- Defined the Join Market Data step.
- Joined the DEMOGRAPHICS_TARGET and GEOGRAPHIES_TARGET tables.
- Created the LOOKUP_MARKET table.
- Defined properties for the LOOKUP_MARKET.

Chapter 9. Testing warehouse steps

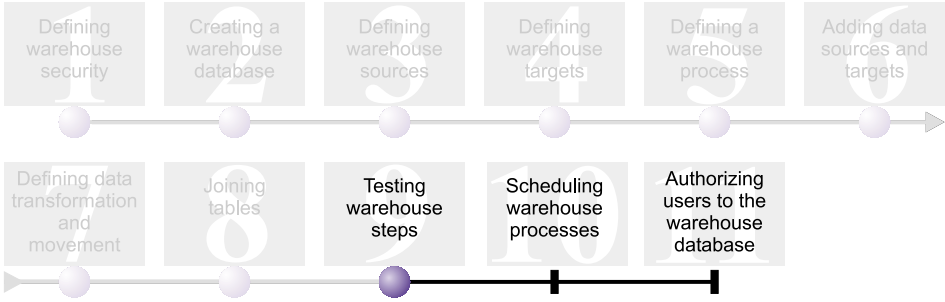

In this lesson, you will populate the LOOKUP_MARKET table by running the steps that you created in previous lessons.

You will then verify the results of running them.

Before you run the steps, you must promote them to test mode. Up to this point, the steps you created were in development mode. When a step is in development mode, you can change any of the specifications for the step. When you promote the step to test mode, the Data Warehouse Center creates the target table for the step. Therefore, after you promote a step to test mode, you can make only those changes that are not destructive to the target table. For example, you can add columns to a target table when its associated step is in test mode, but you cannot remove columns from the target table.

After you promote the steps to test mode, you will run each step individually. In a later lesson, you will specify that the steps run in sequence.

This lesson takes approximately 10 minutes to complete.

Testing the Load Demographics Data step

In this exercise, you will promote and run the Load Demographics Data step. Then you will promote the rest of the steps in the Build Tutorial Market Dimension process.

To test the Load Demographics Data step:

1. From the Process Model window for the Build Tutorial Market Dimension process, right-click the **Load Demographics Data** step.
2. Click **Mode —> Test**.

The Data Warehouse Center starts to create the target table and displays a progress window. Wait until the Data Warehouse Center finishes processing before you start the next procedure.

A lock appears on the step icon to indicate that only nondestructive changes can be made to the step.

To verify that the Demographics target table was created:

1. In the Process Model window, right-click the DEMOGRAPHICS_TARGET table.
2. Click **Sample Contents**.

   If the table exists, the Sample Contents window opens. If the table does not exist, you will receive an error message.
3. Click **Close**.

Repeat the steps in this exercise for the Select Geographies Data step and the Join Market Data step. The target table for the Select Geographies Data step is GEOGRAPHIES_TARGET. The target table for the Join Market Data step is LOOKUP_MARKET.

What you just did

In this lesson, you completed the following tasks:

- Promoted the Load Demographics Data, Select Geographies Data, and Join Market Data steps to test mode.
- Verified the results of promoting the steps.

In Chapter 10, “Scheduling warehouse processes” on page 59, you will schedule these steps to run automatically.

Chapter 10. Scheduling warehouse processes

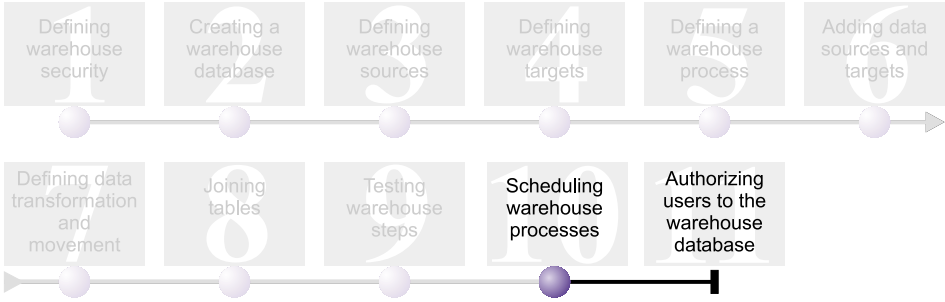

In this lesson, you will specify that the steps in the Build Tutorial Market Dimension process are to run in the following sequence:

1. Load Demographics Data
2. Select Geographies Data
3. Join Market Data

Then you will specify that the Load Demographics Data step is to run at a scheduled time. You will activate the schedule by promoting the steps in the process to production mode.

This lesson takes approximately 10 minutes to complete.

Running steps in sequence

To specify that the steps are to run in sequence:

1. If the Process Model window is closed, from the Data Warehouse window, right-click on the **Build Tutorial Market Dimension** process, and click **Open**. If the Process Model window is already open, proceed to the next step.
2. From the Process Model window, click the **Task Flow** icon. 
3. Click the **On Success** icon (the green arrow).

   On Success indicates that a step is to be started only if the step before it ran successfully. You can also select the following states:

   On Failure
   Indicates that a step is to be started only if the step before it failed. (The red arrow.)

On Completion

Indicates that a step is to be started when the step before it runs successfully or fails. (The blue arrow.)

For more information, see “Scheduling a step” in the Data Warehouse Center online help.

4. Click the Load Demographics Data step, then drag and drop the pointer onto the Select Geographies Data step. A green arrow that represents the task flow is displayed on the canvas between the two steps indicating that the Select Geographies Data step will run if the Load Demographics Data step completes successfully.
5. Click the Select Geographies Data step and drag the pointer to the Join Market Data step. This specifies that the Select Geographies Data step will

run before the Join Market Data step.

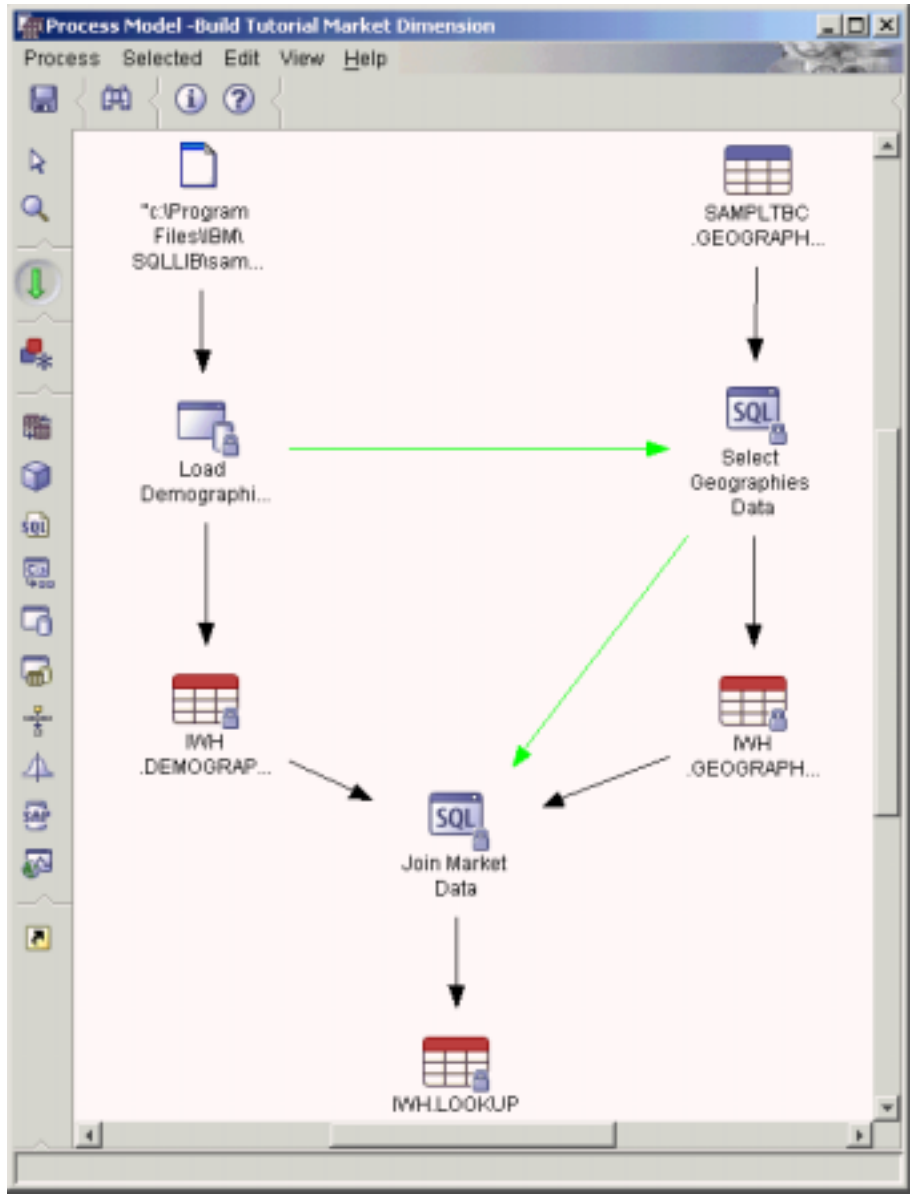

The steps will now run in the order that is listed in the introduction to this lesson.

Scheduling the first step

Now, you will schedule the Load Demographics Data step to start at a specified date and time. The Data Warehouse Center will start the step at the time you specify. When the Load Demographics Data step finishes running, the Data Warehouse Center starts the next step in the sequence that you defined in the previous section.

When you schedule a step, you can specify one or more dates and times on which the step is to run. You can also specify that the step is to run once or is to run at a specified interval, such as every Saturday.

To schedule the Load Demographics Data step:

1. In the Process Modeler window, right-click the Load Demographics Data step, and click **Schedule**.

   The Schedule notebook opens.
2. From the **Interval** list, click **Yearly**.
3. In the **Date** field, accept the default of the current date.
4. In the **Time** field specify a time that is a few minutes from the current time.

   The step will run at the time you select.
5. In the **End** field, accept the default that the schedule is to run indefinitely.
6. Click **Add**.

   The schedule is added to the **Schedule list**.

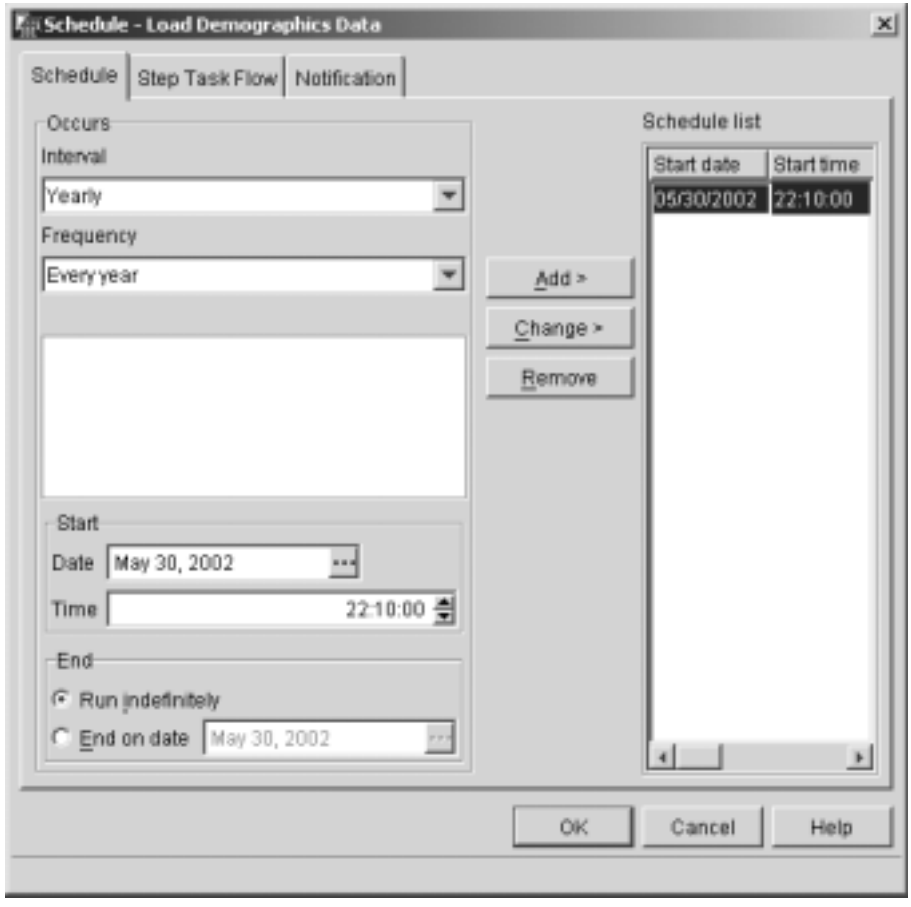

7. Click **OK**.

   The specified schedule is created.

Promoting the steps to production mode

To activate the schedule and the task flow links that you created, you must promote the steps to production mode. Production mode indicates that the steps are in their final format. In production mode, you can change only those settings that will not affect the step's processing, such as changing the description of the step.

To promote the steps to production mode:

1. Right-click the Load Demographics data step.
2. Click **Mode —> Production**, then click **Yes** to save the process.

   The Data Warehouse Center displays a progress window. Wait until the Data Warehouse Center finishes processing before continuing with this

exercise. When the progress window closes, the step icon changes to show two locks indicating that the step is in production mode.

3. Repeat the previous steps for the Select Geographies Data step and Join Market Data steps, in order.
4. Close the Process Model Window.

Monitoring the progress of a step

You can use the Work in Progress window to monitor the progress of a step or process in the Data Warehouse Center.

To monitor the progress of a step with the Work in Progress window, click **Warehouse —> Work in Progress** in the Data Warehouse Center window. The Work in Progress window opens. You can see the status of the step in the window. For example, if the step ran successfully, it has a status of Successful.

For more information about the Work in Progress window, see the topic "Work in Progress – Overview" in the Data Warehouse Center online help.

What you just did

In this lesson, you scheduled the steps that you created to run once a year at the current date and time in the following order:

1. Load Demographics Data
2. Select Geographies Data
3. Join Market Data

Then you promoted the steps to production mode to implement the schedule and monitored the results in the Work in Progress window.

Chapter 11. Authorizing users to the warehouse database

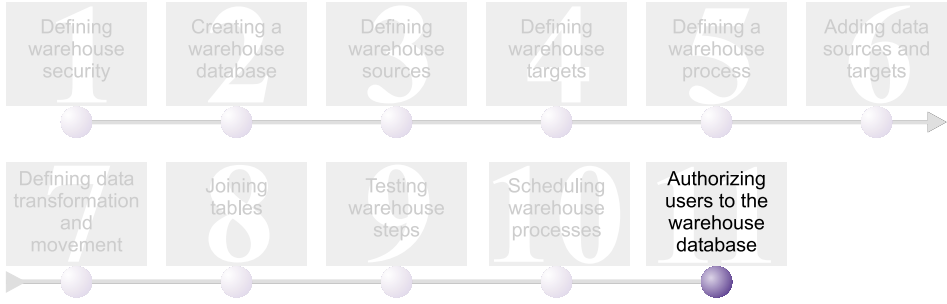

In this lesson, you will define privileges that authorize users to the warehouse database. Security for the warehouse database is managed within DB2 Universal Database. It is separate from the Data Warehouse Center security.

Access within DB2 Universal Database is managed by administrative authorities and user privileges within the database manager. Authorities are usually granted at the database level, and privileges are usually granted for objects within the database (for example, tables).

Privileges are controlled by users with SYSADM or DBADM authority, or by the creator of the object. You can grant privileges to users for the TUTWHS database because you are its creator.

This lesson takes approximately 7 minutes to complete.

Granting privileges

To grant privileges to the TUTWHS database:

1. From the DB2 Control Center, expand the objects in the TUTWHS database until you see the **Tables** folder.
2. Click the **Tables** folder. In the right panel, you will see all the tables in the database.
3. Right-click on the LOOKUP_MARKET table, and click **Privileges**.

   The Table Privileges window opens.
4. Click **Add User**.

   The Add User window opens.
5. Specify TUTUSER in the first field.

6. Click **OK**. The user is added to the User page.
7. Click TUTUSER.
8. Click **Grant All**. The user, TUTUSER is granted all of the available privileges. To grant individual privileges, use the **Privileges** list boxes.
9. Click **Apply** to process your request.
10. Close the Table Privileges window.

What you just did

In this lesson, you authorized users to use the LOOKUP_MARKET table. You can catalog the metadata that you defined in the Data Warehouse Center, so that the users that you authorized in this lesson can find the data that they need more easily.

Chapter 12. Summary

Congratulations! You now have completed the *Business Intelligence Tutorial: Introduction to the Data Warehouse Center*. In this tutorial, you completed the following tasks:

- You defined Data Warehouse Center security by defining a warehouse user and a warehouse group.
- You created a warehouse database.
- You defined warehouse sources.
- You defined warehouse targets.
- You defined warehouse processes.
- You added data sources and targets to a process.
- You defined how data would be transformed and moved in the warehouse.
- You defined the LOOKUP_MARKET table by joining data from two different sources.
- You promoted the steps, tested them, and scheduled them.
- You authorized users to access the Data Warehouse Center.

The *Business Intelligence Tutorial: Extended Lessons in Data Warehousing* is located at http://www.ibm.com/software/data/bi/downloads.html.

Appendix. Related information

This tutorial covers the most common tasks that you can accomplish with the DB2 Control Center and Data Warehouse Center. For more information about related tasks, see the following documents:

Control Center

- *IBM DB2 Universal Database Quick Beginnings for DB2 Clients*
- *IBM DB2 Universal Database Quick Beginnings for DB2 Servers*
- *IBM DB2 Universal Database SQL Reference Volume 1*
- *IBM DB2 Universal Database SQL Reference Volume 2*
- *IBM DB2 Universal Database Administration Guide: Implementation*

Data Warehouse Center

- *IBM DB2 Universal Database Data Warehouse Center Administration Guide*
- *IBM DB2 Warehouse Manager Installation Guide*

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.

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

Printed in U.S.A.