Лекция 9

Тема « Sorting Members of a Level »

Sorting Members of a Level

Members of a level are the members of the attribute that defines that level. For example, the members of the level Country in the Geography hierarchy are actually the members of the attribute English Country Region Name. The member name that is shown in the Dimension Designer Browser is the text associated with the Name of the Country. It is not uncommon for dimension tables to have one column for the descriptive name and one column that is the key column of the table. You can use the descriptive name column to display the name of the attribute and the key column to sort the members in that attribute.
The attributes ' properties help you sort members of a level.
Each attribute in a dimension has two properties: KeyColumns and NameColumn. The KeyColumns property is used to specify the columns that are used for sorting the members, and the NameColumn is used for the descriptive name of the member. By default, the Dimension Wizard and the Dimension Designer set the KeyColumns attribute when an attribute is added to the dimension. They do not set the NameColumn property. If the NameColumn property is empty, Analysis Services will return the KeyColumns value for the descriptive names in response to client requests.
Figure 5 - 34 shows these properties for the attribute English Country Region Name (Country Level in the Geography multilevel hierarchy). The data type of the attribute is also shown in the KeyColumns property. Country is of data type WChar, which means the members are strings. Therefore, when you view the members in the Dimension Browser the members are sorted by the names. The Dim Geography dimension table has the column Country Region Code. You can define the sort order of the countries based on the Country Region Code instead of their names by changing the KeyColumns and NameColumn properties appropriately. The following exercise demonstrates how you can change the order of the countries based on the order of Country Region Code (AU, CA, DE, FR, GB, and US) instead of the country names.

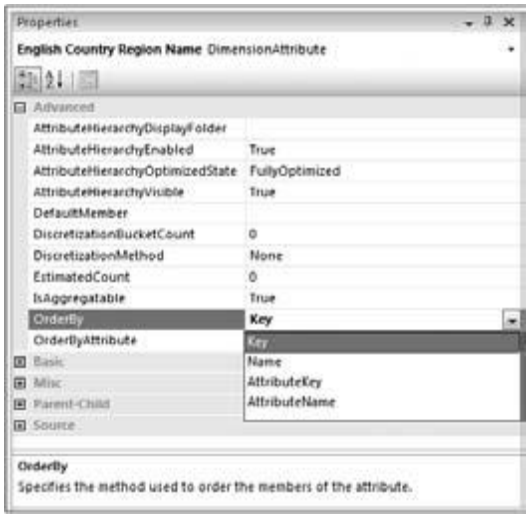

Figure 5-34

1. Click English Country Region Name in the Attributes pane; then in the Properties pane, click the NameColumn property value ellipses. This opens an Object Binding dialog showing all the columns in the Dim Geography table. Select the column EnglishCountryRegionName and click OK.
2. Click the KeyColumns property value (the ellipsis button). This action launches the Key Columns dialog. Remove the column EnglishCountryRegionName from the collection. In the list of available columns, select CountryRegionCode and add it to the Key Columns list. The Key Columns selection dialog should look like Figure 5 - 35 . Click the OK button.

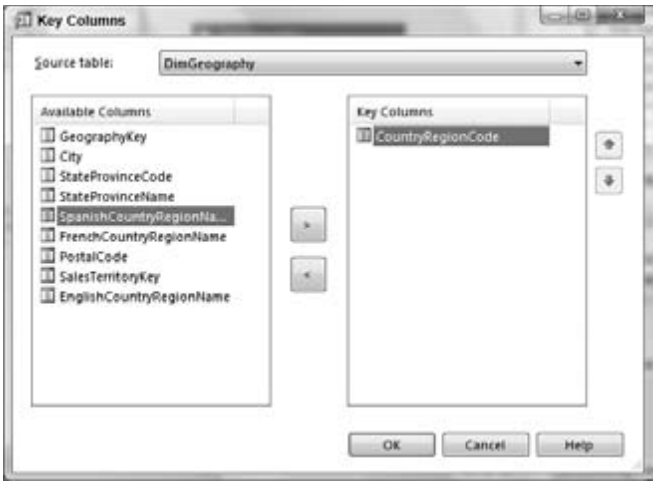

Figure 5-35

3. Click the Advanced Properties for the attribute EnglishCountryRegionName. Make sure the value of the property OrderBy is Key as shown in Figure 5 - 34 . This instructs the server to order this attribute using the Key attribute (CountryRegionCode), which you specified in step 2.

4. Deploy the project to the Analysis Services instance.

Deploying the project to the Analysis Services instance results in sending the new changes defined in steps 1 through 3 followed by processing the dimension. BIDS will switch to the Broswer tab. (If it doesn ' t, switch to the Browser tab and click the Reconnect option to retrieve the latest dimension data.) In the Dimension Browser select the Geography hierarchy. The order of the countries has now changed based on the order of Country Region Code (AU, CA, DE, FR, GB, and US followed by the Unknown members) instead of the country names you viewed in Figure 5 - 32 . The new order of countries is shown in Figure 5- 36.

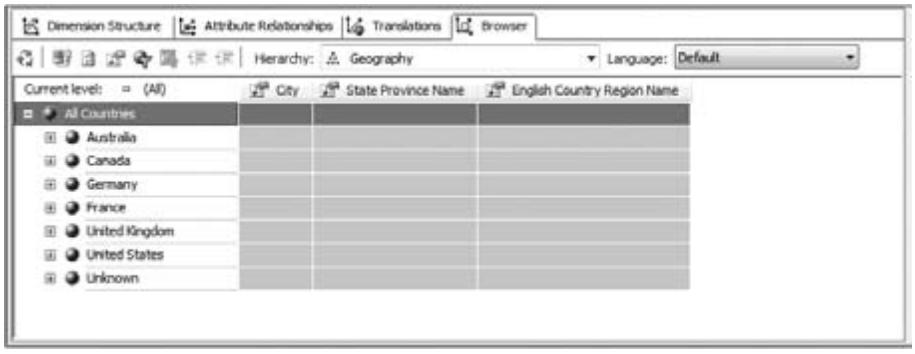

Figure 5-36

Optimizing Attributes

During the design of a dimension you might want to include certain attributes in the dimension, but not want to make the attribute hierarchies available to end users for querying. Two attribute properties allow you to manipulate visibility of attributes to end users. One property, AttributeHierarchyEnabled, allows you to disable the attribute. By setting this property to False you are disabling the attribute in the dimension; you cannot include this attribute in any level of a multilevel hierarchy. This attribute can only be defined as a member property (related attribute) to another attribute. Members of this attribute cannot be retrieved by an MDX query, but you can retrieve the value as a member property of another attribute. If you disable an attribute you might see improvements in processing performance depending on the number of members in the attribute. You need to be sure that there will be no future need to slice and dice on this attribute.

Another property called AttributeHierarchyVisible is useful for setting an attribute hierarchy to invisible for browsing; but even with this set, the attribute can be used as a level within a hierarchy and it can be used for querying. If you set this property to False, you will not see this attribute in the Dimension

Browser. The properties AttributeHierarchyEnabled and AttributeHierarchyVisible are part of the Advanced group in the Properties window, as shown in Figure 5 - 37 .

Figure 5-37

If you want to create a dimension that contains only multilevel hierarchies and no attributes, you can mark the AttributeHierarchyVisible property to False for all the attributes. When you go to the Dimension Browser you will only see the multilevel hierarchies. Even though you have disabled the attribute for browsing, you will still be able to query the attribute using MDX.

Defining Translations in Dimensions

If your data warehouse is to be used globally, you want to show the hierarchies, levels, and members in different languages so that customers in those countries can read the cube in their own language. Analysis Services 2008 provides you with a feature called Translations (not a super - imaginative name but a name that is intuitive) that helps you create and view dimension members in various languages. The benefit of this feature is that you do not have to build a new cube in every language. For the translation feature to be used, you need to only have a column in the relational data source that will have the translated names for the members of a specific attribute in the dimension.
For example, the Dim Geography table has two columns, Spanish Country Region Name and French Country Region Name, which contain the translated names of the countries that are members of the attribute English Country Region Name. The following steps describe how to create a new translation:

1. Switch to the Translations page in the Dimension Designer.
2. Click the New Translation toolbar button shown in Figure 5 - 38 or choose New Translation from the Dimension menu to create a new translation and choose a language. The Select Language dialog now pops up.

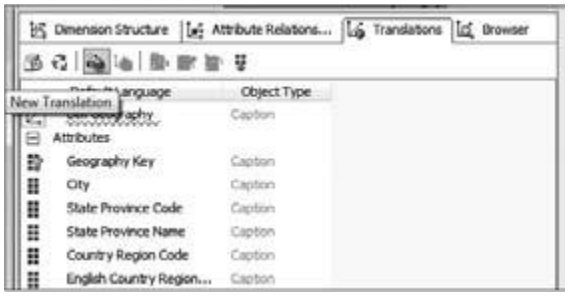

Figure 5-38

3. Select the language French (France) and click OK.

4. A new column with the title French (France) is added as shown in Figure 5 - 39 . Select the cell from the column French (France) in the English Country Region Name row. Then click the button that appears on the right side of the cell. You now see the Attribute Data Translation dialog.

5. Select the French Country Region Name column in the Translation Columns tree view as shown in Figure 5 - 40 and click OK.

6. Repeat steps 2 through 5 for the language Spanish (Spain).

You have now created two translations in French and Spanish languages. In addition to specifying the columns for member names, you can also change the metadata information of each level. For example, if you want to change the level Country in the Geography hierarchy in the French and Spanish languages, you can do that by entering the names in the row that show the Country level. Type **Pays** and **Pais** as shown in Figure 5 - 39 for French and Spanish translations, respectively. You have defined translations for the Country attribute in two languages making use of the columns in the relational data source. To see how this metadata information is shown in the Dimension Browser, first deploy the project to your Analysis Services instance.

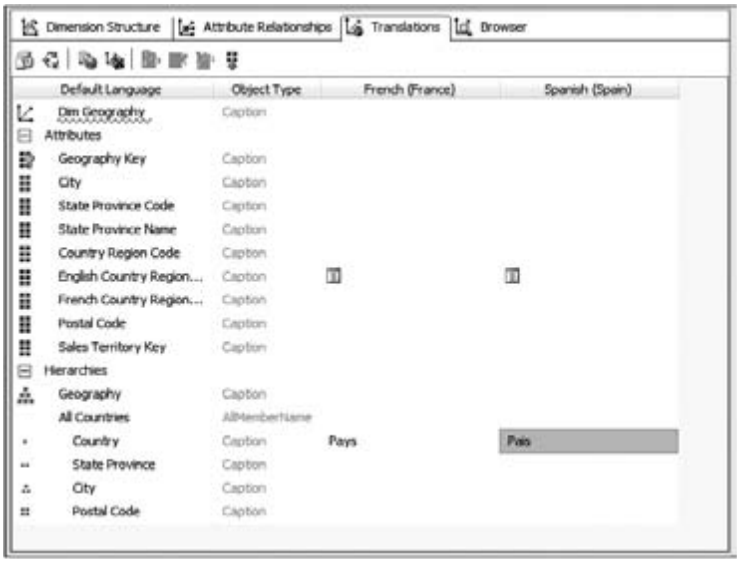

Figure 5-39

Figure 5-40

Next, to see the effect of the translations you have created, select language French (France) from within the Dimension Browser as shown in Figure 5 - 41 . Select the Geography hierarchy and expand the All level. Now you can see all the members in French. If you click any of the countries, the metadata shown for the level is " Pays " (French for country) as shown in Figure 5 - 39 . There is a negligible amount of overhead associated with viewing dimension hierarchies, levels, and members in different languages from your UDM.

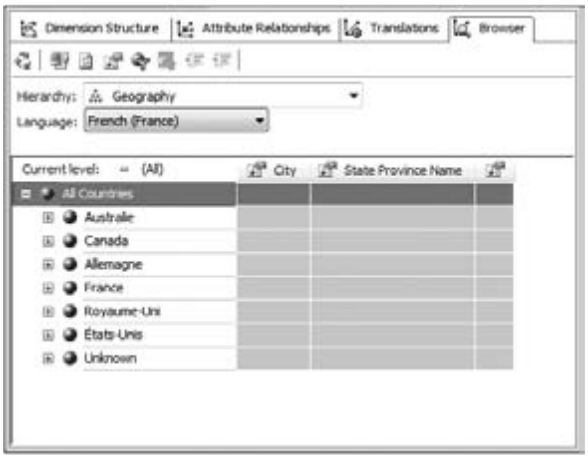

Figure 5-41

Creating a Snowflake Dimension

A snowflake dimension is a dimension that is created using a set of dimension tables. A snowflake dimension normally suggests that the tables in the data source have been normalized. Normalization is the process by which tables of a relational database are designed to remove redundancy and are optimized for frequent updates. Most database design books, including *The Data Warehouse Toolkit* by Ralph Kimball (Wiley, 1996) and *An Introduction to Database Systems* by C. J. Date (Addison Wesley, 2003), talk about the normalization process in detail.

The columns from different tables of a snowflake dimension often result in levels of a hierarchy in the dimension. The best way to understand a snowflake dimension is to create one yourself. To create one you ' re going to need two additional tables added to your DSV. Here is how to add the two tables:

1. Open the AdventureWorksDW DSV and click the Add/Remove Tables button (top - left button in the DSV).

2. Click DimProductCategory, then Control - Click DimProductSubcategory, then click the right arrow > to move the two tables from the source to the DSV and click OK.

The DSV Designer identifies the relationships defined in the relational backend and shows the relationships between the DimProduct, DimProductSubCategory, and DimProductCategory tables within the DSV Designer graphical design pane. Now that you have the necessary tables in the DSV and the relationships and logical primary keys defined, you can create a snowflake Product dimension. You can either delete the existing Product dimension in the AnalysisServices2008Tutorial project and create a snowflake dimension using the Dimension Wizard or refine the existing Product dimension and make it a snowflake dimension. In this illustration you will be refining the existing Product dimension and making it a snowflake dimension. Follow these steps to create a snowflake dimension called Dim Product:

1. Double - click the Dim Product dimension in the Solution Explorer to open the Dimension Designer for the Dim Product dimension.

2. Within the Data Source View pane of the Dimension Designer, right - click and select Show Tables as shown in Figure 5 - 42 .

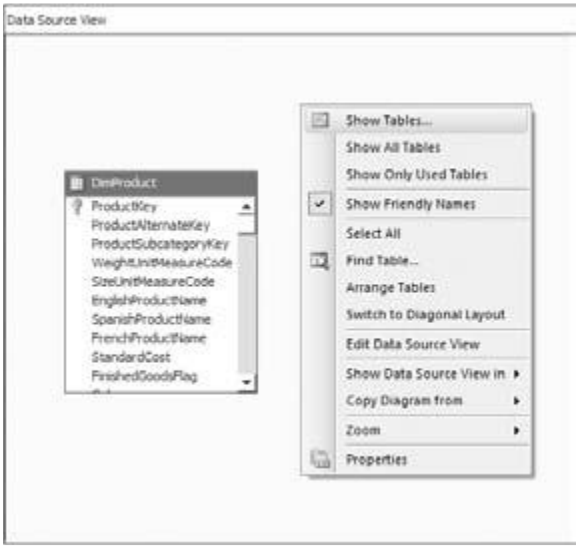

Figure 5-42

3. In the Show Tables dialog select the DimProductSubCategory and DimProductCategory tables as shown in Figure 5 - 43 and click OK.

Figure 5-43

You will now see the DimProductCategory and DimProductSubCategory tables added to the Data Source View pane of the Dimension Designer as shown in Figure 5 - 44 . Notice that the new tables added to the pane have a lighter colored caption bar. This indicates that none of the columns in the tables are included as attributes within the dimension.

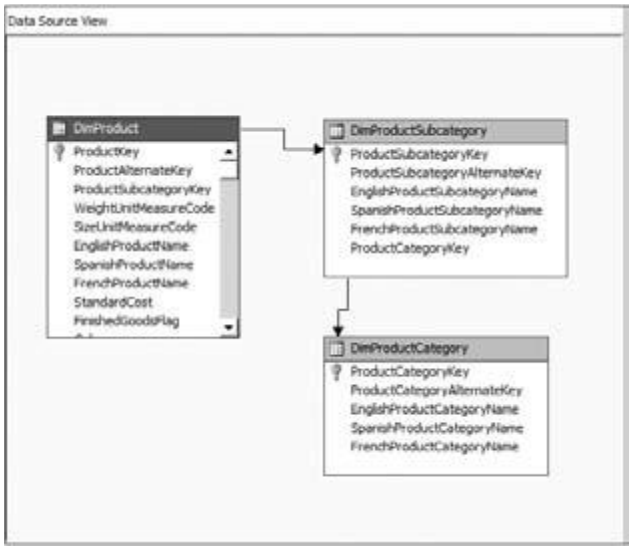

Figure 5-44

4. Drag and drop the column ProductCategoryKey from the DimProductSubCategory table in the DSV pane to the Attributes pane.
5. Launch the Name Column dialog for the ProductCategoryKey attribute by clicking the ellipsis next to the Name Column property in the Properties window.
6. Select EnglishProductCategoryName from the DimProductCategory table as the Name Column and click OK.
7. Select the attribute ProductSubCategoryKey from the Attributes pane.
8. Launch the Name Column dialog for the ProductSubCategoryKey attribute by clicking the

ellipsis next to the Name Column property in the Properties window.

9. Select EnglishProductSubCategoryName from the DimProductSubCategory table as the Name Column and click OK.

10. Launch the Name Column dialog for Dim Product (the key attribute) by clicking the ellipsis next to the Name Column property in the Properties window.

11. Select English Product Name as the Name Column and click OK.

12. Create a Product Categories multilevel hierarchy with levels ProductCategoryKey, ProductSubCategoryKey, and Dim Product by dragging and dropping the attributes to the Hierarchies pane and naming the hierarchy as Product Categories.

13. Rename the level ProductCategoryKey to ProductCategory.

14. Rename the level ProductSubCategoryKey to ProductSubCategory.

15. Rename the level Dim Product to Product Name.

16. Change the EnglishProductName attribute to Product Name.

17. Figure 5 - 45 shows the Dimension Designer after all the refinements to the Product dimension.

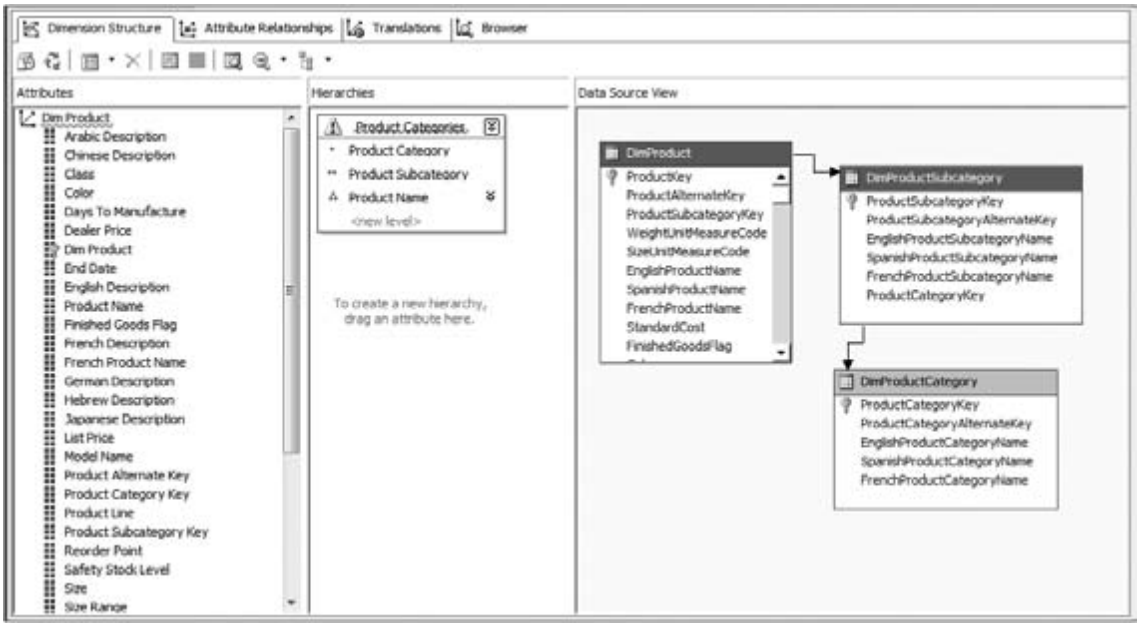

Figure 5-45

You have now successfully created a snowflake Dim Product dimension. You can perform most of the same operations in a snowflake dimension as you can in a star schema dimension, including adding attributes, creating hierarchies, and defining member properties. We recommend you deploy the AnalysisServices2008Tutorial project and browse the snowflake dimension Dim Product.

Creating a Time Dimension

Almost every data warehouse will have a Time dimension. The Time dimension can be comprised of the levels Year, Semester, Quarter, Month, Week, Date, Hour, Minute, and Seconds. Most data warehouses contain the levels Year, Quarter, Month, and Date. The Time dimension helps in analyzing business data across similar time periods; for example, determining how the current revenues or profit of a company compare to those of the previous year or previous quarter.

Even though it appears that the Time dimension has regular time periods, irregularities often exist. The number of days in a month varies across months, and the number of days in a year changes each leap year. In addition to that, a company can have its own fiscal year, which might not be identical to the calendar year. Even though there are minor differences in the levels, the Time dimension is often viewed as having regular time intervals. Several MDX functions help in solving typical business problems related to analyzing data across time periods. ParallelPeriod is one such function, which you learned

about in Chapter 3 . Time dimensions are treated specially by Analysis Services and certain measures are aggregated across the Time dimension uniquely and are called semi - additive measures. You learn more about semi - additive measures in Chapters 6 and 9 .
The AnalysisServices2008Tutorial project has a Dim Date dimension that was created by the Cube Wizard in Chapter 2 . Even though the dimension has been created from the Dim Date table, it does not have certain properties set that would allow Analysis Services to see it as the source for a Time dimension. In the following exercise you first delete the Dim Date dimension and then create a Time dimension. Follow these steps to create a Time dimension on the Dim Date table of the AdventureWorksDW2008 database:

1. In the Solution Explorer right - click the Dim Date dimension and select Delete.
2. In the Delete Objects and Files dialog, Analysis Services requests you to confirm the deletion of corresponding Cube dimensions (you learn about Cube dimensions in Chapter 9). Select OK to delete the Dim Date dimension.
3. Launch the Dimension Wizard by right - clicking Dimensions in the Solution Explorer and selecting New Dimension. When the welcome screen of the Dimension Wizard opens up, click Next.
4. In the Select Creation Method page of the wizard, select the " Use an existing table " option and click Next.
5. In the Specify Source Information page, select DimDate as the main table from which the dimension is to be designed and click Next.
6. In the Select Dimension Attributes page, in addition to the Date Key attribute, enable the checkboxes for the following attributes: Calendar Year, Calendar Semester, Calendar Quarter, English Month Name, and Day Number Of Month.
7. Set the Attribute Type for the " Calendar Year " attribute to Date Calendar Year as shown in Figure 5 - 46 .

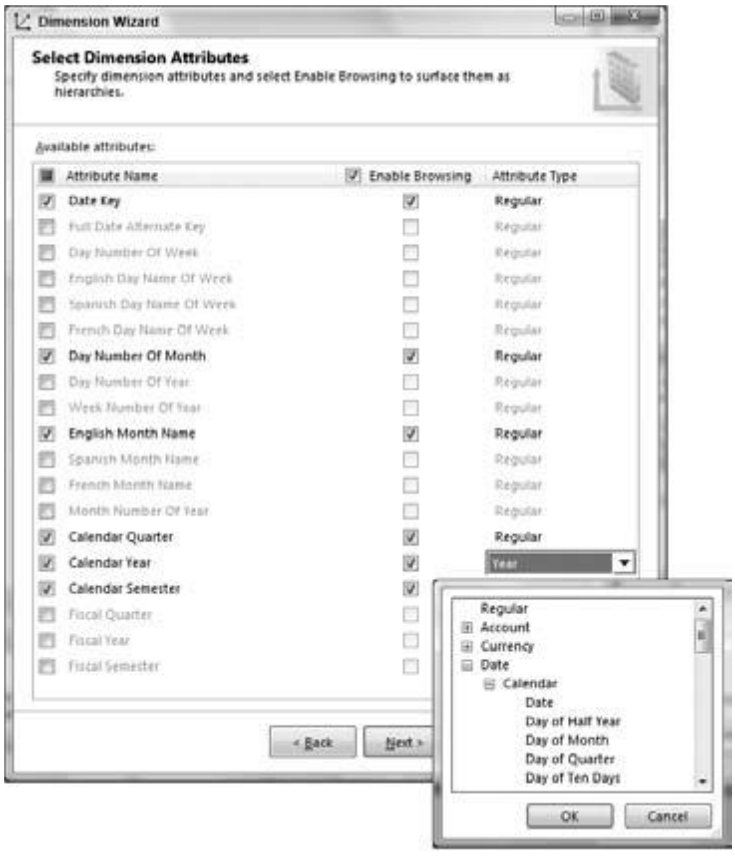

Figure 5-46

8. Set the Attribute Type for the remaining enabled attributes so they match those shown in Figure 5 - 47 and click Next to continue.

Figure 5-47

9. Set the name of the dimension to " Dim Date " and click Finish to close the Dimension Wizard. You have now successfully created a Time dimension using the Dimension Wizard.

10. Create a multilevel hierarchy Calendar Date with the levels Calendar Year, Calendar Semester, Calendar Quarter, Month (rename English Month Name), and Day (rename Day Number Of Month).

11. Save the project and deploy it to the Analysis Services instance.

12. Switch to the Browser pane of the Dim Date dimension.

Figure 5 - 48 shows the Calendar Date hierarchy that you created. Notice that the order of months within a quarter is not the default calendar order. For example, the order of months of CY Q1 of year 2002 is February, January, and March. To change the order, change the KeyColumns, NameColumn, and SortOrder appropriately and redeploy the project. We recommend that you define the necessary attribute relationships and attribute key values as defined by your business needs.

Figure 5-48

You have now successfully created a Time dimension. If you review the properties of the Dim Date dimension you will see the property " Type " set to Time which indicates to Analysis Services that the Dim Date dimension is a Time dimension. If you review the basic properties of each attribute in the Dim Date dimension, you will notice that the property Type has values such as Quarters, HalfYears, Years, DayOfMonth, and Months. You can use the Properties pane to set the right property type for the chosen attribute. Setting the right property type is important because a client application could use this property to apply the MDX functions for a Time dimension.

Creating a Parent - Child Hierarchy

In the real world you come across relationships such as that between managers and their direct reports. This relationship is similar to the relationship between a parent and child in that a parent can have several children and a parent can also be a child, because parents also have parents. In the data warehousing world such relationships are modeled as a Parent - Child dimension and in Analysis Services 2008 this type of relationship is modeled as a hierarchy called a Parent - Child hierarchy. The key difference between this relationship and any other hierarchy with several levels is how this relationship is represented in the data source. Well, that and certain other properties that are unique to the Parent - Child design. Both of these are discussed in this section.

When you created the Geography dimension, you might have noticed that there were separate columns for Country, State, and City in the relational table. Similarly, the manager and direct report can be modeled by two columns, ManagerName and EmployeeName, where the EmployeeName column is used for the direct report. If there are five direct reports for a manager, there will be five rows in the relational table. The interesting part of the Manager - DirectReport relationship is that the manager is also an employee and is a direct report to another manager. This is unlike the columns City, State, and Country in the Dim Geography table.

It is probably rare at your company, but employees can sometimes have new managers due to reorganizations. The fact that an employee ' s manager can change at any time is very interesting when you want to look at facts such as sales generated under a specific manager, which is the sum of sales generated by the manager ' s direct reports. A dimension modeling such a behavior is called a slowly changing dimension because the manager of an employee changes over time. You can learn slowly changing dimensions and different variations in detail in the book *The Microsoft Data Warehouse Toolkit: With SQL Server 2005 and the Microsoft Business Intelligence Toolset* by Joy Mundy et al. (Wiley, 2006).

The DimEmployee table in AdventureWorksDW has a Parent - Child relationship because it has a join from ParentEmployeeKey to the EmployeeKey. You have already created a DimEmployee dimension in the AnalysisServices2008Tutorial project in Chapter 2 using the Cube Wizard. In the following exercise you refine the existing Dim Employee dimension and learn how to create a dimension with a Parent - Child hierarchy using the Dimension Wizard. Note that you will actually be refining, not creating, the Dim Employee dimension in the illustration.

1. Launch the Dimension Wizard by right - clicking Dimensions in the Solution Explorer and selecting New Dimension. If the welcome screen of the Dimension Wizard opens up, click Next.
2. Make sure the " Use an existing table " option is selected and click Next.

3. In the Specify Source Information page, select DimEmployee as the main table from which the dimension is to be designed and click Next.

4. On the Select Related Tables screen, uncheck the DimSalesTerritory table and click Next.
In the Select Dimensions Attributes dialog, the Dimension Wizard has detected three columns of the DimEmployee table to be included as attributes. The Dimension Wizard will select columns if they are either the primary key of the table or a foreign key of the table or another table in the DSV. Figure 5 - 49 shows two of the attributes. The attributes suggested by the Dimension Wizard in this example are the key attribute Employee Key, the parent - child attribute Parent Employee Key, and the Sales Territory Key, which is a foreign key column to the DimSalesTerritory table.

5. Select all the columns of the DimEmployee table as attributes and click Next.

6. Notice in the preview pane of the Completing the Wizard dialog that the Parent Employee Key attribute has a unique icon (see Figure 5 - 50) indicating that Analysis Services detected a parent - child relationship in the DimEmployee table. The wizard was able to identify the parent - child relationship due to the join within the same table in the DSV.

7. Click the Cancel button because you will not be creating another DimEmployee dimension.

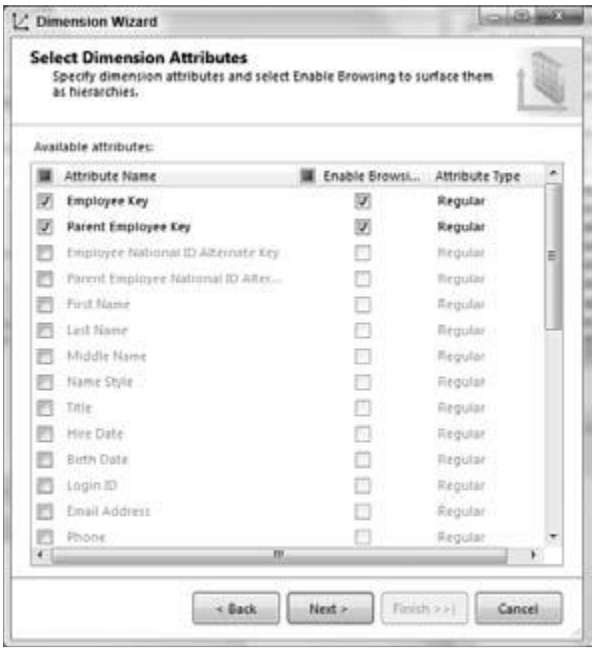

Figure 5-49

Figure 5-50

By default the Dimension Wizard defines the properties for the attribute modeling the Parent - Child hierarchy at the completion of the Dimension Wizard or the Cube Wizard.

8. Double - click the DimEmployee dimension in the Solution Explorer to open the Dimension Designer.

9. See the properties of the Parent Employee Key attribute that indicate that this attribute defines a Parent - Child hierarchy as shown in Figure 5 - 51 .

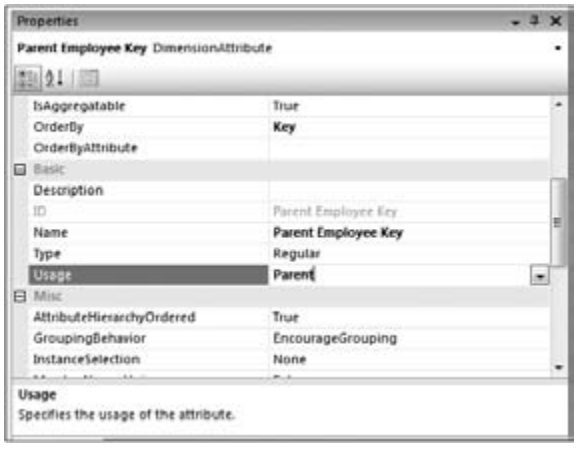

Figure 5-51

Notice that the hierarchy doesn ' t appear in the Hierarchies pane of the Dimension Designer. That ' s because the Parent - Child hierarchy is actually a special type of attribute hierarchy that can contain multiple levels, unlike the other attributes. The Parent - Child hierarchy that the wizard created is on the attribute ParentEmployeeKey. The Usage property for this attribute is set to Parent, which indicates that this attribute is a Parent - Child hierarchy. If you browse the Parent - Child hierarchy of the DimEmployee dimension, you will notice that you see the IDs of parent and employee as a multilevel hierarchy as seen in Figure 5 - 52 .

Figure 5-52

Typically, you would want to see the names of the employees rather than their IDs. You learned earlier that you can use the named column to specify the name that is shown in the browser and use the key column for ordering. Because the Parent - Child hierarchy retrieves all the information from the Key attribute, which is the DimEmployee attribute in this example, you need to modify the named column of the DimEmployee attribute rather than the named column of the Parent - Child hierarchy attribute.

10. Change the NameColumn property of the Key attribute Dim Employee to LastName and deploy the project to your Analysis Services instance.

When you browse the Parent - Child hierarchy, you will see the members of the hierarchy showing the last names of the employees, as shown in Figure 5 - 53 .

Figure 5-53

Summary

Using the Dimension Wizard and other wizards in BIDS is only the starting point for designing objects in

Analysis Services 2008. For optimal results, you will need to fine - tune what those wizards produce. At the beginning of each episode of the television serial *The Outer Limits* the viewers were exhorted to not adjust their television set. Just the opposite is true with Analysis Services 2008, because you do need to refine your objects created by the Dimension and Cube Wizards. A couple of examples are using the Properties window to assign descriptive names to an attribute that might otherwise harbor some obscure name coming from a source database and defining attribute relationships to optimize dimension performance. More profoundly, you can use the Dimension Designer to create translations for the attributes and hierarchies of a dimension into another language.

In addition to learning about dimensions, you learned the necessity of deploying your dimension to the instance of Analysis Services where the dimension is processed by retrieving the data from the data source. Processing is essential to enable the user to browse a dimension. The communication between BIDS and an instance of Analysis Services is accomplished through a SOAP - based XML API called XMLA (XML for Analysis), which is an industry standard. Even more interesting is that dimensions stored in Analysis Services are represented internally as cubes — one - dimensional cubes; and, what a coincidence, cubes are the topic of Chapter 6 .