Лекция 3

Тема « *Creating an Analysis Services Database* »

Creating an Analysis Services Database Using Business Intelligence Development Studio

You are now ready to create a cube. The cube you create in this chapter is based on the relational database Adventure Works DW 2008 that is available at http://www.codeplex.com as part of Microsoft SQL Server 2008 sample databases. Many versions of Adventure Works are available on CodePlex. Download and install the SQL Server 2008 Adventure Works DW 2008 sample database for your machine ' s architecture. For example, if you have an x64 machine, the sample database to install is SQL2008.AdventureWorks_DW_BI_v2008.x64.msi.

Adventure Works DW 2008 contains the sales information of a fictional bicycle company. Figure 2 - 13 shows the structure of the data warehouse you will build in this chapter, which consists of two fact tables and eight dimension tables. FactInternetSales and FactResellerSales are fact tables. They each contain several measures and foreign keys related to their dimension tables. Both fact tables contain three dimension keys, ShipDateKey, OrderDateKey, and DueDateKey, which are joined to the dimension table DimDate. The FactInternetSales and the FactResellerSales fact tables join to the other appropriate dimension tables by a single key as shown in Figure 2 - 13 . The ParentEmployeeKey in the Employee table is joined with EmployeeKey in the same table, which is modeled as a parent - child hierarchy. You learn about parent - child hierarchies in Chapter 5 .

Figure 2-13

Creating a Data Source

Cubes and dimensions of an Analysis Services database must retrieve their data values from tables in a relational data store. This data store, typically part of a data warehouse, must be defined as a data source. An OLE DB data provider or .NET data provider is used to retrieve the data from the data source. OLE DB and .NET data providers are industry standard technologies for retrieving data from relational databases. If your relational database provider does not provide a specific OLE DB data provider or a .NET data provider, you can use the generic Microsoft OLE DB provider to retrieve data. In this chapter you will be using a SQL Server database and hence you can use the Native OLE DB provider for SQL Server also called as the SQL Server Native Client. If you need to use the .NET data provider, you would select SqlClient Data provider.

To create a data source, follow these steps:

1. Select the Data Sources folder in the Solution Explorer.
2. Right - click the Data Sources folder and click New Data Source, as shown in Figure 2 - 14.

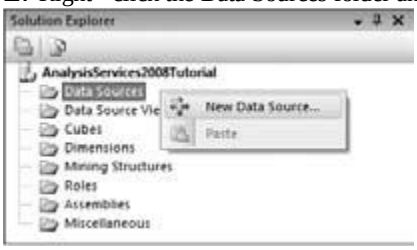

Figure 2-14

This launches the Data Source Wizard. This wizard is self - explanatory and you can easily create a data source by making the appropriate selection on each page of the wizard. The first page of the wizard is the welcome page that provides additional information about a data source. Click Next to continue.

3. You ' re now in the connection definition page of the Data Source Wizard, as shown in Figure 2- 15. In this page, you will provide the connection information about the relational data source that contains the " Adventure Works DW 2008 " database. Click the New button under Data Connection Properties to specify the connection details. The Connection Manager dialog box launches.

Figure 2-15

4. On the page shown in Figure 2 - 16 , specify the connection properties of the SQL Server containing the Adventure Works DW 2008 database. The provider used to connect to any relational database by default is the Native OLE DB\SQL Native Client 10.0 provider. If that provider is not selected, click the Provider drop - down and select SQL Server Native Client 10.0. If you have installed the SQL Server 2008 database engine and the Adventure Works DW 2008 sample database on the same machine, type **localhost** or the machine name in the Server name field as shown in Figure 2 - 16 . If you have restored the sample Adventure Works DW 2008 database on a different SQL Server machine, type that machine name instead. You can choose either Windows Authentication or SQL Server Authentication for connecting to the relational data source. Select Use Windows Authentication. If you choose SQL Server Authentication, you need to specify a SQL Server login name and password. Make sure you check the Save My Password option. Due to security restrictions in Analysis Services 2008, if you do not select this option you will be prompted to key in the password each time you send the definitions of your database to the Analysis Services instance. From the drop - down list box under Select or Enter a Database Name, select AdventureWorksDW2008. You have now provided all the details needed for establishing a connection to the relational data in Adventure Works DW 2008. Click OK.

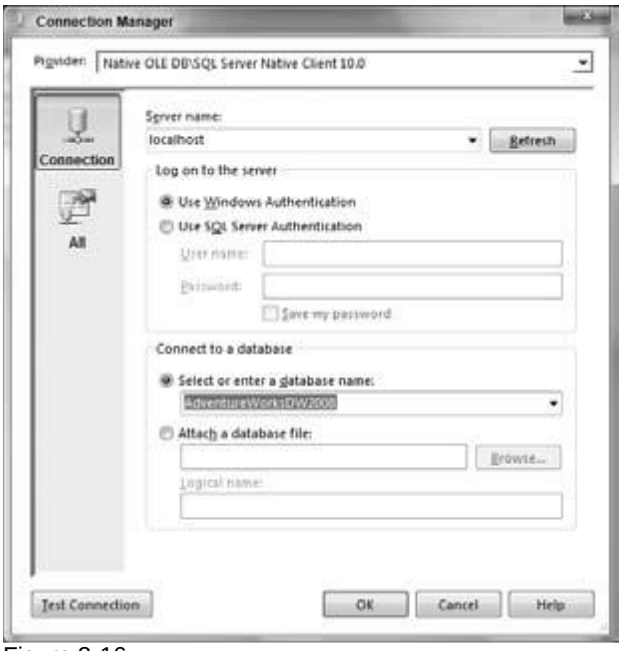

Figure 2-16

5. The connection properties you provided in the connection dialog are now shown in the " Select how to define the connection " page of the Data Source Wizard, as shown in Figure 2 - 17 . Click the Next button.

6. In the Impersonation Information page you need to specify the impersonation details that Analysis Services will use to connect to the relational data source. There are four options as shown in Figure 2 - 18 . You can provide a domain username and password to impersonate or select the Analysis Service instance ' s service account for connection. The option Use the credentials of the current user is primarily used for data mining where you retrieve data from the relational server for prediction. If you use the Inherit option, Analysis Services uses the impersonation information specified for the database. Select the Use the service account option and click Next.

7. On the final page, the Data Source Wizard chooses the relational database name you have selected as the name for the data source object you are creating. You can choose the default name specified or specify a new name here. Specify the name Adventure Works DW as shown in Figure 2 - 19 . The connection string to be used for connecting to the relational data source is shown under Preview. Click Finish.

Figure 2-19

Super! You have successfully created a data source.

Creating a Data Source View (DSV)

The Adventure Works DW database contains 25 tables. The cube you build in this chapter uses 10 tables. Data Source Views give you a logical view of the tables that will be used within your OLAP database. A Data Source View can contain tables and views from one or more data sources. Although you could accomplish the same functionality by creating views in the relational server, Data Source Views provide additional functionality, flexibility, and manageability, especially when you do not have privileges to create views on the relational backend.

To create a Data Source View, follow these steps:

1. Select the Data Source Views folder in the Solution Explorer.
2. Right - click Data Source Views and select New Data Source View, as shown in Figure 2 - 20 .

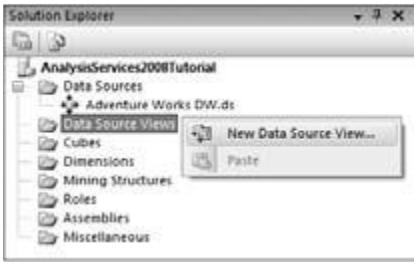

This launches the Data Source View Wizard. Similar to the Data Source wizard, this wizard allows you to create a Data Source View just by choosing an appropriate selection on each page of the wizard. Click the Next button to go to the next page of the wizard.

3. The second page of the DSV Wizard (see Figure 2 - 21) shows the list of data source objects from which you might want to create a view. The New Data Source button allows you to launch the Data Source Wizard so that you can create new data source objects from the wizard. You have already created a data source for the Adventure Works DW 2008 database that you will use for this example. Select this data source and click Next.

Figure 2-20
Figure 2-21

4. When you click the Next button, the DSV Wizard connects to the Adventure Works DW 2008 relational database using the connection string contained in the data source object. The wizard then retrieves all the tables, views, and relationships from the relational database and shows them in the third page. You can now select the tables and views that are needed for the Analysis Services database you are creating. For this tutorial, navigate through the Available Objects list and select the FactInternetSales and FactResellerSales tables. Click the > button so that the tables move to the Included Objects list. Select the two tables in the Included objects list. When you select these tables you will notice that the Add Related Tables button is enabled. This button allows you to add all the tables and views that have relationships with the selected tables in the Included objects list. Click the Add Related Tables button. You will notice that all the related dimension tables mentioned earlier as well as the FactInternetSalesReason table are added to the Included objects list. In this tutorial you will not be using the FactInternetSalesReason table, so you should remove this table. Select the FactInternetSalesReason table in the Included Objects list and click the < button to remove it from the Included Objects. You have now selected all the tables needed to build the cube in this tutorial. Your Included Objects list of tables should match what ' s shown in Figure 2 - 22 .

Figure 2-22

5. Click the Next button and you are at the final page of the DSV Wizard! Similar to the final page of the Data Source Wizard, you can specify your own name for the DSV object or use the default name. Specify the " Adventure Works DW " for the DSV Name in the wizard and click Finish.
You have now successfully created the DSV that will be used in this chapter. The DSV object is shown in the Solution Explorer and a new designer page is created in the main area of the BIDS as shown in Figure 2 - 23 . This is the Data Source View editor. The Data Source View editor contains three main areas: Diagram Organizer, the Tables view, and the Diagram view. The Diagram view shows a graphical representation of the tables and their relationships. Each table is shown with its columns with an indication of the key attribute. The connecting lines show the relationships between tables. If you double - click a connecting line you will find the columns of each table that are used to form the join that defines the relationship. You can make changes to the Data Source View by adding, deleting, or modifying tables and views in the DSV Designer. In addition, you can establish new relationships between tables. You learn more about the DSV Designer in Chapter 4

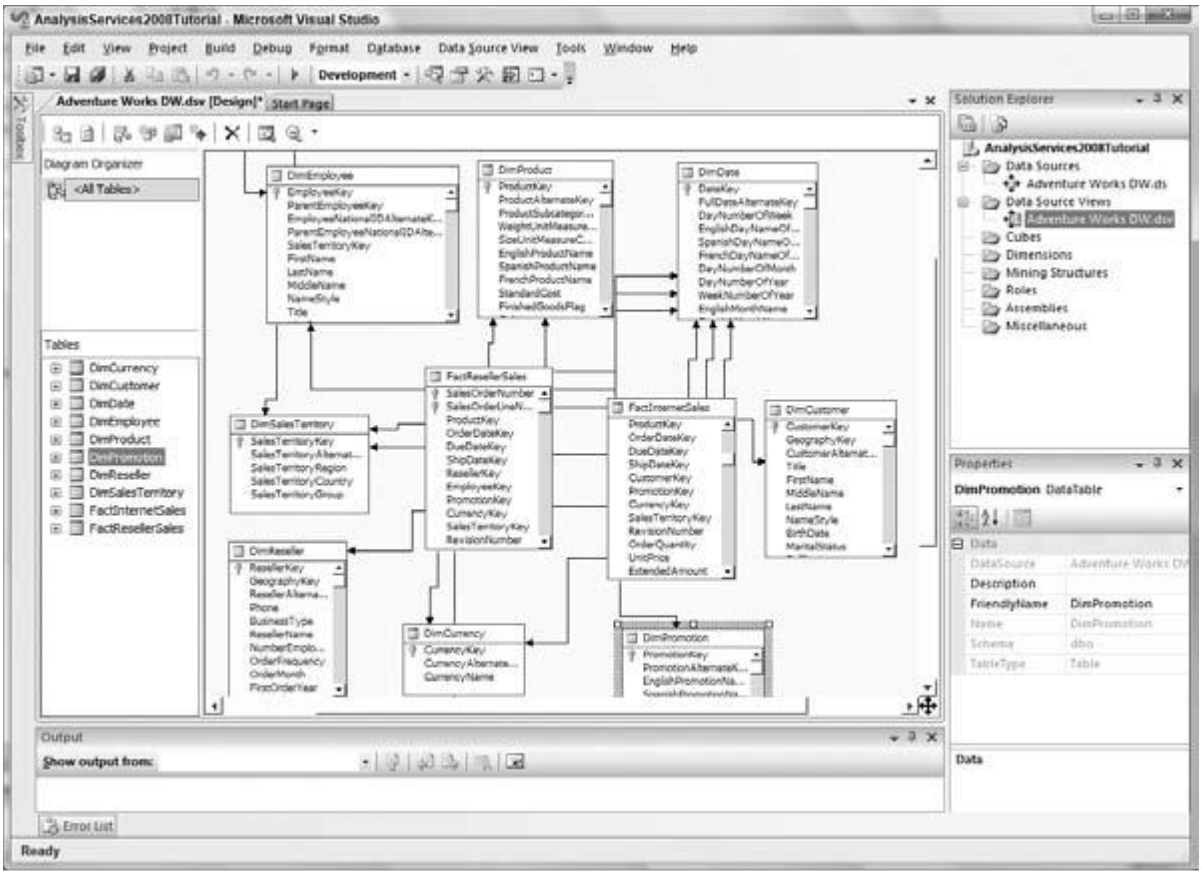

The number of tables you can see in the Diagram view depends on the resolution on your machine. In this view, you can zoom in to see a specific table enlarged or zoom out to see all the tables within the Diagram view. To use the zoom feature, you can right - click anywhere within the Diagram view, select Zoom, and set the zoom percentage you want. Alternatively, you can select View Zoom and then select the zoom percentage. Select a zoom percentage of 150%. Figure 2 - 24 shows a zoomed-in Diagram view so that you can see the FactResellerSales table clearly.

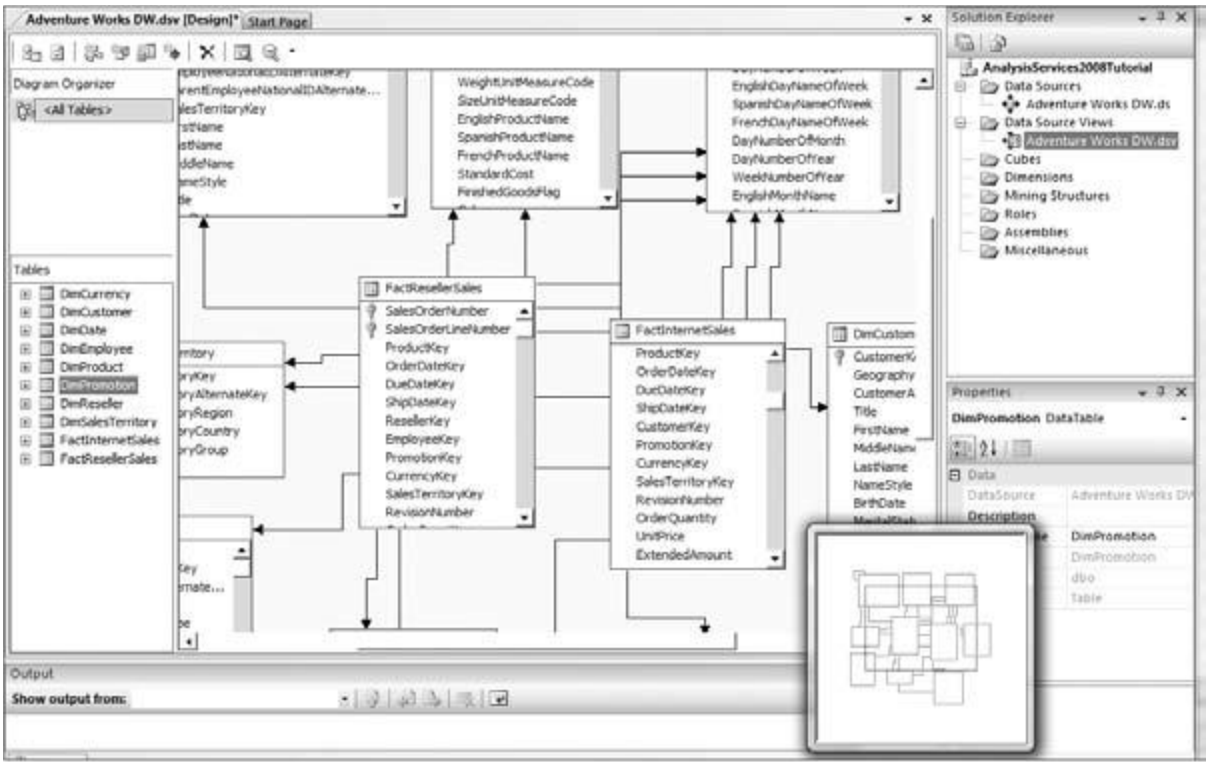

The Diagram view in the DSV arranges the tables to best fit within the view. Sometimes the number of tables in the DSV can be quite large. In such circumstances, navigating to the tables in the Diagram view can be difficult. For easier navigation you can use the Locator window (see Figure 2 - 24). The Locator window shows the full DSV diagram as a thumbnail. You can open it by performing a left mouse click on the 4 - headed arrow in the lower - right corner of the diagram, as highlighted in Figure 2 - 23 . The Locator window remains open while the mouse button is held down. This allows you to navigate through the visible area in the Diagram view by moving the mouse.

You have now learned the basic operations used within a Data Source View. Next, you move on to creating a cube using the Cube Wizard.

Creating a Cube Using the Cube Wizard

In Analysis Services 2008 you can build cubes via three approaches — top - down, bottom - up, or an empty cube. The traditional way of building cubes is bottom - up from existing relational databases. In the bottom - up approach, you need a Data Source View from which a cube can be built. Cubes within a project can be built from a single DSV or from multiple DSVs. In the top - down approach, you create the cube and then generate the relational schema based on the cube design. In Analysis Services 2008 you also have the option to first create an empty cube and then add objects to it. A cube in Analysis Services 2008 consists of one or more measure groups from a fact table (typically you will have one measure group per fact table) and one or more dimensions (such as Product and Time) from the dimension tables. Measure groups consist of one or more measures (for example, sales, cost, count of objects sold). When you build a cube, you need to specify the fact and dimension tables you want to use. Each cube must contain at least one fact table, which determines the contents of the cube. The facts stored in the fact table are mapped as measures in a cube. Typically, measures from the same fact table are grouped together to form an object called a measure group. If a cube is built from multiple fact tables, the cube typically contains multiple measure groups. Before building the cube, the dimensions need to be created from the dimension tables. The Cube Wizard packages all the steps involved in creating a cube into a simple sequential process:

1. Launch the Cube Wizard by right - clicking the Cube folder in the Solution Explorer and selecting New Cube.
2. Click the Next button in the welcome page.
3. You are now asked to select the method to create the cube. Choose the default value (Use existing tables) and click Next (see Figure 2 - 25).

Figure 2-25

4. In the Select Measure Group Tables page, select the Data Source View Adventure Works DW2008 as shown in Figure 2 - 26 .

5. The Suggest button helps you identify the measure group tables. If you click the Suggest button, the Cube Wizard will analyze the relationships between the tables in the Data Source View and select the potential measure group tables. For this example, select Fact Internet Sales and Fact Reseller Sales as measure groups as shown in Figure 2 - 27 and click Next.

6. The Select Measures page allows you to select specific columns from the measure group tables as measures as shown in Figure 2 - 28 . By default all the columns in the selected measure group tables except the key column are shown as measures and selected. Choose the default selection shown by the wizard and click Next.

Figure 2-28

7. In the Select New Dimensions page, the Cube Wizard shows you the potential dimensions along with their attributes. The Cube Wizard by default will include the key attribute in each dimension, which is highlighted in this page as shown in Figure 2 - 29 . Deselect the Fact Internet Sales and Fact Reseller Sales dimensions as shown in Figure 2 - 29 and click Next.

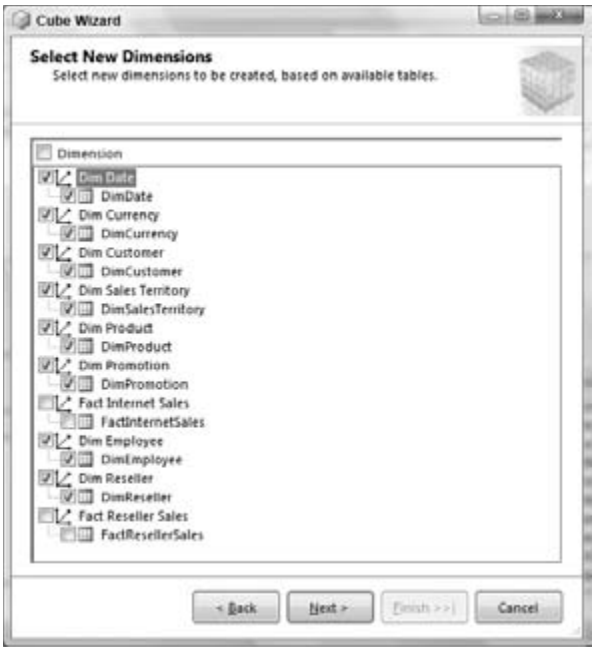

8. In the Final page of the Cube Wizard, provide the cube name Adventure Works DW as shown in Figure 2 - 30 and click Finish.

9. After the wizard completes you will notice that the Adventure Works DW cube and Dim Date, Dim Currency, Dim Customer, Dim Sales Territory, Dim Product, Dim Promotion, Dim Employee, and Dim Reseller dimensions are created in the Solution Explorer as shown in Figure 2 - 31.

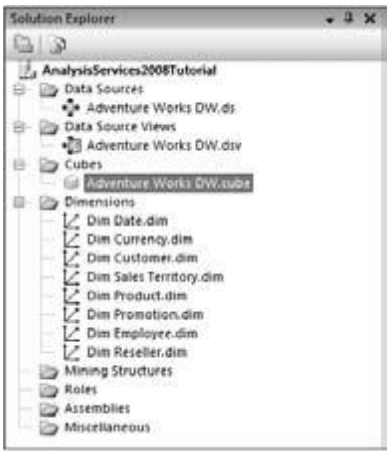

The Adventure Works DW cube is automatically opened in the Cube Editor, as shown in Figure 2 - 32 .

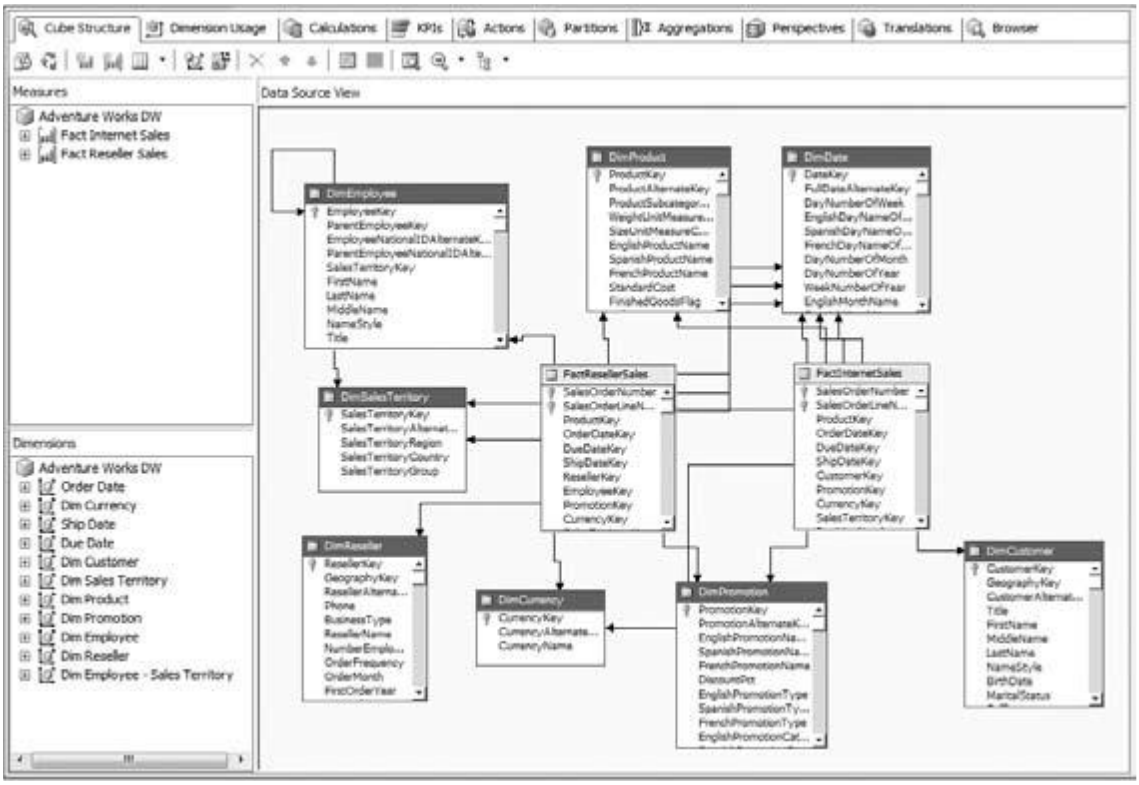

The Cube Editor has several panes that allow you to perform various operations on a cube object. The default pane upon completion of the Cube Wizard is the Cube Structure pane. Other panes of the Cube Editor are Dimension Usage, Calculation, KPIs, Actions, Partitions, Aggregations, Perspectives, Translations, and Browser. In this chapter you will become familiar with basic operations in the Cube Structure and the Browser panes. You learn more about the Cube Editor in Chapters 6 and 9 .

The Cube Structure pane is divided into three windows: Measures, Dimensions, and the Data Source View. If you need to add or modify measure groups or measures you will do that in the Measures window. The Dimensions window is used to add or modify the dimensions relevant to the current cube. The Data Source View shows all the fact and dimension tables used in the cube with appropriate colors (yellow for fact tables and blue for dimension tables). Actions such as zoom in, zoom out, navigation, finding tables, and different diagram layouts of the tables are available in the DSV of the Cube Editor.

If you right - click within the Measures, Dimensions, or Data Source View windows you will be

able to see the various actions that can be accomplished within each window. The actions within the Measures, Dimensions, or DSV windows of a Cube Editor can also be accomplished by clicking the appropriate buttons (see Figure 2 - 32) in the Cube Editor toolbar.
You have now successfully created a cube using Business Intelligence Development Studio. The Cube Wizard has only added the most essential attributes to the dimensions created. This is a change from Analysis Services 2005 to make sure the cube designer includes only the necessary attributes. The default dimensions created by the Cube Wizard need to be refined further in order to analyze the data in the cube. Because this is the first cube you are creating and we wanted to have very simple instructions, the following steps include most of the attributes from the dimensions. In reality when you create a cube based on your needs, you would usually include only the dimension attributes that are required. Continue with the following steps to refine the dimensions created by the Cube Wizard so that you can perform a simple analysis.

10. Double - click the Dim Date dimension (Dim Date.dim object) in the Solution Explorer.

11. You will now be in the Dimension Editor with the Dim Date dimension loaded. The Dimension Editor contains three panes: Attributes, Hierarchies, and Data Source View as shown in Figure 2 - 33 . Select all the columns in the DimDate table in the Data Source View except the key column Date Key.

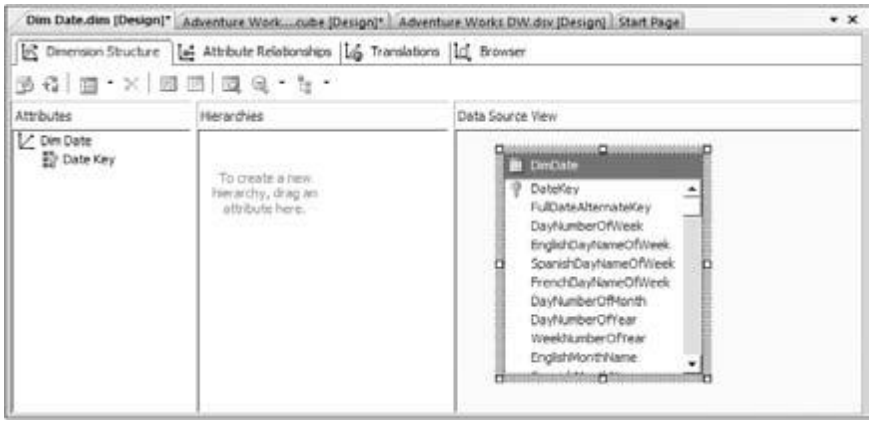

12. Drag and drop the selected columns to the Attributes pane. This action creates an attribute hierarchy for each of the columns in the DimDate table.

13. Rename the key attribute from Date Key to Dim Date.

14. Drag and drop Fiscal Quarter from the Attributes pane to the Hierarchies pane. This creates a new hierarchy called Hierarchy.

15. Drag and drop Month Number of Year onto the Hierarchies pane below the Fiscal Quarter. This creates a second level in the Hierarchy hierarchy.

16. Drag and drop the key attribute Dim Date onto the Hierarchies pane below the Month Number of Year.

17. Right - click the Hierarchy hierarchy and select Rename. Rename the hierarchy to Fiscal Quarter – Month Number Of Year. The Dimension Editor with the Dim Date dimension should appear as shown in Figure 2 - 34 .

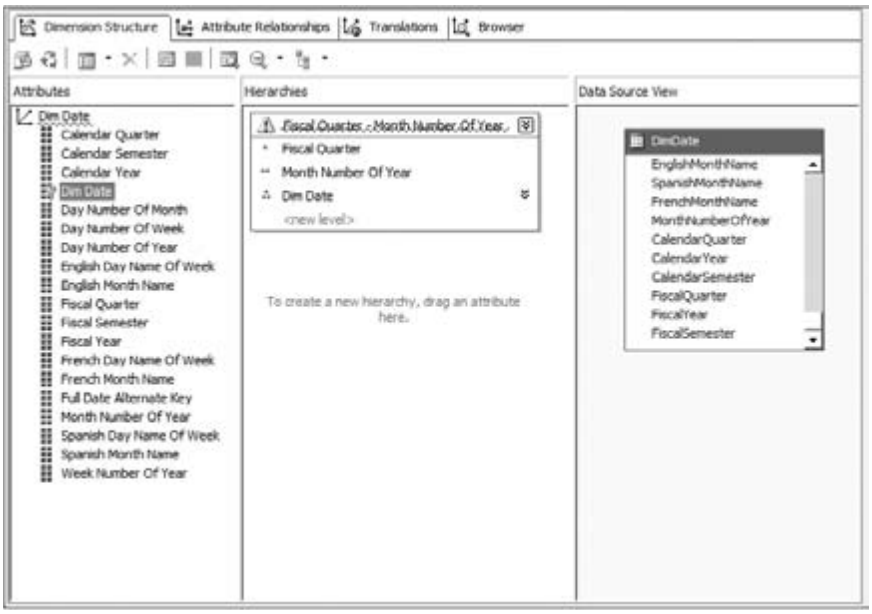

Figure 2-34

18. Double - click the Dim Currency dimension (Dim Currency.dim object) in the Solution Explorer.
19. Drag and drop Currency Alternate Key to the Attributes pane.
20. Rename the key attribute from Currency Key to Dim Currency.
21. Double - click the Dim Customer dimension (Dim Customer.dim object) in the Solution Explorer.
22. Rename the key attribute from Customer Key to Dim Customer.
23. Drag and drop all the columns except Customer Key from the DimCustomer table in the Data Source View pane to the Attributes pane.
24. Double - click the DimSalesTerritory dimension (Dim Sales Territory.dim object) in the Solution Explorer.
25. Drag and drop all the columns from the DimSalesTerritory table in the Data Source View pane except the key attribute SalesTerritoryKey to the Attributes pane.
26. Rename the key attribute from Sales Territory Key to Dim Sales Territory.
27. Double - click the Dim Product dimension (Dim Product.dim object) in the Solution Explorer.
28. Rename the key attribute from Product Key to Dim Product.
29. Drag and drop all the columns of the DimProduct table except ProductKey and LargePhoto from the Data Source View pane to the Attributes pane.
30. Double - click the Dim Promotion dimension (Dim Promotion.dim object) in the Solution Explorer.
31. Rename the key attribute from Promotion Key to Dim Promotion.
32. Drag and drop all the columns of the DimPromotion table except PromotionKey from the Data Source View pane to the Attributes pane.
33. Drag and drop English Promotion Category from the Attributes pane to the Hierarchies pane. This creates a new Hierarchy.
34. Drag and drop the attribute Discount Pct from the Attributes pane to the Hierarchies pane below English Promotion Category. This creates a new level in the Hierarchy hierarchy.
35. Drag and drop the key attribute Dim Promotion from the Attributes pane to the Hierarchies pane below the Discount Pct level.
36. Rename the hierarchy to English Promotion Category – Discount Pct. The Dimension Editor with the Dim Promotion dimension should look like Figure 2 - 35 .

Figure 2-35

37. Double - click the Dim Reseller dimension (Dim Reseller.dim) in the Solution Explorer.

38. Rename the key attribute from Reseller Key to Dim Reseller.

39. Drag and drop all the columns of the DimReseller table except ResellerKey from the Data Source View pane to the Attributes pane.

40. Drag and drop the Annual Revenue attribute from the Attributes pane to the Hierarchies pane. A new hierarchy with the name Hierarchy is created.

41. Drag and drop Number Employees from the Attributes pane to the Hierarchies pane under Annual Revenue. This creates a new level called Number Employees.

42. Drag and drop the Dim Reseller attribute from the Attributes pane to the Hierarchies pane under Number Employees.

43. Rename the hierarchy Hierarchy to Annual Revenue – Number of Employees. Your Dim Reseller Dimension Editor should look like Figure 2 - 36 .

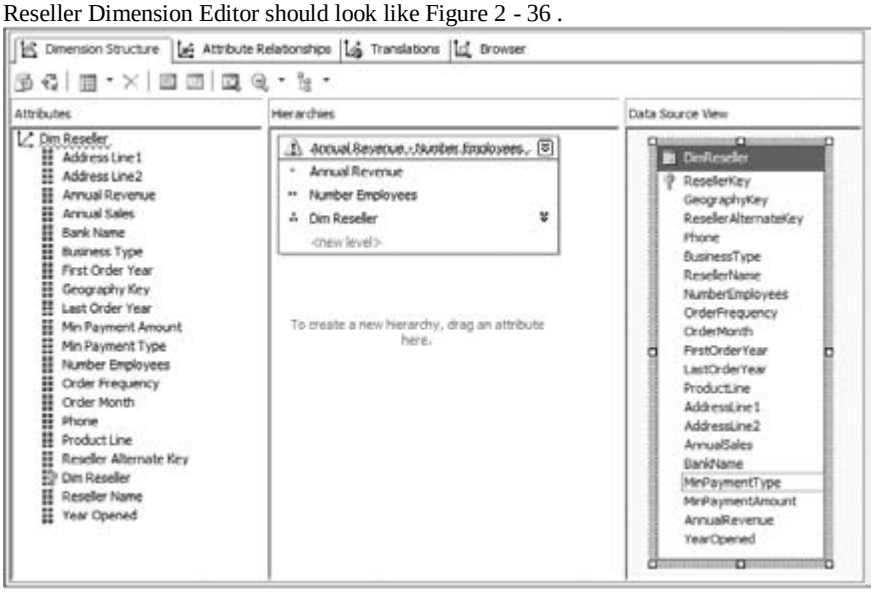

Figure 2-36

44. Double - click the Dim Employee dimension (Dim Employee.dim) in the Solution Explorer. This opens up the Dimension Editor with the Employee dimension loaded.

45. Notice that this dimension has three attributes created by the Cube Wizard compared to the single attribute created for all the other dimensions you opened. This is due to the fact that the Cube Wizard detected a parent - child relationship within the Dim Employee dimension. You learn more about parent - child dimensions in Chapter 5 .

46. Drag and drop all the columns in the DimEmployee table except the three attributes that have already been created by the Cube Wizard from the Data Source View pane to the Attributes pane.

47. Rename the key attribute from Employee Key to Dim Employee.

48. Drag and drop Department Name from the Attribute pane to the Hierarchies pane. This creates a new hierarchy called Hierarchy with a single level.

49. Drag and drop the Title attribute from the Attributes pane to the Hierarchies pane below Department Name.

50. Drag and drop the Dim Employee attribute from the Attributes pane to the Hierarchies pane below Title.

51. Rename the hierarchy to Department Name – Title.

dimensions in order to do simple analysis. You might have noticed warning symbols in the Dimension Editor for the dimensions where you created hierarchies. You learn more about these warnings, the creation of dimensions, attributes, hierarchies, and attribute relationships in Chapters 5 and 9 . All you have done, though, is create the structure of the cube. There has not been any interaction with the Analysis Services instance at this point. This method of creating the cube structure without any interaction with the Analysis Services instance is referred to as project mode. Using BIDS you can also create these objects directly on the Analysis Services instance. That method of creating objects on the Server is called online mode, which is discussed in Chapter 7 . Now you need to send the schema definitions of the newly created cube to the Analysis Services instance. This process is called *deployment* .

Deploying and Browsing a Cube

To deploy the database to the Analysis Server, right - click the project name and select Deploy, as shown in Figure 2 - 37 . You can also deploy the project to the server from the main menu in BIDS by selecting Debug Start or just by pressing the F5 function key on your keyboard.

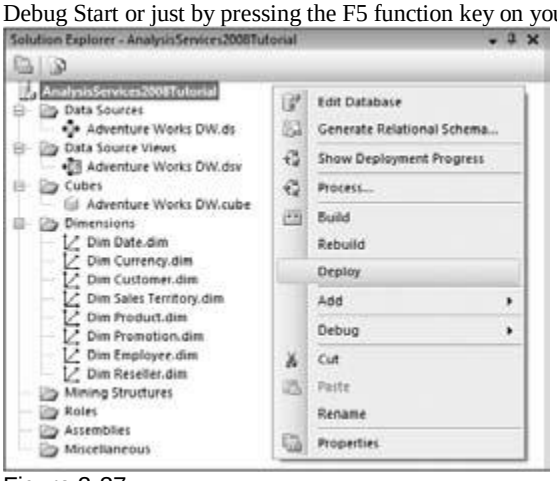

Figure 2-37

When you deploy an Analysis Service project, BIDS first builds the project you have created and checks for preliminary warnings and errors such as invalid definitions. If there are no errors with respect to project definitions, BIDS packages all the objects and definitions you have created in the project and sends them to the Analysis Services instance. By default these definitions are sent to the Analysis Services instance on the same machine (localhost). A database with the name of the project is created and all the objects defined in the project are created within this database. When deploying, BIDS not only sends all the schema definitions of the objects you have created, but also sends a command to process the database. If you want to deploy your project to a different machine that is running Analysis Services 2008, you need to right - click the project and select Properties. This brings up the Properties Pages dialog in which you can specify the Analysis Services instance to deploy the project to. This page is shown in Figure 2 - 38 . Change the Server property to the appropriate machine and follow the steps to deploy the project.

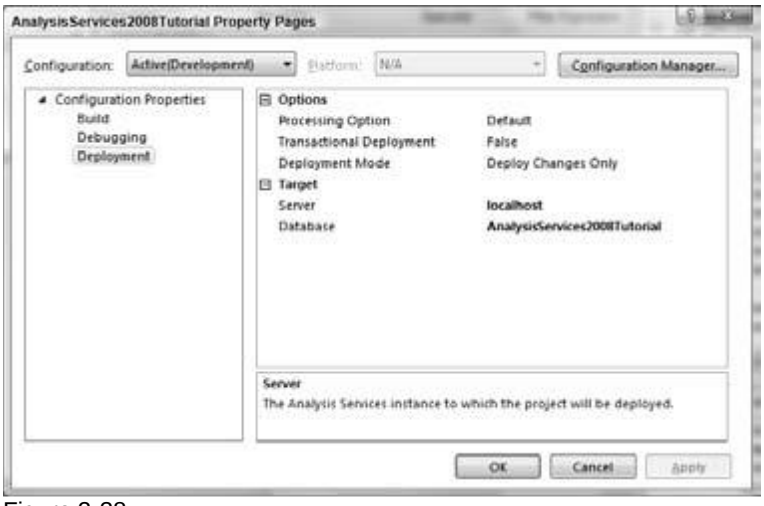

Figure 2-38

After you deploy the project you will see a Deployment Progress window at the location of the Properties window. The Output pane in BIDS shows the operations that occur after selecting Deploy — building the project, deploying the definitions, and the process command that is sent to the server. BIDS retrieves the objects being processed by the Analysis Services instance and shows the details (the object being processed; the relational query sent to the relational database to process that object including the start and end time; and errors, if any) in the Deployment Progress window. Once the deployment is completed, appropriate status will be shown in the Deployment Progress window as well as in the Output pane. If there were errors reported from the server these will be presented to you in the Output pane. You can use the Deployment Progress window to identify which object caused the error. BIDS waits for results from the server. If the deployment succeeded (successful deployment of schema and processing of all the objects), that information is shown as " Deploy: 1 succeeded, 0 failed, 0 skipped " . You will also notice the message " Deployment Completed Successfully " in the Deployment Progress window. If there are any errors reported from Analysis Services, deployment will fail and you will be prompted with a dialog box. The errors returned from the service will be shown in the Output pane. In your current project, deployment will succeed as shown in Figure 2 - 39 and you will be able to browse the cube.

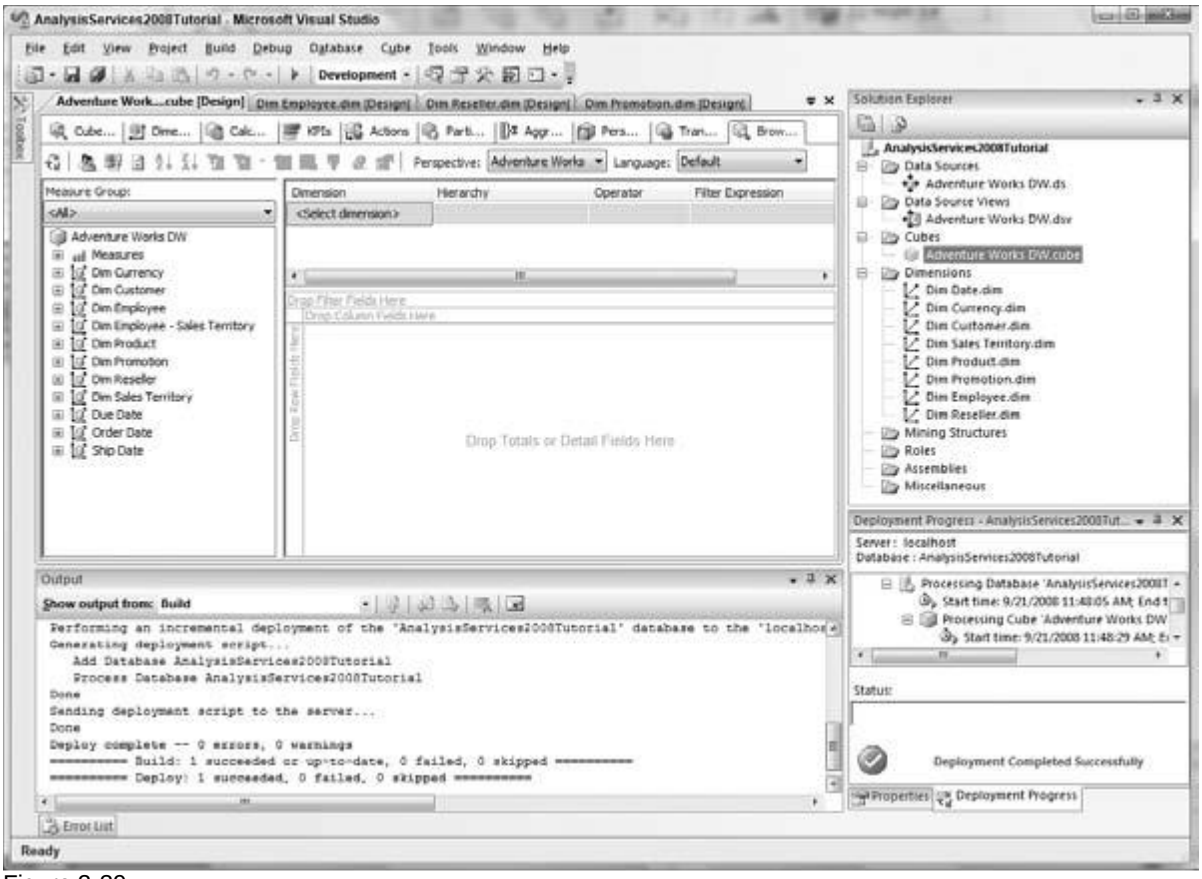

Figure 2-39

To browser your successfully deployed cube, use the following steps:

1. Double - click the Adventure Works DW cube to open the Cube Editor.

2. Switch to the Browser pane. The Browser pane has three main windows, as shown in Figure 2 - 40 . The left window shows the available measures and dimensions. This is called the Measure Group window. You can expand the tree nodes to see the measure groups, measures, dimensions, and hierarchies. On the right side, you have two windows split horizontally. The top pane is referred to as the Filter window because you can specify filter conditions to use while browsing the cube. The bottom pane hosts the Office Web Components (OWC) pivot table control, which is used for analyzing results. You can drag and drop measures and dimensions from the Measure Group pane to the OWC to analyze data.

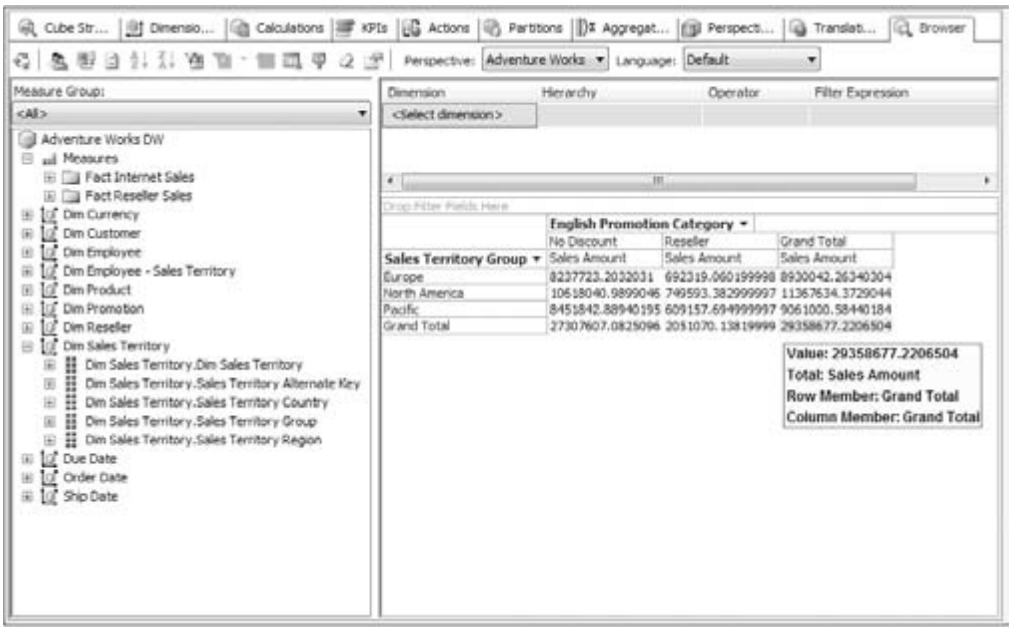

Figure 2-40

Drag and drop the " English Promotion Category " hierarchy of the " Dim Promotion " dimension and the " Sales Territory Group " hierarchy of the " Dim Sales Territory " dimension on to the Column and Row fields, respectively, of the OWC as shown in Figure 2 - 40 .

3. Drag and drop the Sales Amount measure from the Fact Internet Sales measure group to the Data area. You can similarly drag and drop multiple measures within the data area. You will now see the measure values that correspond to the intersection of the different values of the two hierarchies English Promotion Category and Sales Territory Group. As shown in Figure 2 - 40 you will notice " Grand Total " generated for each dimension along the Row and Column axes. The Grand Total values are retrieved by OWC by sending appropriate MDX queries to the server. Each measure value corresponding to the intersection of the dimension values is referred to as a cell. If you hover over each cell you will see a window that shows the properties of that cell. In Figure 2 - 40 you can see the basic cell properties for the cell at the intersection of English Promotion Category = Reseller and Sales Territory Group = North America shown in a tooltip.