Лекция 8

Тема « Working with the Dimension Designer »

Working with the Dimension Designer

The Dimension Designer, shown in Figure 5 - 10 , is an important tool that helps you to refine the dimension created by the Dimension Wizard. You can define the properties such as unary operators, custom roll - ups, and so forth, which help you to define how data should be aggregated for cells referred to by members of hierarchies in the dimension. The Dimension Designer itself is composed of four pages, which can be accessed from the tabs at the top of the designer: Dimension Structure, Attribute Relationships, Translations, and Browser. The first of these pages, Dimension Structure, contains three panes: Attributes, Hierarchies, and Data Source View. In addition to that you have the toolbar, which contains several buttons that help you to enhance the dimension. The Attributes pane shows all the attributes, the Hierarchies pane shows all the hierarchies along with their levels, and the Data Source View pane shows the tables that are used in the dimension. If you hover over each toolbar button you will see a tooltip that describes the functionality of that button. Some of the buttons are the same as the ones you saw in the DSV Designer and are used for operations within the Dimension Designer ' s Data Source View pane. The functionality of some of the other buttons is discussed later in this chapter and in Chapter 8.

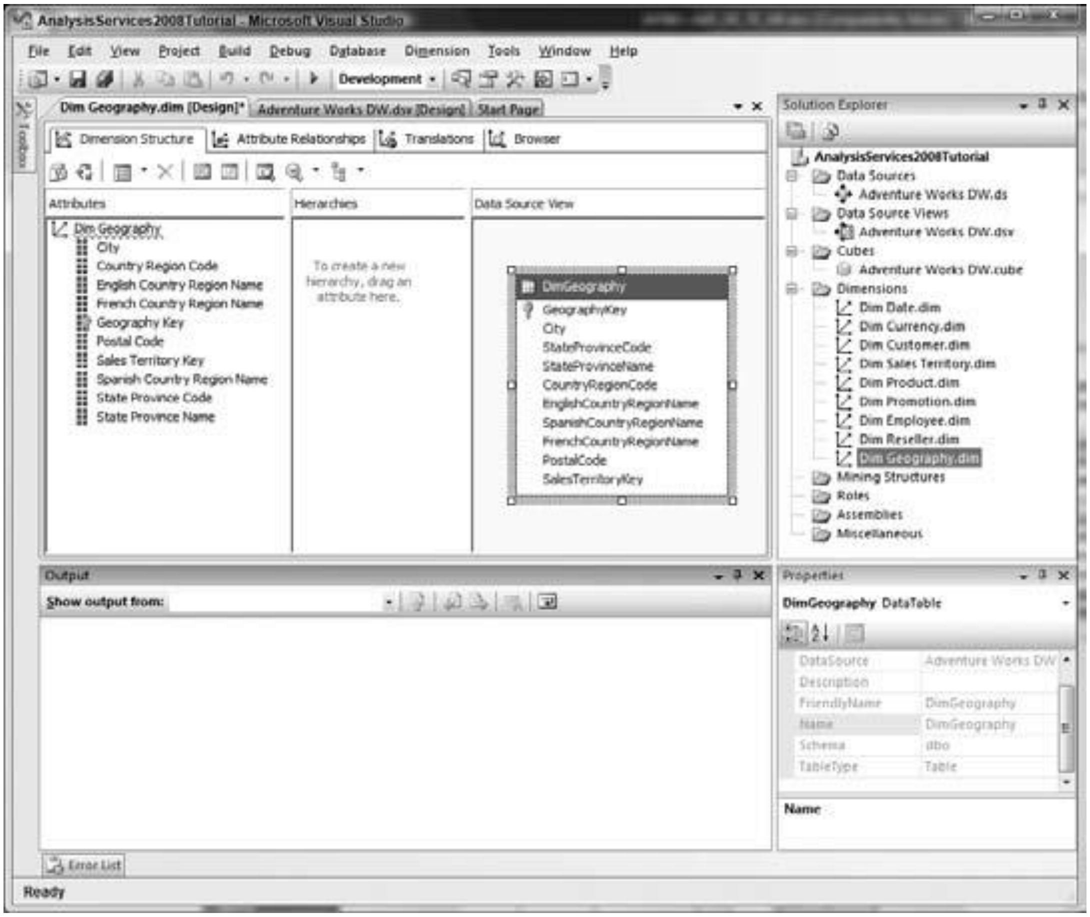

DataSource	Adventure Works DW
Description	
FriendlyName	DimGeography
Name	DimGeography
Schema	dbo
TableType	Table

Figure 5-10

Attributes

Attributes are hierarchies that have only two levels: the leaf level, which contains one member for each distinct attribute value, and the All level, which contains the aggregated value of all the leaf level members. The All level is optional. Each attribute directly corresponds to a table ' s column in the

DSV. The Attributes pane in the Dimension Designer shows all the attribute hierarchies of the dimension. The default view of all the attributes within the Attributes pane window is a Tree view as shown in Figure 5 - 11 . Two additional views are supported in the Dimension Designer: List view and Grid view. These views show the attributes and associated properties in different ways.

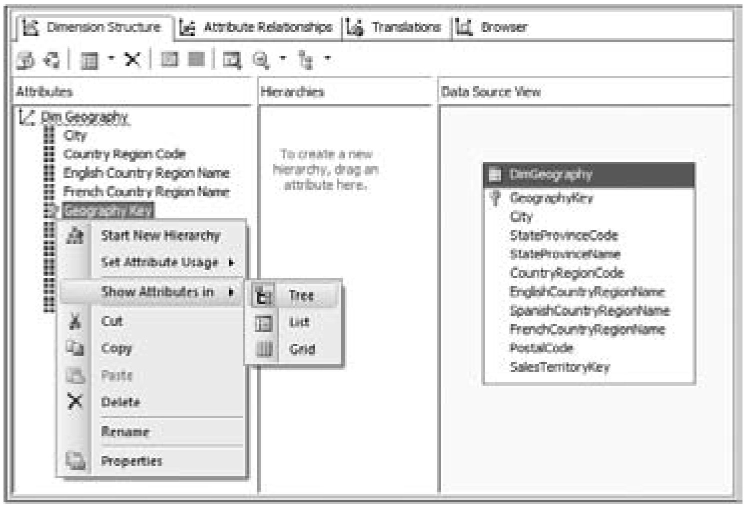

Figure 5-11

The List view repositions the Attributes pane below the Hierarchies pane and shows only the attributes of the dimension in a flat list (it doesn ' t show the dimension name as a parent node). This view is useful when your dimension has a lot of multilevel hierarchies. Because you get a wider area for the Hierarchies pane, you get a visually optimized view for this kind of dimension.

The Grid view is laid out similarly to the List view but includes additional columns that allow you to easily edit some of the important dimension properties right in the Attributes pane. When you ' re working with more than one attribute, editing these properties in the Attributes pane is less cumbersome than having to select each attribute and then switching over to the Properties window to change the attribute ' s value. (All the properties shown in the Grid view are also present in the Properties window.)

You can toggle between the different views by right - clicking in the Attributes pane and selecting the view type you desire, as shown in Figure 5 - 11 . Just choose the view that best suits you to visualize and design your dimension easily.

Figure 5 - 12 shows the List view and Grid view of the attributes shown in Figure 5 - 11.

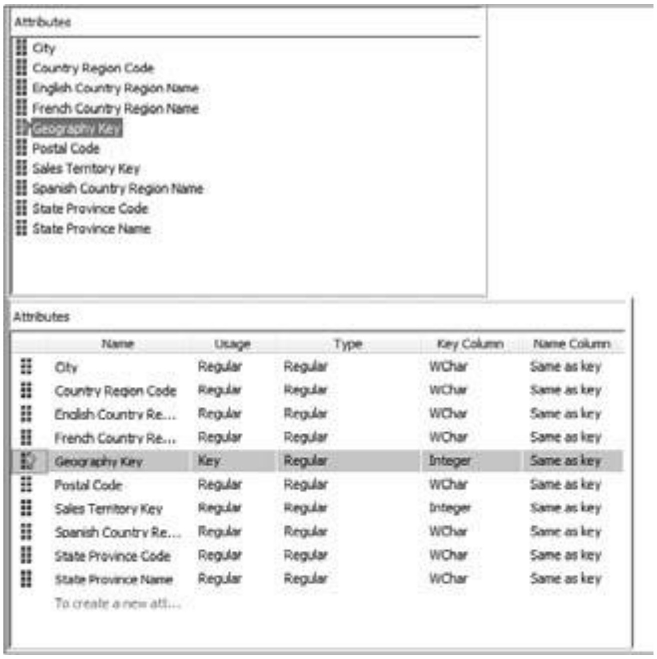

Attributes

- City
- Country Region Code
- English Country Region Name
- French Country Region Name
- Geography Key
- Postal Code
- Sales Territory Key
- Spanish Country Region Name
- State Province Code
- State Province Name

Attributes

Name	Usage	Type	Key Column	Name Column
City	Regular	Regular	WChar	Same as key
Country Region Code	Regular	Regular	WChar	Same as key
English Country Re...	Regular	Regular	WChar	Same as key
French Country Re...	Regular	Regular	WChar	Same as key
Geography Key	Key	Regular	Integer	Same as key
Postal Code	Regular	Regular	WChar	Same as key
Sales Territory Key	Regular	Regular	Integer	Same as key
Spanish Country Re...	Regular	Regular	WChar	Same as key
State Province Code	Regular	Regular	WChar	Same as key
State Province Name	Regular	Regular	WChar	Same as key
To create a new att...				

Figure 5-12

Attribute Relationships

Attribute relationships can be defined when attributes within the same dimension have a one - to - many relationship with each other. For example, if you have the attributes Country, State, and City, you have one - to - many relationships between country and state, as well as between state and city. Each dimension has to have at least one attribute that is defined as the key attribute. By definition, the key attribute has a one - to - many relationship with every attribute in the dimension. The Dimension Wizard automatically establishes relationships, such that all attributes of the dimension are related to key attributes.

If you are aware of one - to - many relationships between attributes, we highly recommend that you specify these relationships in the Dimension Designer with attribute relationships. Specifying the attribute relationship helps improve query performance as well as changing the aggregation design for multilevel hierarchies so as to include the attributes that are part of a hierarchy. Because the Dim Geography dimension contains one - to - many relationships, you need to specify attribute relationships to get query performance improvements. You learn more about the benefits of attribute relationships in Chapter 14 .

If you are familiar with Analysis Services 2005, you will notice that attributes in the Tree view no longer expand to show attribute relationships for the attributes in the tree. Analysis Services 2008 has added a separate page in the Dimension Designer to make definition of attribute relationships easier for the user.

To view and edit the attribute relationships in Analysis Services 2008, you use the Dimension Designer ’ s Attribute Relationships page shown in Figure 5 - 13 . The Attribute Relationships page contains three panes. The top pane graphically shows attribute relationships, the Attributes pane on the lower left shows the list of attributes in the dimension, and the Attribute Relationships pane in the lower right shows the list of defined relationships. The attributes shown in the top pane below the Geography Key attribute are all the attributes of the dimension that have one - to - many relationships with the key attribute. All these attributes are also referred to as member property attributes or related attributes.

Because the attribute relationship is defined for each member in the Geography Key attribute, you can retrieve its properties from the related attributes using member property MDX functions.

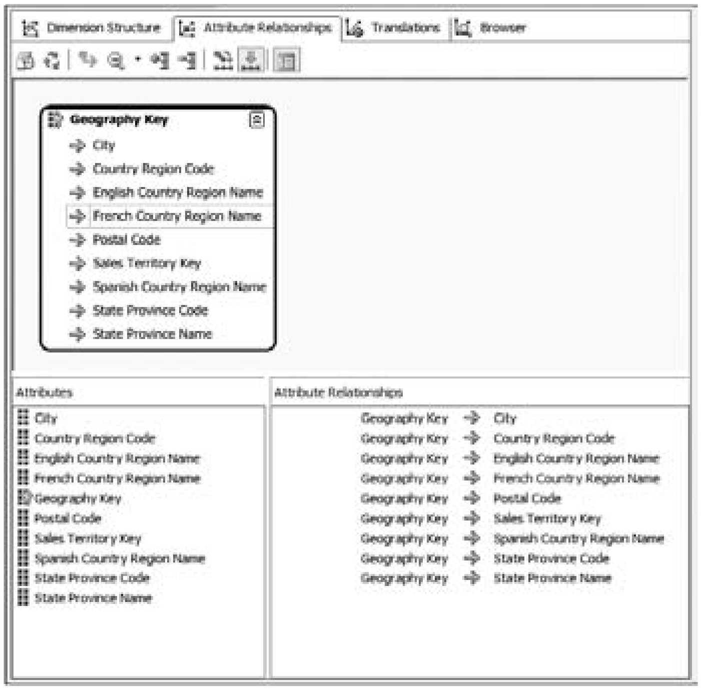

Attributes

- City
- Country Region Code
- English Country Region Name
- French Country Region Name
- Geography Key
- Postal Code
- Sales Territory Key
- Spanish Country Region Name
- State Province Code
- State Province Name

Attribute Relationships

Geography Key	→	City
Geography Key	→	Country Region Code
Geography Key	→	English Country Region Name
Geography Key	→	French Country Region Name
Geography Key	→	Postal Code
Geography Key	→	Sales Territory Key
Geography Key	→	Spanish Country Region Name
Geography Key	→	State Province Code
Geography Key	→	State Province Name

Figure 5-13

You can create new or modify existing attribute relationships using various methods within each of these panes. Follow these steps to update the Geography dimension' s attribute relationships:

1. In the visualization (top) pane, modifying attribute relationships is accomplished by dragging and dropping member property attributes onto the attribute to which they are related.
For example, if a State - Province Name has a one - to - many relationship with City, you would create the relationship by dragging the City attribute onto the State - Province Name attribute as follows: In the Attribute Relationships page ' s visualization pane, select the City attribute and drag and drop it onto the State - Province Name attribute. This creates a new attribute relationship node as shown in Figure 5 - 14 . Note the change in the Attribute Relationships pane to reflect the newly defined relationship. In the visualization pane, editing and deleting attribute relationships is accomplished by right - clicking a relationship ' s line and selecting the desired action from the context menu.

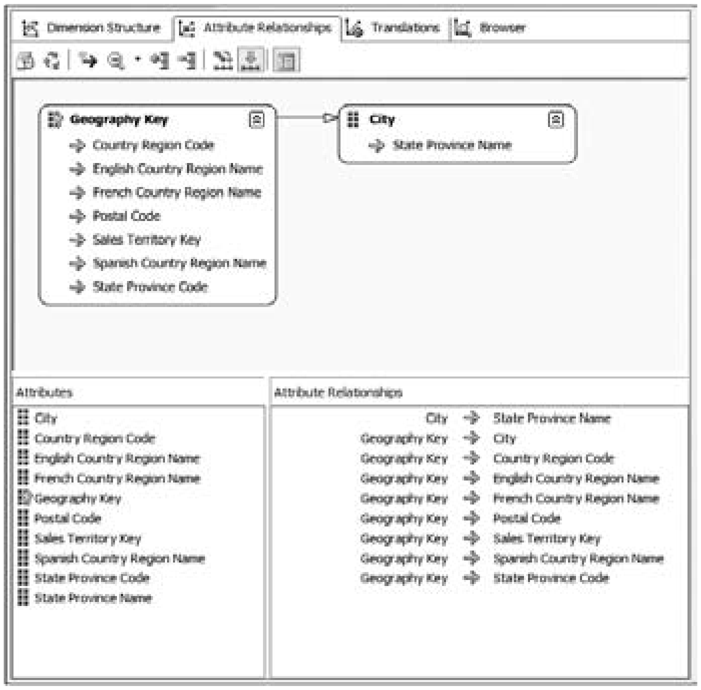

Geography Key
- Country Region Code
- English Country Region Name
- French Country Region Name
- Postal Code
- Sales Territory Key
- Spanish Country Region Name
- State Province Code

City
- State Province Name

Attributes

- City
- Country Region Code
- English Country Region Name
- French Country Region Name
- Geography Key
- Postal Code
- Sales Territory Key
- Spanish Country Region Name
- State Province Code
- State Province Name

Attribute Relationships

City	→	State Province Name
Geography Key	→	City
Geography Key	→	Country Region Code
Geography Key	→	English Country Region Name
Geography Key	→	French Country Region Name
Geography Key	→	Postal Code
Geography Key	→	Sales Territory Key
Geography Key	→	Spanish Country Region Name
Geography Key	→	State Province Code

Figure 5-14

2. To define a new relationship using the Attributes pane, you right - click the attribute that makes up the many side of the one - to - many relationship and select New Attribute Relationship. This will launch the Create Attribute Relationship dialog. The Source Attribute is the attribute corresponding to the " many " side of the relationship, and the Related Attribute is the attribute corresponding to the " one " side. You can also set the relationship type in this dialog to either Flexible (default) or Rigid. By default, Analysis Services tools define all the relationships to be flexible. A relationship is flexible if the value can change over time. A relationship is rigid if the relationship does not change over time. For example, the birth date of a customer is fixed and hence the relationship between a customer ' s key/name and the birth date attribute would be defined as rigid. However, the relationship between a customer and his city is flexible because the customer can move from one city to another. You can ' t delete or alter existing relationships from the Attributes pane. You have a one - to - many relationship between English Country Region Name and State - Province Name. To specify that the attribute English Country Region Name has a one - to - many relationship or is a member property of State - Province Name, perform the following steps in the Attribute Relationships tab of the Dimension Designer.

3. Right - click the State - Province Name attribute (source attribute) in the Attributes pane of the Attribute Relationships page and select New Attribute Relationship from the context menu.

4. In the Create Attribute Relationship dialog select the English Country Region Name as the Related Attribute (see Figure 5 - 15).

5. Click OK to create the relationship.

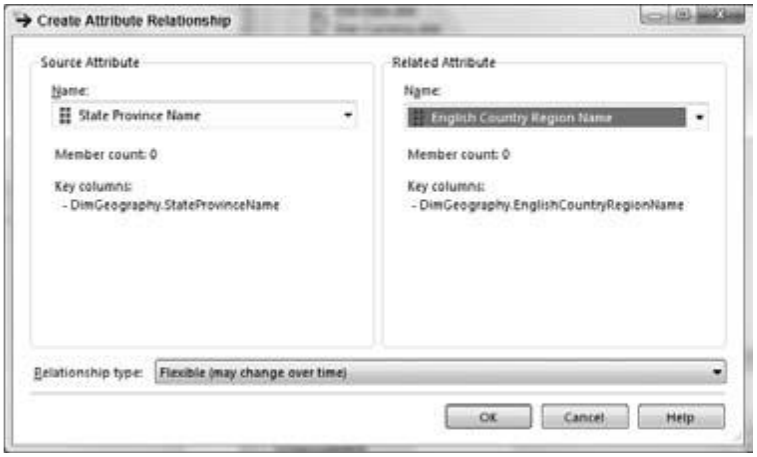

Create Attribute Relationship

Source Attribute

Name: State Province Name

Member count: 0

Key columns:
- DimGeography.StateProvinceName

Related Attribute

Name: English Country Region Name

Member count: 0

Key columns:
- DimGeography.EnglishCountryRegionName

Relationship type: Flexible (may change over time)

Figure 5-15

Your Attribute Relationships display should be similar to that shown in Figure 5 - 16 . Establishing attribute relationships is important for two reasons: It can improve processing performance (as you learn in Chapter 14), and it affects calculations that are aggregated across these attributes. You can view and modify the cardinality of the attribute relationships you establish using the Cardinality property in the Properties window. Click the relationship in the visualization pane or in the Attribute Relationships pane to bring up its properties. By default the Cardinality is set to many. If you know that the relationship between the attributes is one - to - one you can change the cardinality to one. For example, the cardinality between a customer ’ s ID and social security number is one - to - one, however the cardinality between the English Country Region Name attribute and the State - Province Name attribute is one - to - many.

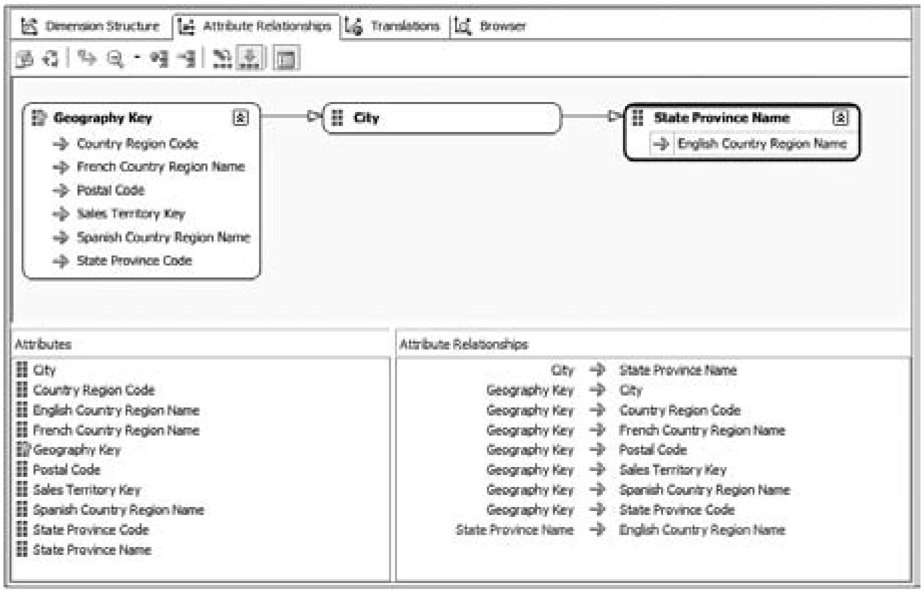

Attributes

- City
- Country Region Code
- English Country Region Name
- French Country Region Name
- Geography Key
- Postal Code
- Sales Territory Key
- Spanish Country Region Name
- State Province Code
- State Province Name

Attribute Relationships

```
               City  →  State Province Name
      Geography Key  →  City
      Geography Key  →  Country Region Code
      Geography Key  →  French Country Region Name
      Geography Key  →  Postal Code
      Geography Key  →  Sales Territory Key
      Geography Key  →  Spanish Country Region Name
      Geography Key  →  State Province Code
State Province Name  →  English Country Region Name
```

Figure 5-16

To define a new relationship using the Attribute Relationships pane you must first select an attribute from either the visualization or Attributes pane and then right - click in the Attribute Relationships pane in the area not occupied by existing attribute relationships. This can be a bit cumbersome. We recommend you use the visualization pane or the Attributes pane for creating new attribute relationships. Editing and deleting existing relationships in the Attribute Relationships pane, however, is as simple as right - clicking the relationship and choosing the desired action from the context menu as shown in Figure 5 - 17 . Editing a relationship launches the Relationship Editor dialog, whose functionality and layout is identical to the Create Attribute Relationships dialog shown in Figure 5 - 15 ; only the title is different.

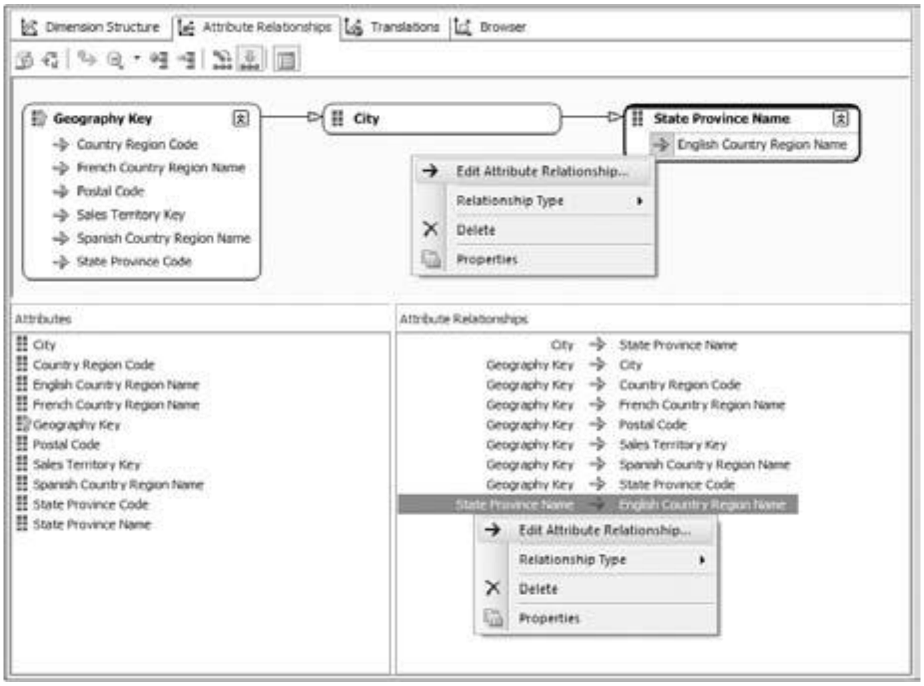

Attributes
City
Country Region Code
English Country Region Name
French Country Region Name
Geography Key
Postal Code
Sales Territory Key
Spanish Country Region Name
State Province Code
State Province Name

Attribute Relationships
City → State Province Name
Geography Key → City
Geography Key → Country Region Code
Geography Key → French Country Region Name
Geography Key → Postal Code
Geography Key → Sales Territory Key
Geography Key → Spanish Country Region Name
Geography Key → State Province Code
State Province Name → English Country Region Name

Figure 5-17

Use the Attribute Relationships pane to edit the existing relationships for the French and Spanish Country Region Name attributes using the following steps:

6. Click the Geography Key to French Country Region Name relationship in the Attribute Relationships pane.

7. Right - click and select Edit Attribute Relationship.

8. In the Edit Attribute Relationship dialog select the English Country Region Name as the Source Attribute (see Figure 5 - 18).

9. Click OK to save the change to the relationship.

10. In the Properties window, change the Cardinality property corresponding to the relationship between English Country Region Name and French Country Region Name from Many to One.

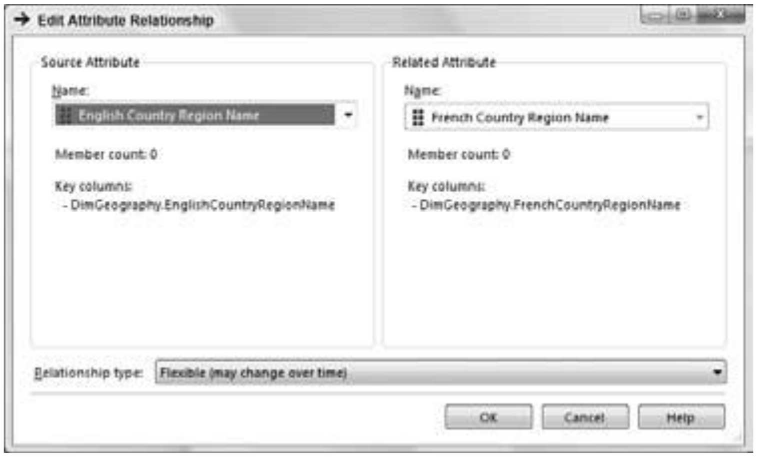

Source Attribute
Name: English Country Region Name
Member count: 0
Key columns:
- DimGeography.EnglishCountryRegionName

Related Attribute
Name: French Country Region Name
Member count: 0
Key columns:
- DimGeography.FrenchCountryRegionName

Relationship type: Flexible (may change over time)

Figure 5-18

Repeat the previous five steps for the Spanish Country Region Name attribute relationship.
You have now used three different methods for working with attribute relationships. Often in business

analysis when you are analyzing a specific member of a dimension, you need to see the properties of the dimension member to understand it better. In such circumstances, instead of traversing the complete hierarchy you can retrieve the member by querying the member properties. This once again is a performance improvement from the end user ' s perspective. A wide variety of client tools support the ability to retrieve member properties of a specific member when needed by the data analyst. You can add additional attributes by dragging and dropping a column from the DSV to the Attributes pane or delete an existing attribute by right - clicking that attribute and selecting Delete.

Hierarchies and Levels

Hierarchies (also called multilevel hierarchies) are created from attributes of a dimension. Each multilevel hierarchy contains one or more levels, and each level is an attribute hierarchy itself. Based on the attributes of the Geography dimension you created, the logical multilevel hierarchy to create would be Country - State - City - Postal Code. You can create this hierarchy using the following steps:

1. Switch to the Dimension Structure tab of the Geography dimension and drag and drop the attribute English Country Region Name from the Attributes pane to the Hierarchies pane. This creates a multilevel hierarchy called Hierarchy with one level: English Country Region Name. This level actually corresponds to a country. To make this name more user friendly, rename the English Country Region Name to " Country " by right - clicking the attribute within the multilevel hierarchy and selecting Rename.

2. Drag and drop State - Province Name from the Attributes pane to the Hierarchies pane such that the State - Province Name attribute is below Country in the multilevel hierarchy. Rename State - Province Name to " State - Province " by right - clicking the attribute and selecting Rename.

3. Drag and drop City and Postal Code attributes to the multilevel hierarchy in that order so that you now have a four - level hierarchy Country - State - City - Postal Code.

4. The default name of the hierarchy you have created is " Hierarchy. " Rename it to " Geography " by right - clicking its name and selecting Rename (see Figure 5 - 19). You can also rename hierarchy and level names in the Hierarchies pane by selecting the item and changing the value of its Name property in the Properties pane.

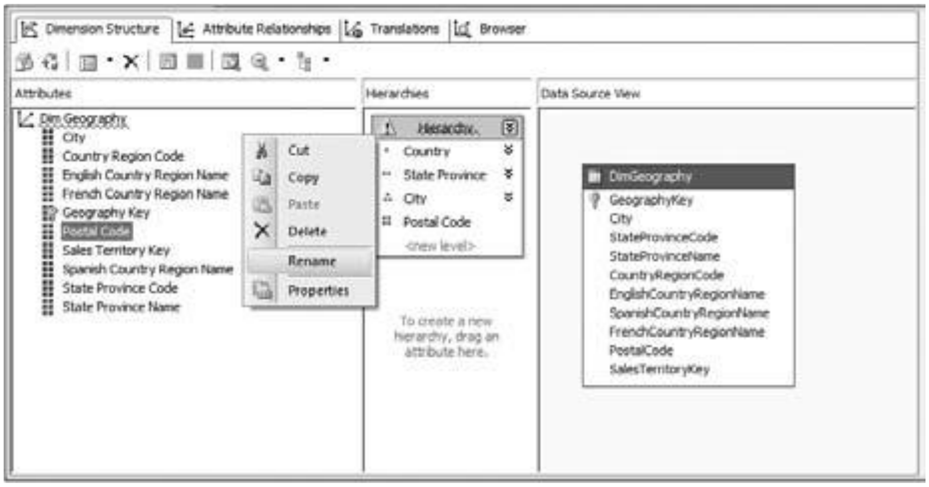

Figure 5-19

You have now created a multilevel hierarchy called Geography that has four levels, as shown in Figure 5 - 20 . You can click the arrows to expand the attribute in each level to see all the member properties. You can create additional hierarchies in the Hierarchies pane.

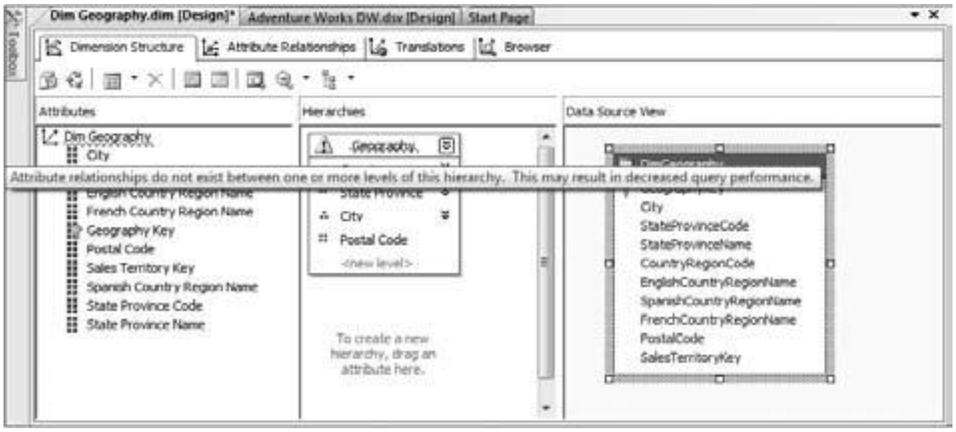

Attribute relationships do not exist between one or more levels of this hierarchy. This may result in decreased query performance.

Figure 5-20

Notice the warning icon next to the Geography hierarchy name and squiggly line under the name of the hierarchy. If you place your mouse over this icon or the hierarchy name, you will see a tooltip message indicating that attribute relationships do not exist between one or more levels of the hierarchy and could result in decreased performance as shown in Figure 5 - 20 .
The current hierarchy design is what is called an unnatural hierarchy. An unnatural hierarchy exists when knowing the attribute value of one level of the hierarchy is not sufficient to know who its parent is in the next level up the hierarchy. Another example of an unnatural hierarchy would be a Customer Gender - Age hierarchy, where Gender is the top level of the dimension and Age is the second level. Knowing that a customer is 37 years old does not give any indication of their gender.

Conversely, in a natural hierarchy, knowing the value of an attribute at one level clearly indicates who its parent is on the next level of the hierarchy. An example of a natural hierarchy would be a Product Department hierarchy with Category and Sub - Category levels. By knowing that a product is in the Mountain Bike Sub - Category, we would know that it belongs to the Bike Category. This relationship between attribute values is defined through attribute relationships. In order for a hierarchy to be considered natural, attribute relationships must exist from the bottom level of the hierarchy all the way to the top. Analysis Services 2008 will only materialize hierarchies that are considered natural. Use the following steps to refine the current Geography hierarchy so that it is natural:

5. Switch back to the Attribute Relationships page. The page should look similar to Figure 5 - 21 .

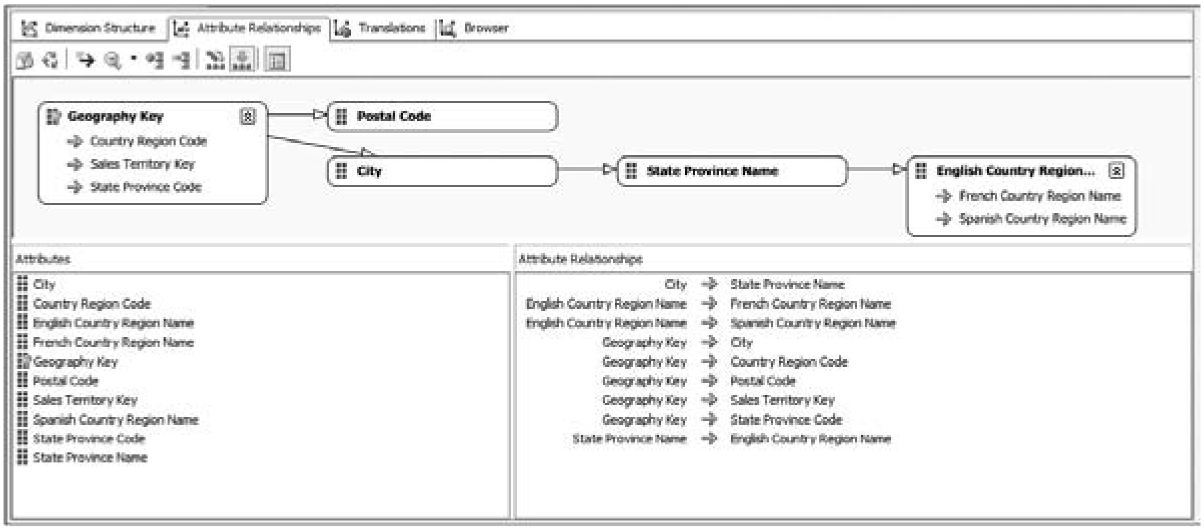

City → State Province Name
English Country Region Name → French Country Region Name
English Country Region Name → Spanish Country Region Name
Geography Key → City
Geography Key → Country Region Code
Geography Key → Postal Code
Geography Key → Sales Territory Key
Geography Key → State Province Code
State Province Name → English Country Region Name

Figure 5-21

6. There is no relationship between Postal Code and City. In the visualization pane drag and drop the Postal Code attribute to the City attribute.

An attribute relationship between the Postal Code attribute and the City attribute is created as shown in Figure 5 - 22 . Notice that the visualization of the attribute relationships extends beyond the visualization pane. (Depending on the resolution of your monitor, you might be able to view the all the attribute relationships.) You can zoom in or zoom out using the Zoom item in the context menu of the visualization pane to adjust the size of the content of the visualization pane to see all the attribute relationships. Sometimes the attribute relationships view can be quite large depending on the number of attributes and the relationships you have established. You can easily navigate to the area of the visualization pane you ' re interested in by clicking the " + " symbol at the far right of the horizontal scrollbar and using the locator window (as shown in Figure 5 - 23).

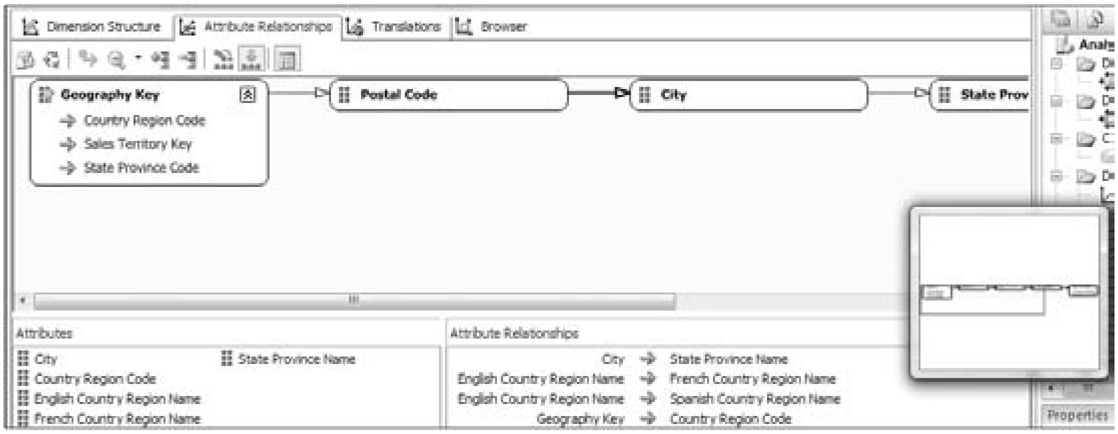

Attributes

City
Country Region Code
English Country Region Name
French Country Region Name
State Province Name

Attribute Relationships

City → State Province Name
English Country Region Name → French Country Region Name
English Country Region Name → Spanish Country Region Name
Geography Key → Country Region Code

Figure 5-22

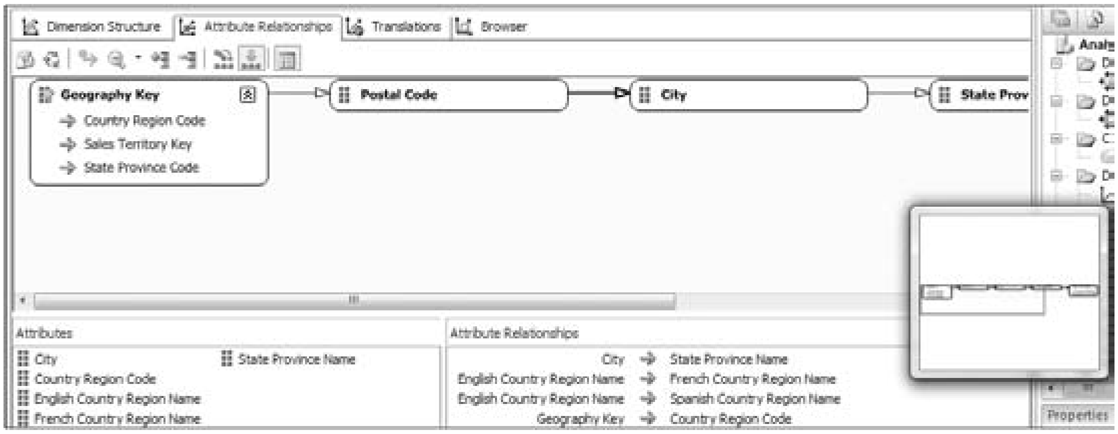

Attributes

City
Country Region Code
English Country Region Name
French Country Region Name
State Province Name

Attribute Relationships

City → State Province Name
English Country Region Name → French Country Region Name
English Country Region Name → Spanish Country Region Name
Geography Key → Country Region Code

Figure 5-23

7. Switch back to the Dimension Structure tab, verify the warning is gone, and save the dimension as shown in Figure 5 - 24 .

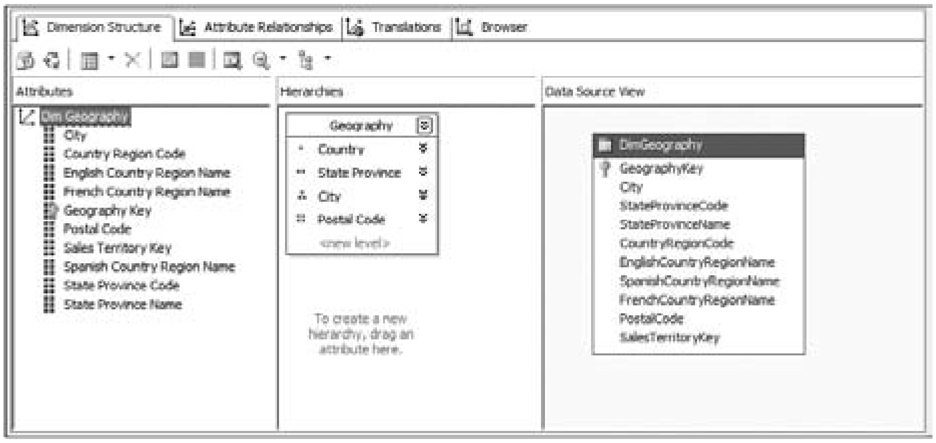

Figure 5-24

Browsing the Dimension

After successfully creating the Dim Geography dimension, you definitely would like to see the results of what you have created and find out how you can see the members of the dimension. So far, the dimension has been designed but not deployed to the server. Indeed, there has been no interaction with the instance of Analysis Services yet. To see the members of the dimension, Analysis Services needs to receive the details of the dimension (the attributes, member properties, and the multilevel hierarchies you have created). The Analysis Services 2008 tools communicate to the instance of Analysis Services via XMLA (XML for Analysis).

XMLA is an industry - standard, Simple Object Access Protocol (SOAP) – based XML Application Programming Interface (API) that is designed for OLAP and Data Mining. The XMLA specification defines two functions, Execute and Discover, which are used to send actions to and retrieve data from the host instance. The Execute and Discover functions take several parameters that define the actions the instance of Analysis Services will perform. One of the parameters of the Execute function is the command sent to an instance of Analysis Services. Note that in addition to supporting the functionality defined in the XMLA specification, Analysis Services supports extensions to the standard. Following is a sample Execute request sent to an instance of Analysis Services using XMLA. The Execute request is a modified version of the one in XMLA specification available at http://www.xmla.org .

```
< Execute xmlns="urn:schemas-microsoft-com:xml-analysis"
SOAP-ENV:encodingStyle="http://schemas.xmlsoap.org/soap/encoding/" >
< Command >
< Statement > select [Measures].members on Columns from Adventure
Works < /Statement >
< /Command >
< Properties >
< PropertyList >
< DataSourceInfo > Provider=SQL Server 2008;Data Source=local; < /DataSourceInfo >
< Catalog > AnalysisServices2008Tutorial < /Catalog >
< Format > Multidimensional < /Format >
< AxisFormat > ClusterFormat < /AxisFormat >
< /PropertyList >
< /Properties >
< /Execute >
< /SOAP-ENV:Body >
< /SOAP-ENV:Envelope >
```

In the preceding XMLA, a request is sent to execute an MDX query that is specified within the command Statement on the catalog AnalysisServices2008Tutorial. The XML request shown results in the query being executed on the server side and the results returned to the client via XMLA.

Several different commands are used to communicate to Analysis Server 2008. Some of the common ones are Create, Alter, Process, and Statement. These commands are used to change the structure of objects referenced in the command. Each object in Analysis Services 2008 has a well - defined set of properties. The definition of the objects is accomplished by commands referred to as Data Definition Language (DDL) commands in this book. Other commands work with data that has already been defined. Those commands are referred to as Data Manipulation Language (DML) commands. You learn some of the DML and DDL commands used in Analysis Services 2008 in various chapters of the book through examples. For in - depth understanding of DML and DDL commands, we recommend you read the Analysis Services 2008 documentation.

You might recall that you deployed the AnalysisServices2008Tutorial project in Chapter 2 . When you deploy a project, the BIDS packages all the design change information in the project into a single XMLA request and sends it to the server. In this case, you want to see the contents of the dimension you have

the entire project using BIDS to Analysis Services, several XMLA requests are sent by BIDS. They are:

1. Request a list of the databases from Analysis Services to determine if the current project already exists on the instance. The project name you specified while creating the object will be used as the database name. Based on the deployment settings in your project, BIDS either sends the entire definition of all the objects or only the changes you have made since the last deploy. BIDS will send either a Create or Alter command based upon whether the database already exists on the Analysis Services instance. We have not included the Create/Alter XMLA request in the following code because it is quite large. You can use the SQL Server Profiler to analyze the XMLA request (you learn to use SQL Server Profiler in Chapter 15).

2. BIDS then sends an XMLA request to process the objects on the instance of Analysis Services. Following is the request that would be sent to the server to process the dimension Dim Geography:

```
< Batch xmlns="http://schemas.microsoft.com/analysisservices/2003/engine" >
< Parallel >
< Process xmlns:xsd="http://www.w3.org/2001/XMLSchema"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xmlns:ddl2="http://schemas.microsoft.com/analysisservices/2003/engine/2"
xmlns:ddl2_2="http://schemas.microsoft.com/analysisservices/2003/engine/2/2"
xmlns:ddl100_100="http://schemas.microsoft.com/analysisservices/2008/engine/100/100" >
< Object >
< DatabaseID > AnalysisServices2008Tutorial < /DatabaseID >
< DimensionID > Dim Geography < /DimensionID >
< /Object >
< Type > ProcessDefault < /Type >
< WriteBackTableCreation > UseExisting < /WriteBackTableCreation >
< /Process >
< /Parallel >
< /Batch >
```

BIDS performs certain validations to make sure your dimension design is correct. If there are errors, BIDS will show those errors using red squiggly lines. In addition to that, a set of error handling properties in the Analysis Services instance helps in validating errors in your dimension design when data is being processed. BIDS sends the default error handling information to the Analysis Services instance for the server to raise any referential integrity errors as part of the deployment process. The default error handling mode in BIDS has been changed in SQL Server 2008 to make sure the developer is aware of all the warnings and errors by default. Follow these steps to deploy your Analysis Services2008Tutorial project:

1. Deploy the AnalysisServices2008Tutorial project to your Analysis Services 2008 instance by either hitting the F5 key or right - clicking the project in the Solution Explorer window and selecting Deploy. BIDS deploys the project to the Analysis Services 2008 instance. You will get a dialog indicating that you have deployment errors as shown in Figure 5 - 25 .

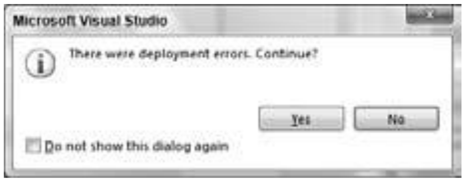

Microsoft Visual Studio

There were deployment errors. Continue?

Figure 5-25

2. Click the No button in the deployment errors dialog shown in Figure 5 - 25 .

BIDS will now report all the warnings as well as errors identified by BIDS and from the Analysis Services instance using the Errors tab as shown in Figure 5 - 26 .

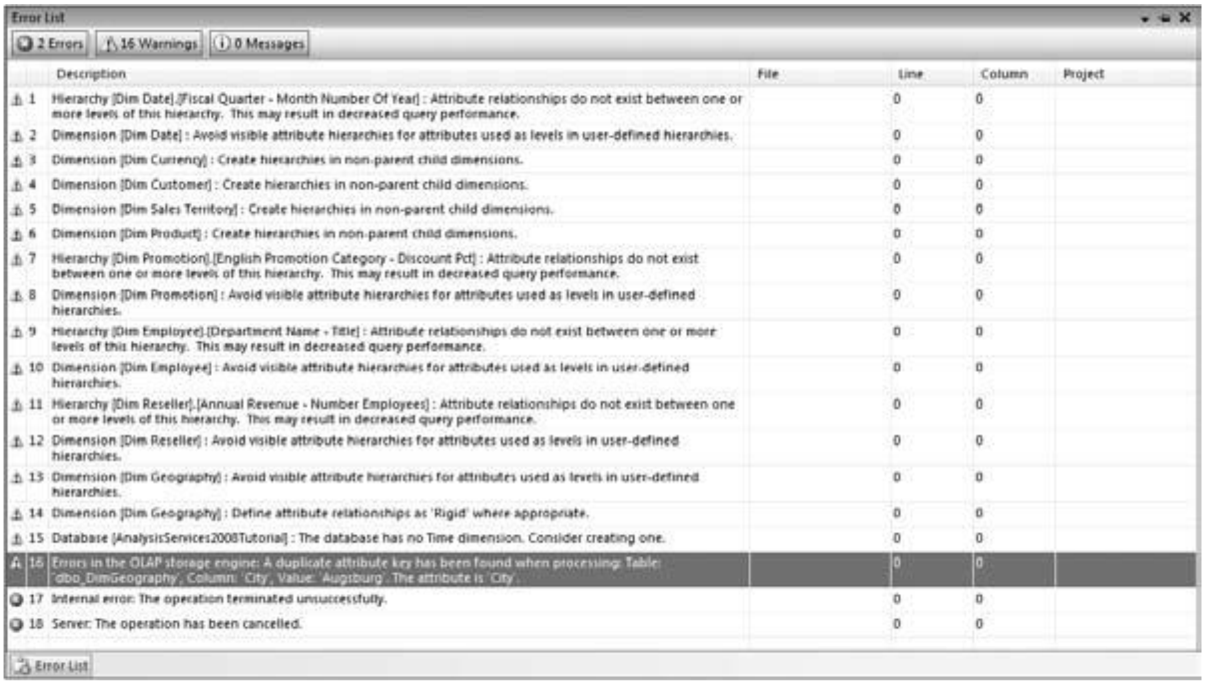

	Description	File	Line	Column	Project
⚠ 1	Hierarchy [Dim Date].[Fiscal Quarter - Month Number Of Year] : Attribute relationships do not exist between one or more levels of this hierarchy. This may result in decreased query performance.		0	0	
⚠ 2	Dimension [Dim Date] : Avoid visible attribute hierarchies for attributes used as levels in user-defined hierarchies.		0	0	
⚠ 3	Dimension [Dim Currency] : Create hierarchies in non-parent child dimensions.		0	0	
⚠ 4	Dimension [Dim Customer] : Create hierarchies in non-parent child dimensions.		0	0	
⚠ 5	Dimension [Dim Sales Territory] : Create hierarchies in non-parent child dimensions.		0	0	
⚠ 6	Dimension [Dim Product] : Create hierarchies in non-parent child dimensions.		0	0	
⚠ 7	Hierarchy [Dim Promotion].[English Promotion Category - Discount Pct] : Attribute relationships do not exist between one or more levels of this hierarchy. This may result in decreased query performance.		0	0	
⚠ 8	Dimension [Dim Promotion] : Avoid visible attribute hierarchies for attributes used as levels in user-defined hierarchies.		0	0	
⚠ 9	Hierarchy [Dim Employee].[Department Name - Title] : Attribute relationships do not exist between one or more levels of this hierarchy. This may result in decreased query performance.		0	0	
⚠ 10	Dimension [Dim Employee] : Avoid visible attribute hierarchies for attributes used as levels in user-defined hierarchies.		0	0	
⚠ 11	Hierarchy [Dim Reseller].[Annual Revenue - Number Employees] : Attribute relationships do not exist between one or more levels of this hierarchy. This may result in decreased query performance.		0	0	
⚠ 12	Dimension [Dim Reseller] : Avoid visible attribute hierarchies for attributes used as levels in user-defined hierarchies.		0	0	
⚠ 13	Dimension [Dim Geography] : Avoid visible attribute hierarchies for attributes used as levels in user-defined hierarchies.		0	0	
⚠ 14	Dimension [Dim Geography] : Define attribute relationships as 'Rigid' where appropriate.		0	0	
⚠ 15	Database [AnalysisServices2008Tutorial] : The database has no Time dimension. Consider creating one.		0	0	
⚠ 16	Errors in the OLAP storage engine: A duplicate attribute key has been found when processing: Table: 'dbo_DimGeography', Column: 'City', Value: 'Augsburg'. The attribute is 'City'.		0	0	
● 17	Internal error: The operation terminated unsuccessfully.		0	0	
● 18	Server: The operation has been cancelled.		0	0	

Figure 5-26

BIDS in Analysis Services 2008 has added a new feature by which you can see warnings identified by BIDS in addition to the errors that were shown in previous versions of Analysis Services. The first 15 warnings shown in Figure 5 - 26 are an example of this new feature in action. Some of the warnings detected by BIDS might be the result of valid design decisions you made for your Analysis Services database. Hence BIDS supports a warning infrastructure by which you can disable warnings for specific objects or even disable specific warnings from reappearing in future deployments. When you right - click one of the first 15 warnings, you can see an option to Dismiss the warning. Note, however, that you cannot dismiss a warning reported by the Analysis Services instance or any error. If you click the sixteenth warning, you can see that the warning cannot be dismissed. This is the first warning reported by the Analysis Services instance followed by two errors that fail the deployment of your AnalysisServices2008Tutorial project. You learn more about the warning feature in Chapter 9 .

The Analysis Services instance warning (warning 16 shown in Figure 5 - 26) indicates that there was an issue while processing the City attribute and duplicate attribute keys were identified. This warning indicates that there are several cities with the same name. This warning is, in fact, an error raised by the Analysis Services instance. The reason why this error is raised is because you have defined and guaranteed a one - to - many attribute relationship between the City attribute and the State - Province Name attribute. The Analysis Services instance identifies that there are several different Cities with the same name and is unable to decide which State - Province name has the relationship to a specific City. If you query the City column in the DimGeography table, you will see that City names are not unique. For example, London, Paris, and Berlin all appear in excess of a dozen times in the DimGeography table of the AdventureWorksDW database. Hence the Analysis Services instance raises the error with the text “ Errors in OLAP Storage Engine. ” Due to this error, the Analysis Services instance fails the processing of the City attribute and subsequently the Dim Geography dimension and the deployment fails.

To correct this issue, you need to make sure each City is unique. You can do this by creating a surrogate key for the City that makes each city unique and use that key as the key column for the City attribute. Alternatively you can use composite keys to uniquely identify a City attribute. In the following steps, you use a collection of columns to define the keys for several attributes of the Dim Geography dimension. To uniquely identify a City, you need to know the State - Province Name to which it belongs. Therefore, the key collection for the City attribute should be City and State - Province Code. Follow these steps to make the City attribute have unique members:

1. Open the Dim Geography dimension in the Dimension Designer and click the City attribute in the Attributes pane.

2. In the Properties pane, locate the KeyColumns property and click the ellipses (as shown in Figure 5 - 27) to open the Key Columns selection dialog.

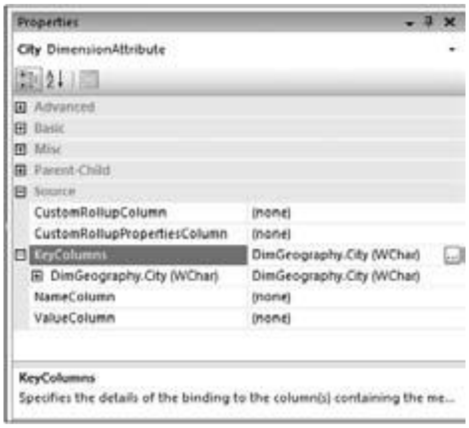

CustomRollupColumn	(none)
CustomRollupPropertiesColumn	(none)
KeyColumns	DimGeography.City (WChar)
DimGeography.City (WChar)	DimGeography.City (WChar)
NameColumn	(none)
ValueColumn	(none)

KeyColumns
Specifies the details of the binding to the column(s) containing the me...

Figure 5-27

3. In the Key Columns selection dialog, add the StateProvince Code to the list of columns and click OK as shown in Figure 5 - 28 .

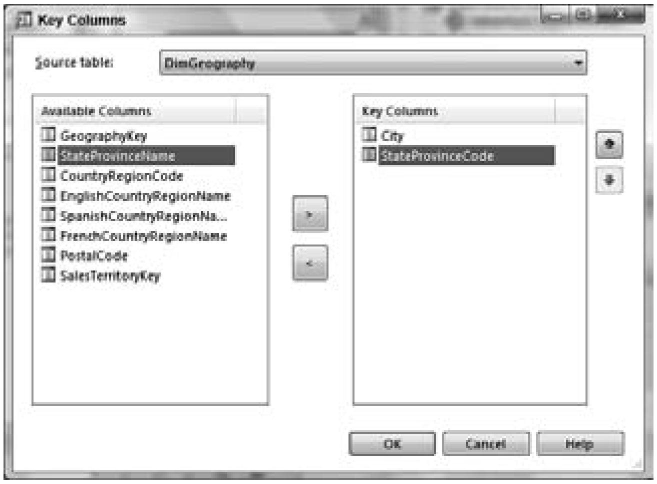

Source table: DimGeography

Available Columns
- GeographyKey
- StateProvinceName
- CountryRegionCode
- EnglishCountryRegionName
- SpanishCountryRegionNa...
- FrenchCountryRegionName
- PostalCode
- SalesTerritoryKey

Key Columns
- City
- StateProvinceCode

Figure 5-28

By default, the Dimension Wizard uses the column name as the key column for the attribute. The Analysis Services instance automatically infers the same column to be the name column (the column that is used to display the member names for the attribute). Whenever you define a composite key, you need to define a name column for the attribute because BIDS and the Analysis Services instance do not know which of the composite key columns should be used as the name column for the attribute.

4. In the NameColumn property for the City attribute, click the ellipses to open the Name Column selection dialog (shown in Figure 5 - 29) and select City as the source for the name of the City attribute.

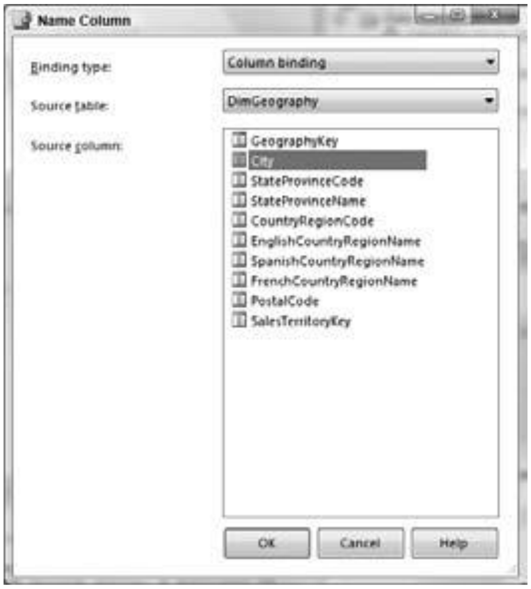

Binding type: Column binding

Source table: DimGeography

Source column:

- GeographyKey
- City
- StateProvinceCode
- StateProvinceName
- CountryRegionCode
- EnglishCountryRegionName
- SpanishCountryRegionName
- FrenchCountryRegionName
- PostalCode
- SalesTerritoryKey

Figure 5-29

The DimGeography table in the data source also contains duplicate PostalCodes. As you just did for the City attribute, you need to make the PostalCode attribute members unique.

5. Select the PostalCode attribute in the Dimension Designer ' s Attributes pane.

6. In the Properties pane, locate the KeyColumns property and click the ellipses to open the Key Columns selection dialog.

7. Change the KeyColumns for the PostalCode attribute to include the StateProvinceCode, City, and PostalCode columns from the data source. Click OK.

8. Change the NameColumn property by clicking the ellipses next to the NameColumn property in the Properties window.

9. In the Name Column dialog, set the NameColumn property to PostalCode. Click OK.

10. Deploy the AnalysisServices2008Tutorial database to the Analysis Services instance.

The AnalysisServices2008Tutorial database now successfully deploys. Now that you have successfully deployed the database, you can browse the data for the Dim Geography structure by switching to the Browser tab of the Dimension Designer as shown in Figure 5 - 30 . To display the data in the browser, BIDS obtains schema information through several Discover requests to retrieve information such as the hierarchies and levels available for the dimension. Finally, an MDX query is sent to the server by BIDS to retrieve dimension data. The MDX query is:

```
SELECT HEAD( [Dim Geography].[Geography].LEVELS(0).MEMBERS, 1000 ) on 0
FROM [$Dim Geography]
```

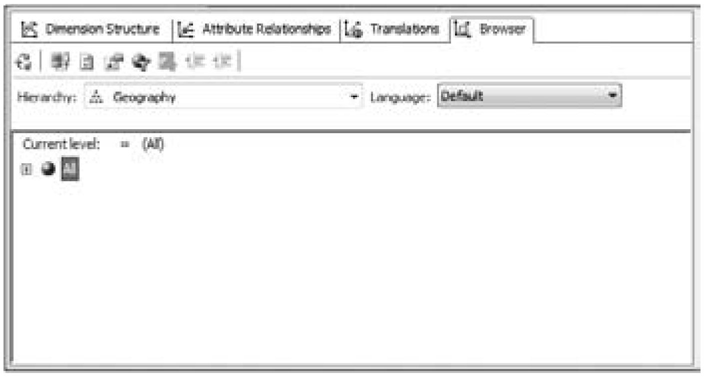

Hierarchy: Geography

Language: Default

Current level: = (All)

All

Figure 5-30

Because you have some familiarity with MDX by now you might have deciphered most of the query. This query uses the HEAD function to request the first 1,000 members from Level 0 of the hierarchy Geography in dimension [Dim Geography]. In the FROM clause you see [$ Dim Geography]. Though you have not created any cube in your data warehouse project yet, you know that the FROM clause should contain a cube name, so how does this MDX query work? When a dimension is created, the server internally stores the values of the dimension as a cube. This means that *every dimension is internally represented as a cube* with a single dimension that holds all the attribute values. The dimension you have created is part of the Analysis Services database AnalysisServices2008Tutorial and is called a database dimension. Because each database dimension is a one - dimensional cube, they can be queried using MDX using the special character $ before the dimension name. This is exactly what you see in the query, [$Dim Geography].

The hierarchy first shown in the hierarchy browser is the most recently created multilevel hierarchy (in this case, Geography). You can choose to browse any of the multilevel hierarchies or attribute hierarchies by selecting one from the drop - down list labeled Hierarchy. This list contains the multilevel hierarchies followed by the attribute hierarchies. Each attribute hierarchy and multilevel hierarchy within a dimension has a level called the All level. In Figure 5 - 30 you can see the All level for the hierarchy Geography. The All level is the topmost level of most hierarchies (the All level can be removed in certain hierarchies) and you can change the name of the All level by changing the property of the hierarchy. It makes sense to call the level " All " because it encompasses all of the sub - levels in the hierarchy. If a hierarchy does not contain the All level, the members of the topmost level would be displayed as the first level in the Dimension Designer ' s Browser page.

Assume you want to change the All level of the Geography hierarchy to " All Countries. " The following steps show how to do this:

1. Go to the Dimension Structure view of the Dimension Designer.

2. Click the Geography hierarchy in the Hierarchies pane.

3. The Properties window now shows all the properties of this hierarchy. The first property is AllMemberName and it displays no value. Add a value by typing **All Countries** in the text entry box to the right of AllMemberName as shown in Figure 5 - 31 .

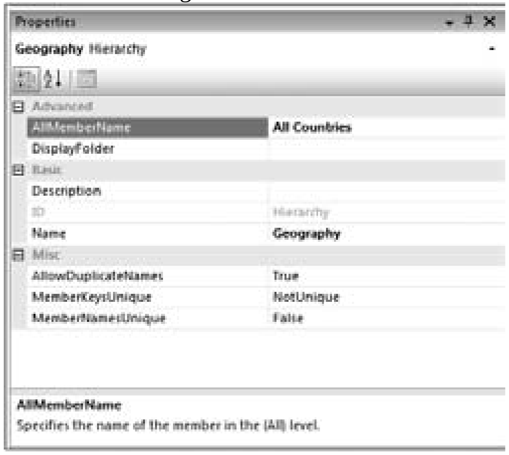

Properties

Geography Hierarchy

Advanced	
AllMemberName	All Countries
DisplayFolder	
Basic	
Description	
ID	Hierarchy
Name	Geography
Misc	
AllowDuplicateNames	True
MemberKeysUnique	NotUnique
MemberNamesUnique	False

AllMemberName

Specifies the name of the member in the (All) level.

Figure 5-31

4. Deploy the project once again.

5. After successful deployment, BIDS switches from the Dimension Structure page to the Browser page. If your Deployment " Processing Option " has been set to " Do Not Process, " you will see a warning in the Dimension Browser pane asking you if you need to process the dimension. If you see this message, click the Process link to launch the Process dialog. Click OK to process the Dim Geography dimension and click the Reconnect button in the Dimension Designer Browser toolbar to view your All member name changes. If your Deployment mode processing has been set to Full or Default, the Dim Geography dimension would be processed along with the deployment and you will see a message in the Dimension Browser to click Reconnect in order to see the latest changes. Click the Reconnect link.

You can now see that the All level of the Geography hierarchy has changed to All Countries, as shown in Figure 5 - 32 . You can also see in the figure that the All Countries level has been expanded to show all members in the next level.

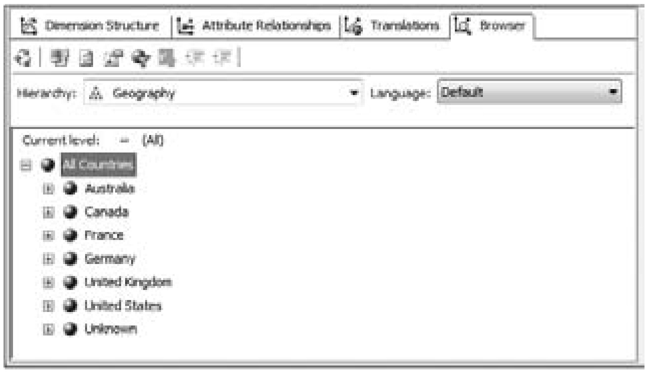

Hierarchy: Geography Language: Default

Current level: (All)

Figure 5-32

When you expand the All Countries level, the following MDX query is sent to the Analysis Services instance to retrieve the members in the next level:

```
WITH MEMBER [Measures].[-DimBrowseLevelKey 0-] AS
‘[Dim Geography].[Geography].currentmember.properties(“key0”, TYPED)’
SELECT { [Measures].[-DimBrowseLevelKey 0-] } ON 0,
HEAD( [Dim Geography].[Geography].[All Countries].Children, 1000) ON 1
FROM [$Dim Geography]
CELL PROPERTIES VALUE
```

The goal of the MDX query is to retrieve all the members that are children of the All Countries level. Similar to the MDX query that was sent to retrieve members in level 0, this query only retrieves the first 1,000 children of All Countries level. This is accomplished by use of the HEAD function as seen in the MDX query. This query includes a calculated measure called Measures.[- DimBrowseLevelKey 0 -], which is selected in the MDX query. The calculated measure expression in this query retrieves the key of the current member by using the MDX function *Properties*. The MDX function Properties returns a string value based on the parameters passed to it. The Properties function returns the value of the member property that is specified as the first argument to the expression. In this query the value requested is the Key of the current member.

Other parameters that can be passed to the Properties function are NAME, ID, and CAPTION, or the name of a member property or related attribute. The properties NAME, ID, KEY, and CAPTION are called *intrinsic member properties* because all attributes and hierarchies will have these properties. The second argument passed to the Properties function is optional and the only value that can be passed is TYPED. If the Properties function is called without the second parameter, the function returns the string representation of the property. If the second argument TYPED is passed to the Properties function, the function returns the data type of the property (data type that was defined in the data source) requested. For example, if the first argument is Key and if the Key of this attribute is of type integer, the Properties function returns integer values. Typically the second parameter TYPED is useful if you want to filter the results based on a member property. For example, if the key of the Geography hierarchy is an integer and if you want to see only the children of member United States, you can use the FILTER function along with the calculated measure that has been created using the parameter TYPED.

The result of the preceding MDX query is shown in the following table. The Dimension Browser retrieves this information and shows the names of the members in the hierarchical format shown in Figure 5 - 32 .

Australia Australia
Canada Canada
France France
Germany Germany
United Kingdom United Kingdom
United States United States
Unknown Unknown

When you defined an attribute relationship between the State - Province Name and City attribute earlier in the chapter, you implicitly set member properties for those attributes. Now you can see these member properties in the Dimension Designer Browser. To do that you can either click the Member Properties button in the Dimension Designer Browser toolbar (highlighted in Figure 5 - 33) or choose Member Properties from the Dimension menu. A dialog appears that has all the attributes of the dimension that participate in attribute relationships. Select the attributes English Country Region Name, State - Province, and City and click OK. The member properties you have selected are now shown in the Dimension Browser as shown in Figure 5 - 33 .

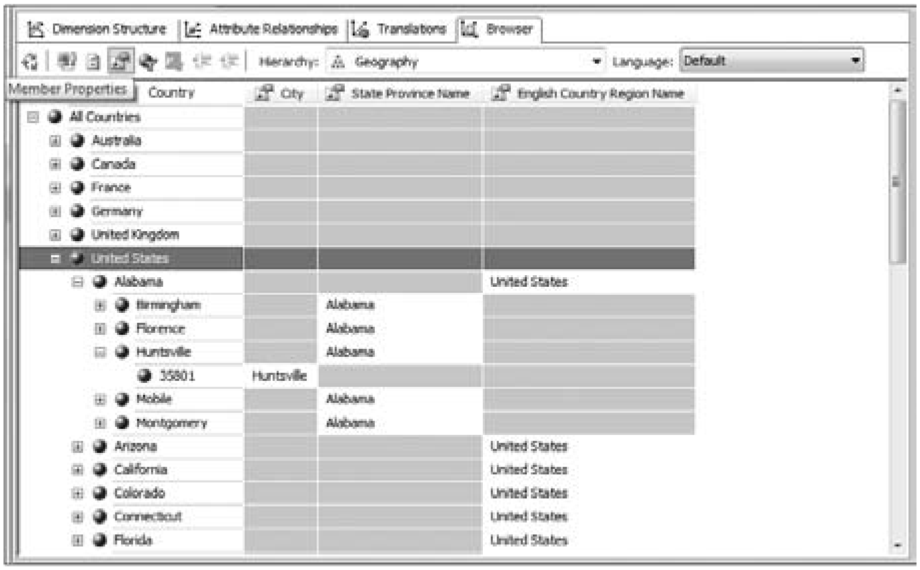

Country	City	State Province Name	English Country Region Name
All Countries			
Australia			
Canada			
France			
Germany			
United Kingdom			
United States			
Alabama			United States
Birmingham		Alabama	
Florence		Alabama	
Huntsville		Alabama	
35801	Huntsville		
Mobile		Alabama	
Montgomery		Alabama	
Arizona			United States
California			United States
Colorado			United States
Connecticut			United States
Florida			United States

Figure 5-33

Expand the members of United States to see the member properties of the States and Cities under United States. The member properties of a member are also retrieved with the help of an MDX query. For example, when you want to see all the cities in Alabama, the following MDX query is sent to the server:

```
WITH MEMBER [Measures].[-DimBrowseLevelKey 0-] AS
‘[Dim Geography].[Geography].currentmember.properties(“key0”, TYPED)’
MEMBER [Measures].[-DimBrowseLevelKey 1-] AS
‘[Dim Geography].[Geography].currentmember.properties(“key1”, TYPED)’
MEMBER [Measures].[-DimBrowseProp State Province Name-] AS
‘[Dim Geography].[Geography].currentmember.properties(“State Province Name”, TYPED)’
SELECT { [Measures].[-DimBrowseLevelKey 0-], [Measures].[-DimBrowseLevelKey 1-],
[Measures].[-DimBrowseProp State Province Name-] } ON 0,
HEAD( [Dim Geography].[Geography].[State-Province]. & [Alabama].Children, 1000) ON 1
FROM [$Dim Geography]
CELL PROPERTIES VALUE
```

Similar to the MDX query you analyzed earlier to retrieve all the members of the All level, this query retrieves all the Alabama City members. The City attribute ’ s member property State - Province Name is retrieved (along with the values that make up the City attribute ’ s composite key, City, and State – Province Code) with the same query as calculated members using the WITH MEMBER clause as seen in the preceding query.