

Лекция 2

Тема «Business Intelligence Development Studio»

You will familiarize yourself with the Analysis Services development environment by working through a tutorial based on a sample relational database for SQL Server Analysis Services 2008 called *Adventure Works DW 2008*, which you can download from www.codeplex.com. This tutorial covers many basic Analysis Services concepts by taking you through the process of building and browsing a cube. The tutorial shows you how to use the tools and also provides you insight into what the product is doing behind the scenes.

In the management environment, you learn the basic operations associated with managing Analysis Services 2008. Further, you learn about the constituent objects that make up an Analysis Services 2008 database and what actions can be taken against them in the management environment. Finally, you are introduced to using the MDX Query Editor to query cube data. *MDX, which stands for Multi Dimensional eXpressions, is the query language used to retrieve data from multi dimensional databases.*

By the end of this chapter you will be familiar with the key components that constitute the Analysis Services Tools, the process of building Analysis Services databases, and how to use MDX to retrieve data from Analysis Services databases. So, snap on your seatbelt and let 's get started!

Differences between Analysis Services 2000, Analysis Services 2005, and Analysis Services 2008

Analysis Services 2005 was not just an evolutionary step up from Analysis Services 2000, but a quantum leap forward in functionality, scalability, and manageability. Analysis Services 2008 builds on the Analysis Services 2005 architecture and enhances its functionality to make it easy and efficient for Analysis Services database developers and administrators to do their jobs. Some of the key enhancements include improvements in cube and dimension wizards to help build your multi - dimensional database to perform more effectively; added guidance in the Analysis Services Tools for improving design; query performance enhancements in the Analysis Services engine; and Analysis Services features such as a read - only database that help in scalability. You learn more about these key enhancements in Analysis Services 2008 throughout the book. Relational databases provide a flexible, well - known model for storing data optimized for rapid incremental updates. They also provide the end user with access to data that can be easily condensed into information - rich reports. OLAP databases, on the other hand, are typically used because of their high - end performance and rich analytic and exploration capabilities. Analysis Services 2008 merges the capabilities of the relational and OLAP worlds, providing a unified view of data to the end user. This unified model is called the *Unified Dimensional Model* (UDM). In sum, Analysis Services 2008 is a powerful, enterprise - class product that you can use to build large - scale OLAP databases and implement strategic business analysis applications against those databases. You learn more about the UDM and the advanced analytic capabilities of Analysis Services 2008 in Chapters 6 , 9 , 21 , 22 , 23 , 24 , and 25 . This chapter gives you hands - on experience with both the development and management environments.

Development, Administrative, and Client Tools

If you have used Analysis Services 2000, you have used the Analysis Manager. The Analysis Manager was implemented as a *Microsoft Management Console* (MMC) snap - in. It served as both the development environment and the management environment for Analysis Services 2000. This tool had limited functionality but did allow you to browse Analysis Services data. A sample application called *MDX Sample* was also shipped in the product and provided the capability of building and sending queries to Analysis Services databases and viewing the results.

Analysis Services 2005 and Analysis Services 2008 have separate environments for development and management. The development environment is called *Business Intelligence Development Studio* (BIDS) and is integrated with Microsoft Visual Studio. Similar to building a Visual Basic or C++ project, you can build a Business Intelligence project. The management environment is called *SQL Server Management Studio* (SSMS) . SSMS is a complete, integrated management environment for several services (including SQL Server itself, Analysis Services, Reporting Services, Integration Services, and SQL Server Compact Edition). SSMS was built to provide ease of use and manageability for database administrators in one single environment. The capability of analyzing and retrieving data from Analysis Services 2008 is integrated into both BIDS and SSMS. You can browse source data from both of these environments as well. In SSMS you are provided with a query builder for writing queries to retrieve data from Analysis

Services. The query builder replaces the MDX Sample application that came with Analysis Services 2000. In addition, the query builder provides IntelliSense support for the MDX language including auto completion and syntax coloring.

If you have used Microsoft SQL Server 2000, you might also be familiar with the SQL Server Profiler. In the SQL Server 2005 release, the capability of tracing, or profiling, Analysis Services queries was added. The SQL Server 2008 SQL Server Profiler also supports Analysis Services profiling. Analysis Services Profile information can be utilized to analyze and improve performance. You learn more about the SQL Server Profiler in Chapter 15 .

Analysis Services Version Differences

Analysis Services 2000 provided a rich feature set that helped in building solid data warehouses. The features combined with the MDX query language provided customers with rich analytic capabilities. As with any software package, though, Analysis Services 2000 had limitations. Some of the limitations of Analysis Services 2000 were:

Even though Analysis Services 2000 had a rich feature set, modeling certain scenarios either resulted in significant performance degradation or simply could not be accomplished. There were size limitations on various database objects such as dimensions, levels, and measures.

Analysis Services 2000 loaded all databases at startup. If there were a large number of databases and/or very large databases, starting the server could take a long time.

Analysis Services 2000 was implemented using a thick client model that helped in achieving very good query performance but did not scale very well in 3 - tier applications (for example, Web scenarios).

The metadata of the databases was either stored in an Access or SQL Server relational database. Maintenance of data and metadata had to be done carefully.

The backup format used by Analysis Services limited the file size to 2GB.

Analysis Services 2008 and Analysis Services 2005, in addition to providing the best of the relational and OLAP worlds, overcame most of the limitations of Analysis Services 2000. The following are some of the improvements implemented:

The thin client architecture improves scalability of 2 - tier and 3 - tier applications.

XML/A (XML for Analysis) was implemented as the native protocol for communication with the server.

Several new OLAP and Data Mining features were added to facilitate easy and optimal design of data warehouses.

Most of the size limits of objects have been greatly increased; or for all practical purposes, eliminated.

Better manageability, scalability, extensibility, fine - grain security, and higher reliability are provided by supporting fail - over clustering.

Native support of Common Language Runtime (CLR) stored procedures with appropriate security permissions is included.

Metadata information is represented as XML and resides in Analysis Services along with the data. This allows for easier maintainability and control.

Analysis Services 2008 uses a different backup format (you learn about backup in Chapters 7 and 13) than the one used in Analysis Services 2000. The 2GB backup file limit from Analysis Services 2000 has been eliminated. The backup format used in Analysis Services 2005 is compatible with Analysis Services 2008.

Analysis Services 2008 significantly enhances the scalability of Analysis Services 2005 backups for databases larger than 20GB. Analysis Services 2008 builds on top of Analysis Services 2005 and provides the following additional benefits:

Analysis Services 2008 enhances your design experience in BIDS by making it easy and efficient to design your databases right from the beginning. BIDS provides informative warnings based on Analysis Services best practices that will help you make optimal choices when designing your Analysis Services databases. You see this in Chapters 5 , 6 , and 9 .

Analysis Services 2008 provides several trace events and performance counters that help you monitor and understand query performance bottlenecks. Several performance enhancements are built into the server that will automatically improve query performance significantly in certain scenarios (which you learn more about in Chapter 15) compared to Analysis Services 2005.

Analysis Services 2008 has much improved database backup performance as compared to Analysis Services 2005. You will notice the improved backup performance in databases that are larger than 20GB. You learn more about backup in Chapter 7 .

Analysis Services 2008 provides you with dynamic management views (DMVs) of all current users and activities that will help you manage your Analysis Services instance efficiently. These

help you in understanding operations within Analysis Services with such things as number of queries and memory consumption. You learn about DMVs in Chapter 13 .

Analysis Services 2008 provides you with shared scalable databases (also called read - only databases) that enable enterprise scale - out scenarios that can handle concurrent requests of several hundreds or thousands of users. You learn about read - only databases and shared scalable databases in Chapters 7 and 15 , respectively.

Two fundamental changes in Analysis Services 2005 that are still applicable in Analysis Services 2008 are the thin client architecture and support for the native XML/A (XML for Analysis) protocol for communication between client and server.

Overall, Analysis Services 2008 provides you with a great combination of functionality and ease of use that enables you to analyze your data and make strategic business decisions. You will see these capabilities emerge step-by-step as you advance through this book.

Upgrading to Analysis Services 2008

You can upgrade to Analysis Services 2008 from Analysis Services 2000 or Analysis Services 2005. If you currently do not have a requirement of upgrading your previous Analysis Services instances to Analysis Services 2008 or if you are a first time user of Analysis Services you can jump to the next section. The Analysis Services upgrade process in general is not a seamless process and not without its share of gotchas. This is especially true when much of the product has been redesigned, such as you are faced with going from Analysis Services 2000 to Analysis Services 2008. Fortunately, Analysis Services 2008 provides a tool called Upgrade Advisor to prepare you to upgrade databases from Analysis Services 2000 and Analysis Services 2005 to Analysis Services 2008. Upgrade Advisor is available as a redistributable package with SQL Server 2008. You need to install Upgrade Advisor from the <processor architecture >\redist\Upgrade Advisor folder on your CD/DVD. Install the Upgrade Advisor on your machine. When you run Upgrade Advisor on your existing Analysis Services 2000 or 2005 instance, Upgrade Advisor informs you whether or not your database(s) will be upgraded successfully without any known issues. Errors and warnings are provided by Upgrade Advisor in cases where upgrade of some of the objects/definitions is not feasible or when there are potential changes in the names of dimensions or cubes during the upgrade process due to the Analysis Services 2008 architecture. Once you have reviewed all the information from Upgrade Advisor, you are ready to start the upgrade. Follow these steps to use Upgrade Advisor for analyzing the effects of upgrading your Analysis Services 2000 or 2005 instance to Analysis Services 2008:

1. Choose Start All Programs SQL Server 2008 SQL Server 2008 Upgrade Advisor on your machine. The welcome screen appears, as shown in Figure 2 - 1 . Click the Launch Upgrade Analyzer Analysis Wizard link at the bottom of the page.



Figure 2-1

2. You will now see the Welcome to Upgrade Advisor for Microsoft SQL Server 2008 page. Click the Next button.
3. In the SQL Server Components selection page, shown in Figure 2 - 2 , enter the name of a machine that contains the Analysis Services 2000 or 2005 instance you want to upgrade. In this illustration, an Analysis Services 2000 server name is specified. If you click the Detect button, Upgrade Advisor will populate the SQL Server Components page with the services running on the server whose name you provided. You can also manually select which services you want Upgrade Wizard to analyze. Select the Analysis Services component as shown in Figure 2 - 2 and click Next.



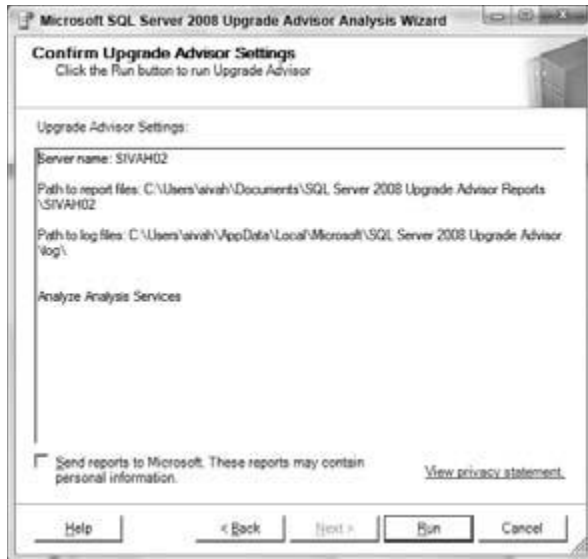
4. In the Analysis Services Parameters page, as shown in Figure 2 - 3 , you can select the Analysis Services instance name. Analysis Services only supports Windows Authentication. Analysis Services 2000 only supports a single instance on one machine, whereas Analysis Services 2005 supports multiple instances. Select the instance name and click Next.

Figure 2-2

Figure 2-3



5. In the Confirm Upgrade Advisor Settings page, as shown in Figure 2 - 4 , you can review your selections. If your selections are not correct, go back to the previous page and make the appropriate changes. Click the Run button for upgrade analysis.



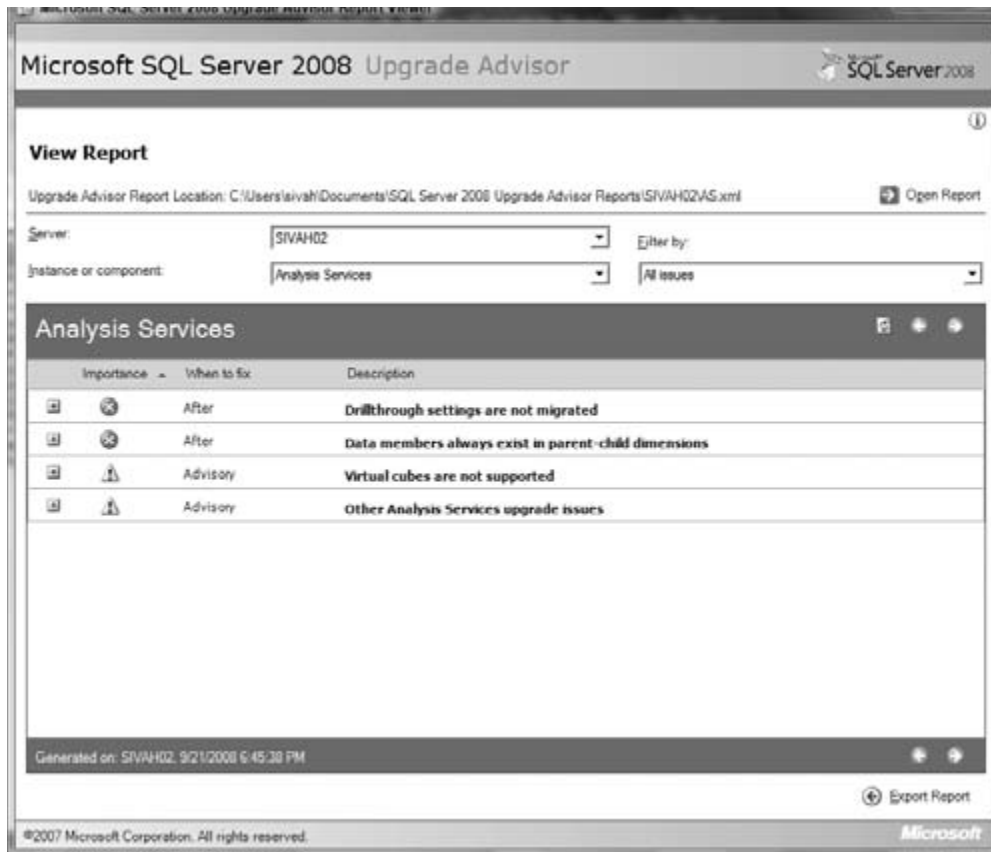
In the next screen you see the Upgrade Advisor analyzing the databases on your Analysis Services instance. You should be aware that the Upgrade Advisor needs the DSO component to connect to your Analysis Services instance. Hence, you need to make sure you install the backward compatibility MSI (SQLServer2005_BC.msi) available with the SQL Server 2008 setup. At the end of the analysis you see the errors and warnings reported by the Upgrade Advisor, as shown in Figure 2 - 5 .



Figure 2-4

Figure 2-5

6. Click the Launch Report button to see the detailed report of the analysis and the actions you need to take for a smooth migration of your databases, as shown in Figure 2 - 6 .



We strongly recommend that you run the Upgrade Advisor utility, analyze all the errors and warnings reported, and take the appropriate actions. In certain cases, you might have to perform some operations on your existing Analysis Services database. For example, if you have a writeback partition in your Analysis Services 2000 database that contains data, the recommended approach is to convert the writeback partition to a MOLAP partition, upgrade the database to Analysis Services 2008, reprocess the partition, and then re - create a new writeback partition. Similarly, you might have to perform several steps either before or after the upgrade on your Analysis Services database to ensure your existing applications will work correctly. Similar to the example shown for analyzing your Analysis Services 2000 database, you need to utilize the Upgrade Advisor to analyze the Analysis Services 2005 database. Because Analysis Services 2008 builds upon Analysis Services 2005 architecture, you may not see a significant number of errors or warnings reported by the Upgrade Advisor for an Analysis Services 2005 database. Even so, you still should test your applications on your Analysis Services 2008 database before proceeding with the upgrade process.

Once you have analyzed the Upgrade Advisor report on your Analysis Services 2000 or Analysis Services 2005 databases you are ready for upgrade. Install the product and select the option to upgrade your Analysis Services 2000 or 2005 databases. Analysis Services 2008 only upgrades the metadata of your Analysis Services 2000 databases, but it upgrades both metadata and data for your Analysis Services 2005 databases. Hence, when you upgrade your Analysis Services 2000 databases you will need your corresponding relational data source available so that source data can be repopulated into your cubes. You need to process the databases that have been upgraded from Analysis Services 2000. Once this is completed, all your cubes and dimensions will be available for querying. If warnings in Upgrade Advisor indicate that names of dimensions or hierarchies will be changed, your applications might also have to be updated accordingly. Please plan to spend time to verify that all your applications are working for your customers after the upgrade process. We have an additional experienced - based recommendation — perform the entire upgrade process on a test machine. In this way, you can verify if your existing applications are working as expected using the Analysis Services 2008 instance. Finally, with confidence, you can perform the upgrade on your production machine. The upgrade process from an Analysis Services 2005 instance to Analysis Services 2008 should be relatively simple. The Upgrade Advisor will report warnings for the issues that affect the upgrade of your Analysis Services 2005 databases that you need to be aware of and handle appropriately.

If you do not have a test machine for upgrading your Analysis Services 2000 instance, you should install Analysis Services 2008 as a named instance and then run the Analysis Services Migration Wizard to migrate your databases from an Analysis Services 2000 server to an Analysis Services 2008 instance. For testing the upgrade process for your Analysis Services 2005 databases, we recommend that you install Analysis Services 2008 as a named instance. You then need to back up your Analysis Services 2005 databases and restore them on your Analysis Services 2008 instance. You then need to test the databases. Once you have confirmed that your applications work against your Analysis Services 2008 instance as expected, you can upgrade your Analysis Services 2005 instance to Analysis Services 2008 using SQL Server 2008 setup 's upgrade path. Analysis Services 2008 provides you with an integrated environment to manage all SQL Server 2008 products using SQL Server Management Studio (SSMS). SSMS is the newer version of the famous Query Analyzer, which is available in SQL Server 2000. Because Analysis Services 2008 builds upon the Analysis Services 2005 architecture, the upgrade process from Analysis Services 2005 to Analysis Services 2008 should be fairly smooth. However, the upgrade process from Analysis Services 2000 to Analysis Services 2008 is bit more involved. Hence, we are including step - by - step instructions. In general we recommend you re - design your Analysis Services 2000 databases in Analysis Services 2008. However, if you do need to upgrade, the tutorial in this section will be helpful to you. If you do not have Analysis Services 2000 databases to upgrade, you can skip the rest of this section.

Using the following tutorial you learn to upgrade from Analysis Services 2000 to Analysis Services 2008. In the following short tutorial, we will reference FoodMart2000 as a sample database and you can use your own databases where appropriate. To migrate your Analysis Services 2000 databases to an Analysis Services 2008 instance, follow these steps:

1. Launch SQL Server Management Studio, which comes with Analysis Services 2008, by choosing Start All Programs Microsoft SQL Server2008 SQL Server Management Studio. Connect to the Analysis Services 2008 instance using SQL Server Management Studio 's Object Explorer. Right - click the server name and select Migrate Database as shown in Figure 2 - 7 . This takes you to the welcome screen of the wizard. If someone else had used this wizard and disabled the welcome page you might not see the welcome page. If you are in the welcome page, click the Next button to proceed to step 2.
2. In the Specify Source and Destination page, the wizard pre - populates the name of your Analysis Services 2008 instance. Enter the machine name of your Analysis Services 2000 server as shown in Figure 2 - 8 and click Next.



Figure 2-7

3. In the Select Databases to Migrate page you will see the list of databases in your Analysis Services 2000 instance itemized and pre - selected for migration as shown in Figure 2 - 9 . A column on the right side provides you with the names of the Destination Databases in your Analysis Services 2008 instance. You have the option of selecting all the databases or just a few databases from your Analysis Services 2000 instance to migrate. Deselect all the databases and select the FoodMart 2000 database; this is the sample database that is shipped with Analysis

Services 2000.

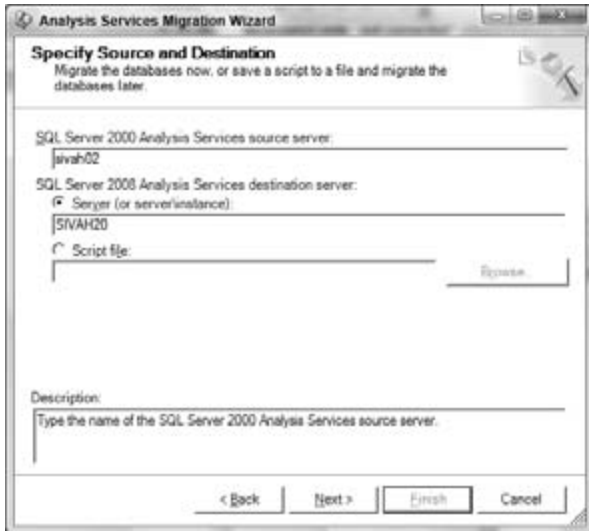
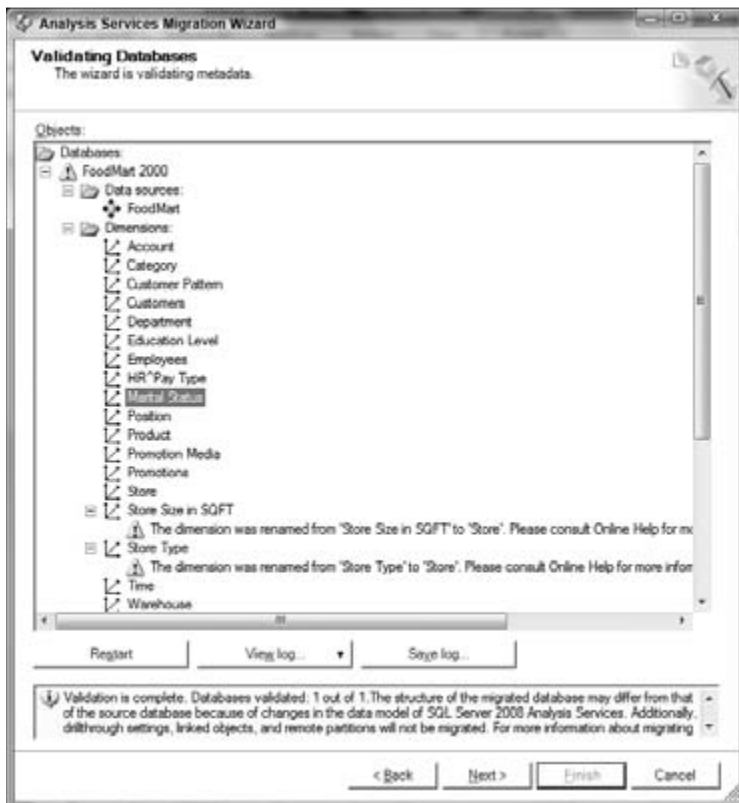


Figure 2-9

4. The Migration Wizard now validates the selected databases and contained objects for migration. As the Migration Wizard does this, it provides a report including warnings for objects that will be changed during the migration process, as shown in Figure 2 - 10 . You can save the logs to a file for future reference. Once you have analyzed the entire report, click Next to deploy the migrated database to your Analysis Services 2008 instance.



5. The Migration Wizard now sends the metadata of the migrated database to the Analysis Services 2008 instance. The new database with migrated objects is created on your Analysis Services 2008 instance and the Migration Wizard reports the status of the migration. Once the migration process is complete, click the Next button.

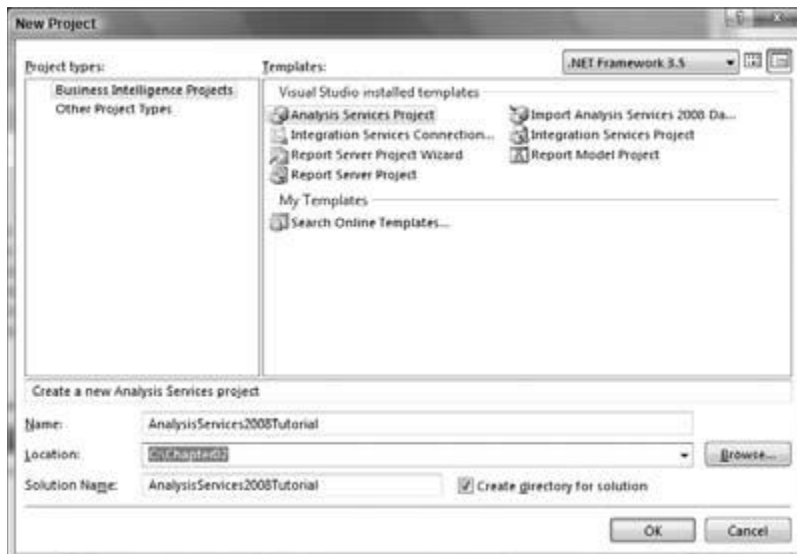
6. In the completion page, the Migration Wizard shows the new databases that have been migrated in a tree view. Click Finish to complete the migration. You should be aware that the Migration Wizard will only migrate databases from Analysis Services 2000. In addition, the wizard only migrates the metadata of an Analysis Services 2000 database and not the data. Hence the migrated cubes and dimensions are not accessible for querying until you reprocess them. Process all the databases that have been migrated, and test your applications against the migrated databases on your Analysis Services 2008 instance. You need to direct your applications to the new Analysis Services 2008 instance name. Once you have verified that all applications are working as expected, you can uninstall Analysis Services 2000 and then rename your Analysis Services 2008 named instance to the default instance using the instance rename utility, ASInstanceRename.exe, which you can find in the \Program Files\Microsoft SQL Server\100\Tools\Binn\VSShell\Common7\IDE directory.

Using Business Intelligence Development Studio

Business Intelligence Development Studio (BIDS) is the development environment for designing your Analysis Services databases. To start Business Intelligence Development Studio, click the Windows Start button and go to All Programs Microsoft SQL Server 2008 SQL Server Business Intelligence Development Studio. If you're familiar with Visual Studio, you might be thinking that BIDS looks a lot like the Visual Studio environment. You're right; in Analysis Services 2008, you create Analysis Services projects in an environment that is Visual Studio. Working in Visual Studio offers many benefits, such as easy access to source control and support for multiple projects within the same Visual Studio solution (a solution within Visual Studio is a collection of projects such as an Analysis Services project, a C# project, an Integration Services project, or a Reporting Services project).

Creating a Project in the Business Intelligence Development Studio

To design your Analysis Services database you need to create a project using BIDS. Typically you will design your database within BIDS, make appropriate changes, and finally send the database to your Analysis Services instance. Each Analysis Services project within BIDS becomes a database on the Analysis Services instance when all the definitions within the project are sent to the server. You can also use BIDS to directly connect to an Analysis Services database and make changes to the database. Follow these steps to create a new project. To start BIDS, click the Start button and go to All Programs Microsoft SQL Server 2008 SQL Server Business Intelligence Development Studio. In BIDS, select File New Project. You will see the Business Intelligence Projects templates as shown in Figure 2 - 11 . Click the Analysis Services Project template. Type **AnalysisServices2008Tutorial** as the project name and select the directory in which you want to create this project. Click OK to create the project.



You are now in an Analysis Services project, as shown in Figure 2 - 12 . When you create a Business Intelligence project, it is created inside a solution with the same name. (A Visual

Studio solution is a container for one or more projects.) When you create a new project with a solution open in Visual Studio, you have the option of adding the project to the existing solution or creating a new one. BIDS contains several panes; of most concern here are the Solution Explorer, Properties, and Output panes.

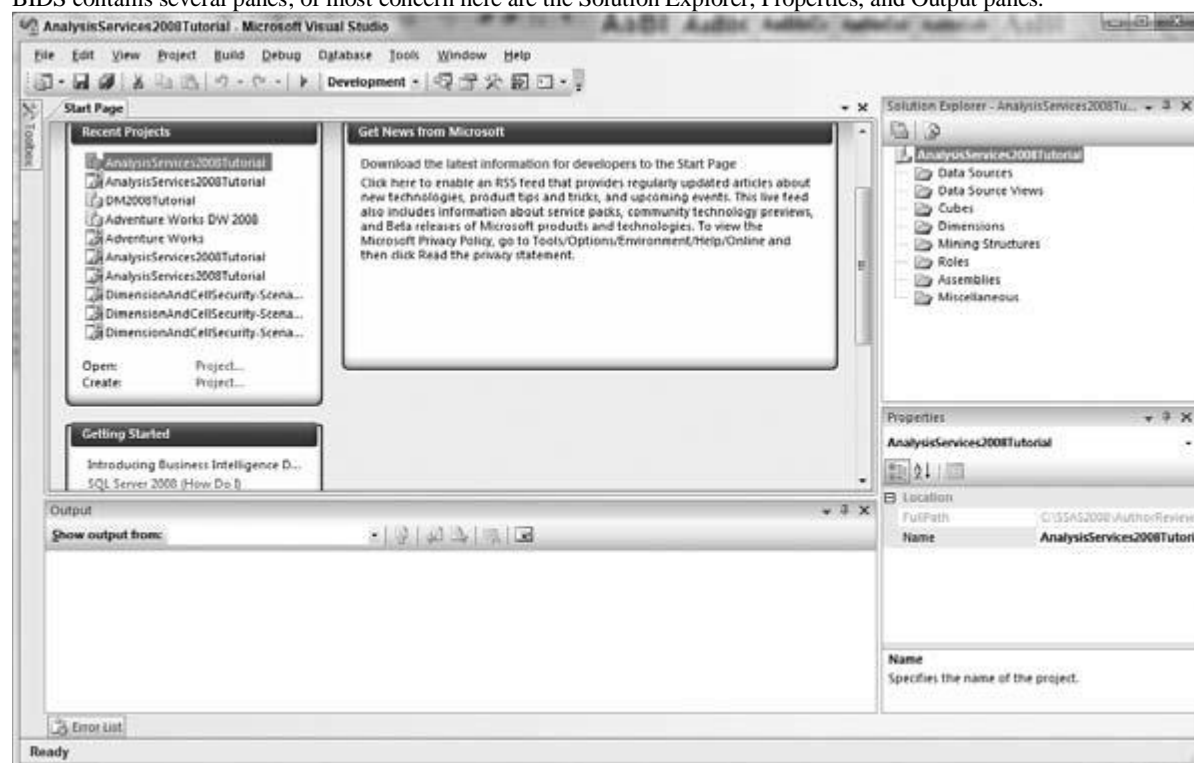


Figure 2-12

The Solution Explorer Pane

The Solution Explorer in Figure 2 - 12 shows eight folders:

Data Sources: Your data warehouse is likely made up of disparate data sources such as Microsoft SQL Server, Oracle, DB2, Teradata, and so forth. Analysis Services 2008 can easily deal with retrieving relational data from various relational databases. Data source objects contain details of a connection to a data source, which include server name, catalog or database name, and login credentials. You establish connections to relational servers by creating a data source for each one.

Data Source Views: When working with a large operational data store you don't always want to see all the tables in the database. With Data Source Views (DSVs), you can limit the number of visible tables by including only the tables that are relevant to your analysis. DSVs allow you to create a logical data model upon which you build your Unified Dimensional Model. A DSV can contain tables from one or more data sources, and one of these data sources is called a primary data source. Data sources and DSVs are discussed in Chapter 4 .

Cubes: Cubes are the foundation for analysis. A collection of *measure groups* (discussed later in this chapter) and a collection of dimensions form a cube. Each measure group is composed of a set of *measures* . Cubes can have more than three dimensions; they are mathematical constructs and not necessarily the three - dimensional objects their name suggests. You learn more about cubes later in this chapter and throughout the book.

Dimensions: Dimensions are the categories by which you slice your data to view specific quantities of interest. Each dimension contains one or more *hierarchies* . Two types of hierarchies exist: attribute hierarchies and user hierarchies. In this book, attribute hierarchies are referred to as attributes, and user or multi level hierarchies are referred to as hierarchies. Attributes correspond to columns in a dimension table, and hierarchies are formed by grouping several related attributes. For example, most cubes have a Time dimension. A Time dimension typically contains the attributes Year, Month, Date, and Day and a hierarchy for Year - Month - Date. Sales cubes often contain Geography dimensions, Customer dimensions, and Product dimensions. You learn about dimensions in Chapter 5 .

Mining Structures: Data mining (covered in Chapter 16) is the process of analyzing raw data using algorithms that help discover interesting patterns not typically found by ad - hoc analysis.

Mining Structures are objects that hold information about a data set. A collection of mining models form a mining structure. Each mining model is built using a specific data mining algorithm and can be used for analyzing patterns in existing data or predicting new data values. Knowing these patterns can help companies make their business processes more powerful. For example, the book recommendation feature on Amazon.com relies on data mining.

Roles: Roles are objects in a database that are used to control access permissions to the database objects (read, write, read/write, process). If you want to provide only read access to a set of users you could create a single role that has read access and add all the users in that set to this role. There can be multiple roles within a database. If a user is a member of several roles, the user inherits the permissions of those roles. If there is a conflict in permissions, Analysis Services grants the most liberal access to the user. You learn more about roles in Chapters 7 and 22 .

Assemblies: Assemblies are user - defined functions that can be created using a .NET language such as Visual Basic.NET, Visual C# .NET, or through languages such as Microsoft Visual Basic or Microsoft C++ that can produce Component Object Model (COM) binaries. These are typically used for custom operations that are needed for specific business logic and are executed on the server for efficiency and performance. Assemblies can be added at the server instance level or within a specific database. The scope of an assembly is limited to the object to which the assembly has been added. For example, if an assembly is added to the server, that assembly can be accessed within every database on the server. On the other hand, if an assembly has been added within a specific database, it can only be accessed within the context of that database. In BIDS you can only add .NET assembly references. You learn more about assemblies in Chapter 11 .

Miscellaneous: This object is used for adding any miscellaneous objects (design or meeting notes, queries, temporary deleted objects, and so on) that are relevant to the database project. These objects are stored in the project and are not sent to the Analysis Services instance.

The Properties Pane

If you click an object in the Solution Explorer, the properties for that object appear in the Properties pane. Items that cannot be edited are grayed out. If you click a particular property, the description of that property appears in the Description pane at the bottom of the Properties pane.

The Output Pane

The Output pane (seen later in this chapter) is used to report warnings and errors during builds. When a project is deployed to the server, progress reporting and error messages are displayed in this pane.

