



Global Knowledge®

Expert Reference Series of White Papers

Microsoft SQL Server 2008 R2: What's New with PowerPivot?

Microsoft SQL Server 2008 R2: What's New with PowerPivot?

Brian D. Egler, Global Knowledge Course Director, MCITP, MCT 2010

Introduction

One of the most eagerly awaited features of Microsoft SQL Server 2008 R2 is the "Self Service Business Intelligence" technology known as PowerPivot. Working in conjunction with Office Excel 2010, data from multiple sources can be combined and analyzed together with the underlying complexities hidden from the analyst. This white paper highlights the new PowerPivot features.

Your Data – Any Place, Any Time

The Microsoft Data Platform Vision is summed up by the phrase: "Your Data Any Place, Any Time."

Four underlying themes categorize the Microsoft SQL Server 2008 release.

- Enterprise Data Platform
- Beyond Relational
- Dynamic Development
- Pervasive Insight

This white paper explores a few of the new features of SQL Server 2008 R2 that were first announced at the SQL PASS Summit Conference in Seattle back in November 2008 and subsequently released in April 2010. The major new features were categorized under three main projects: Project Gemini, Project Madison, and Project Kilimanjaro. Project Gemini produced the "Self Service Business Intelligence" feature set, which is supported by PowerPivot and is the subject of this white paper.

Other features, too numerous to describe here, are documented on the Microsoft web site at <http://www.microsoft.com/sql>. You can also download a trial version of the software from the same location.

Project Gemini: Self-Service Business Intelligence

Microsoft's Business Intelligence tool-of-choice on the client-side is Office Excel. This makes sense since the typical Business Analyst uses Excel on a day-to-day basis so why have them learn a new tool? Using the OLE DB Provider for SQL Server/Analysis Services, the end-user can connect to a Data Source directly and build a Pivot Table or Chart using standard Excel functionality. The importance of Excel 2010 in conjunction with SQL Server 2008 R2 lies in the PowerPivot feature that incorporates extensive column compression technology named "VertiPaq." It was not too long ago that Excel had a meager 100,000 row limit. Now Excel 2010 will support

hundreds of millions of rows with sub-second sorting capability on the client. In addition to this functionality, there is tight integration with Sharepoint 2010 using the server-side version of PowerPivot for publishing dynamic results using Excel Services.

PowerPivot Versions

PowerPivot functionality exists in two versions.

PowerPivot for Excel: the “Add-In” for Excel 2010 is the client-side version for developing powerful results using Pivot Tables and Charts within Microsoft Office 2010 to produce PowerPivot workbooks. Once ready for publishing, the dynamic results can be shared using PowerPivot for Sharepoint.

PowerPivot for Sharepoint: the Add-In for Microsoft Sharepoint 2010 is the server-side version of PowerPivot and works in conjunction with SQL Server 2008 R2. PowerPivot workbooks can be published using Excel Services for dynamic access via Sharepoint.

Prerequisites

PowerPivot for Excel is a free download from <http://PowerPivot.com>. It requires Microsoft Office Professional Plus 2010 and can run on Windows XP (SP3), Vista (SP1), Windows 7, or Windows Server 2008 SP2 with the Platform Update that includes the .NET Framework 3.5 SP1.

PowerPivot Walkthrough

After installing the PowerPivot for Excel Add-In, we see a new tab or “ribbon” in Excel 2010 for PowerPivot.




Figure 1: PowerPivot ribbon in Excel 2010

Within this ribbon, we can launch the PowerPivot Window, which is the main user interface for developing PowerPivot models.

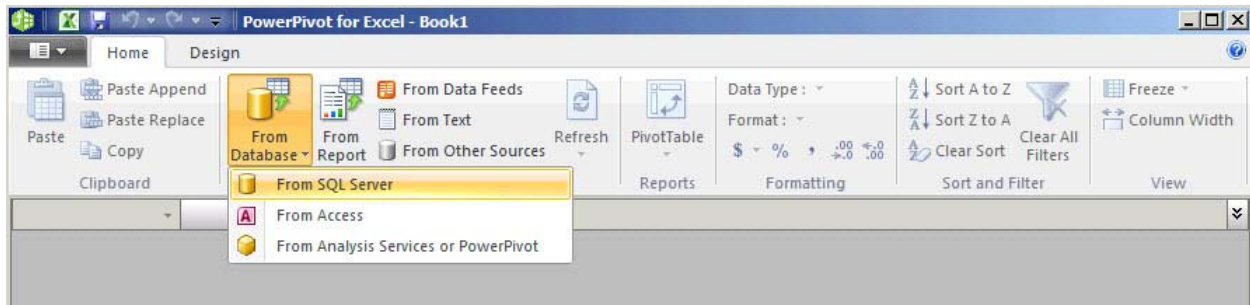


Figure 2: PowerPivot Window

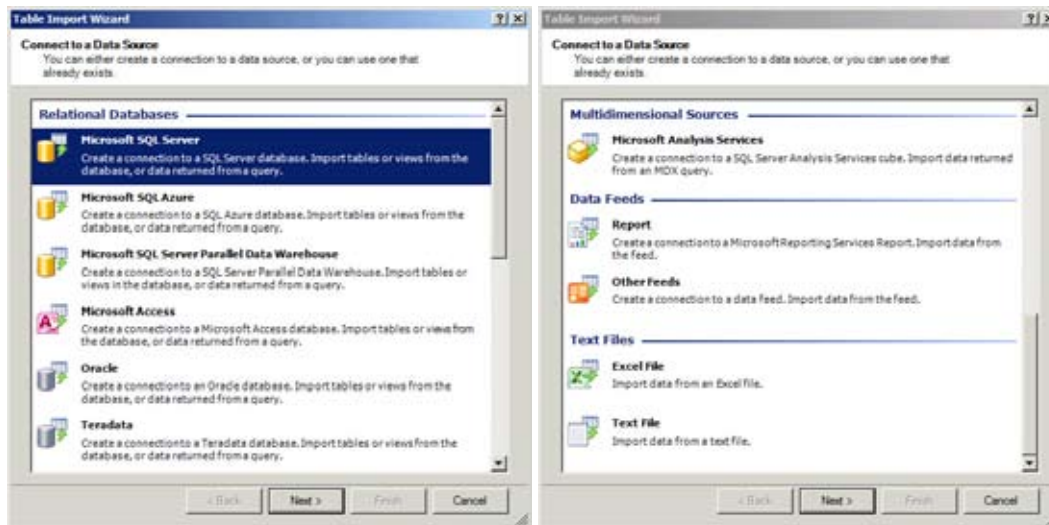


Figure 3: PowerPivot Data Sources

A PowerPivot model consists of one or more data sources and data relationships that form the basis of the PowerPivot application. Once constructed, the PowerPivot model can be used within Excel using Pivot Tables or Charts to analyze the data. The data sources can range from relational databases to multi-dimensional cubes or, if necessary, from text files or data embedded in reports. When using SQL Server 2008 R2, the data is dramatically compressed using a column compression technology called "VertiPaq" providing fast access and sorting.

InventoryKey	DateKey	StoreKey	ProductKey	CurrencyKey	OnHandQuantity	OnOrderQuantity
1	02/09/2008	308	2086	1	16	3
2	06/09/2007	239	643	1	21	1
3	01/31/2009	205	18	1	17	1
4	06/02/2007	199	1587	1	21	1
5	11/21/2009	29	2269	1	45	13
6	10/20/2007	155	1480	1	7	2
7	08/09/2008	61	2379	1	18	0
8	08/09/2008	268	736	1	9	0
9	02/07/2009	232	1457	1	9	0
10	05/19/2007	167	904	1	7	0
11	06/09/2007	171	2172	1	10	0
12	07/28/2007	236	1217	1	22	0
13	07/21/2007	230	638	1	10	0
14	11/10/2007	103	1509	1	10	0
15	01/24/2009	154	1658	1	19	0
16	11/29/2008	199	2504	1	220	170
17	10/27/2007	214	1522	1	10	2
18	02/10/2007	134	2109	1	6	0
19	02/03/2007	193	1962	1	6	0

Figure 4: PowerPivot Window with linked tables

InventoryKey	DateKey	StoreKey	ProductKey	CurrencyKey	OnHandQuantity	OnOrderQuantity
1	02/09/2008	308	2086	1	16	3
2	06/09/2007	239	643	1	21	1
3	01/31/2009	205	18	1	17	1
4	06/02/2007	199	1587	1	21	1
5	11/21/2009	29	2269	1	45	13
6	10/20/2007	155	1480	1	7	2
7	08/09/2008	61	2379	1	18	0
8	08/09/2008	268	736	1	9	0
9	02/07/2009	232	1457	1	9	0
10	05/19/2007	167	904	1	7	0
11	06/09/2007	171	2172	1	10	0
12	07/28/2007	236	1217	1	22	0
13	07/21/2007	230	638	1	10	0
14	11/10/2007	103	1509	1	10	0
15	01/24/2009	154	1658	1	19	0
16	11/29/2008	199	2504	1	220	170
17	10/27/2007	214	1522	1	10	2
18	02/10/2007	134	2109	1	6	0
19	02/03/2007	193	1962	1	6	0

Figure 5: PowerPivot Dynamic Sort

The PowerPivot Window displays data sources as tables within Excel worksheets. Each table is represented by a worksheet displayed using tabs at the bottom of the window using familiar Excel controls. Dynamic sorting can

be performed rapidly within the PowerPivot Window. In the example shown, there are over 8 million rows which sort with sub-second response time.

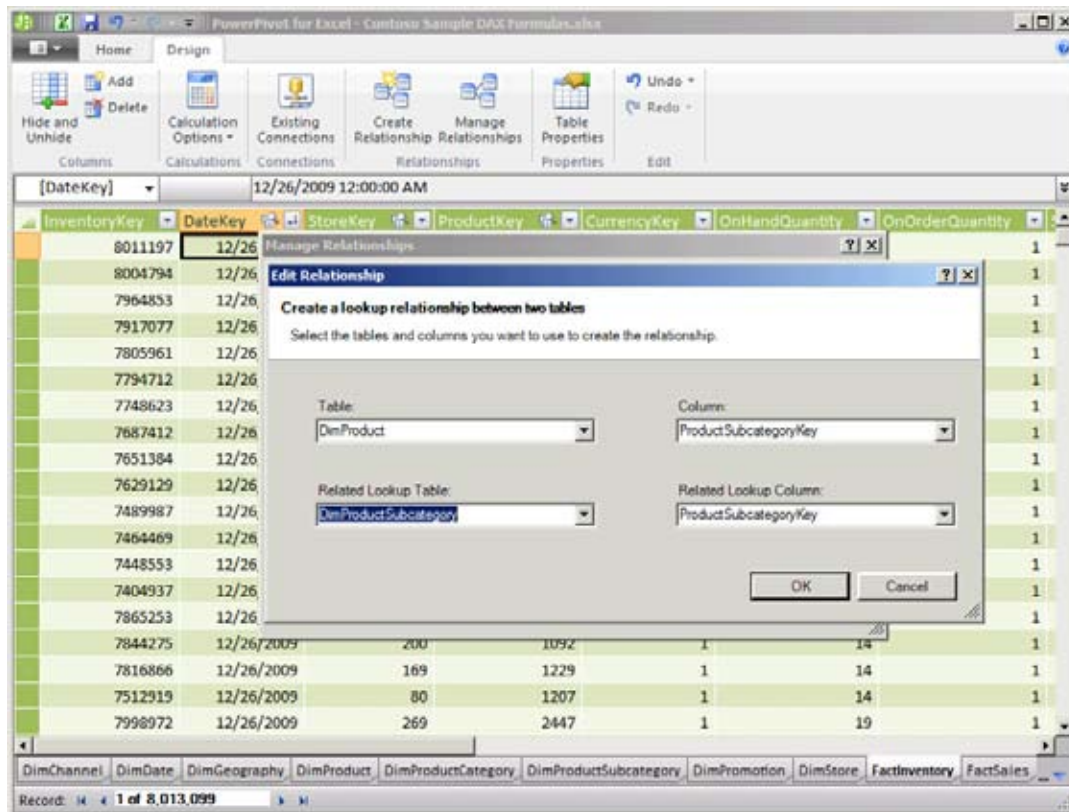


Figure 6: Manage Relationships

Using the Design ribbon within the PowerPivot Window, we can define relationships among the data tables, if necessary. Within a Data Source, data relationships are normally predefined, but this user interface allows us to create relationships between multiple data sources or if a particular required relationship is missing.

TotalCost	SalesAmount	ETLLoadID	LoadDate	UpdateDate	Margin	MarginPct	ProductColor	StoreCity	CountRows
\$728.40	\$1,544.40	1	1/1/2010	1/1/2010	\$816.00	52.8%	Black	Baldon	
\$40.60	\$78.61	1	1/1/2010	1/1/2010	\$38.01	48.3%	Black	Seattle	
\$1,881.27	\$3,628.50	1	1/1/2010	1/1/2010	\$1,747.23	48.2%	White	Cambridge	
\$1,063.20	\$2,254.20	1	1/1/2010	1/1/2010	\$1,191.00	52.8%	Grey	Berlin	
\$3,468.48	\$10,207.08	1	1/1/2010	1/1/2010	\$6,738.60	66.0%	Blue	Berlin	
\$6,621.84	\$13,999.65	1	1/1/2010	1/1/2010	\$7,377.81	52.7%	Silver	North Harford	
\$408.36	\$843.60	1	1/1/2010	1/1/2010	\$435.24	51.6%	Green	Beijing	
\$2,069.37	\$4,399.91	1	1/1/2010	1/1/2010	\$2,330.54	53.0%	Black	Beijing	
\$799.68	\$1,702.58	1	1/1/2010	1/1/2010	\$902.90	53.0%	Black	Bethesda	
\$908.46	\$1,702.80	1	1/1/2010	1/1/2010	\$794.34	46.6%	Black	Beijing	
\$65.24	\$128.00	1	1/1/2010	1/1/2010	\$62.76	49.0%	Blue	Bethesda	
\$1,167.50	\$2,290.00	1	1/1/2010	1/1/2010	\$1,122.50	49.0%	Black	Buffalo	
\$938.28	\$2,832.00	1	1/1/2010	1/1/2010	\$1,893.72	66.9%	Green	South Portland	
\$331.11	\$639.49	1	1/1/2010	1/1/2010	\$308.38	48.2%	Blue	Everett	
\$219.20	\$430.00	1	1/1/2010	1/1/2010	\$210.80	49.0%	Silver	Bethesda	
\$642.33	\$1,245.91	1	1/1/2010	1/1/2010	\$603.58	48.4%	Black	Anchorage	
\$864.72	\$2,609.91	1	1/1/2010	1/1/2010	\$1,745.19	66.9%	Silver	Holyoke	
\$85.65	\$156.80	1	1/1/2010	1/1/2010	\$71.15	45.4%	Silver	Bethesda	
\$2,461.08	\$7,428.00	1	1/1/2010	1/1/2010	\$4,966.92	66.9%	Black	Thimphu	
\$1,529.00	\$2,999.00	1	1/1/2010	1/1/2010	\$1,470.00	49.0%	Black	Moscow	
\$2,298.75	\$5,158.71	1	1/1/2010	1/1/2010	\$2,859.96	55.4%	White	Manitowoc	
\$525.24	\$1,129.41	1	1/1/2010	1/1/2010	\$604.17	53.5%	Silver	Framingham	

Figure 7: Simple PowerPivot DAX Example

Calculated columns can be created by using DAX (Data Analysis Expressions), a special new syntax for PowerPivot that extends normal Excel functions to provide powerful dynamic aggregation capabilities. The example above shows a simple Margin calculation as the difference between the Cost and the SalesAmount. The similarity with Excel expressions will allow current Excel users to use their existing skills while providing simplicity of the user interface.

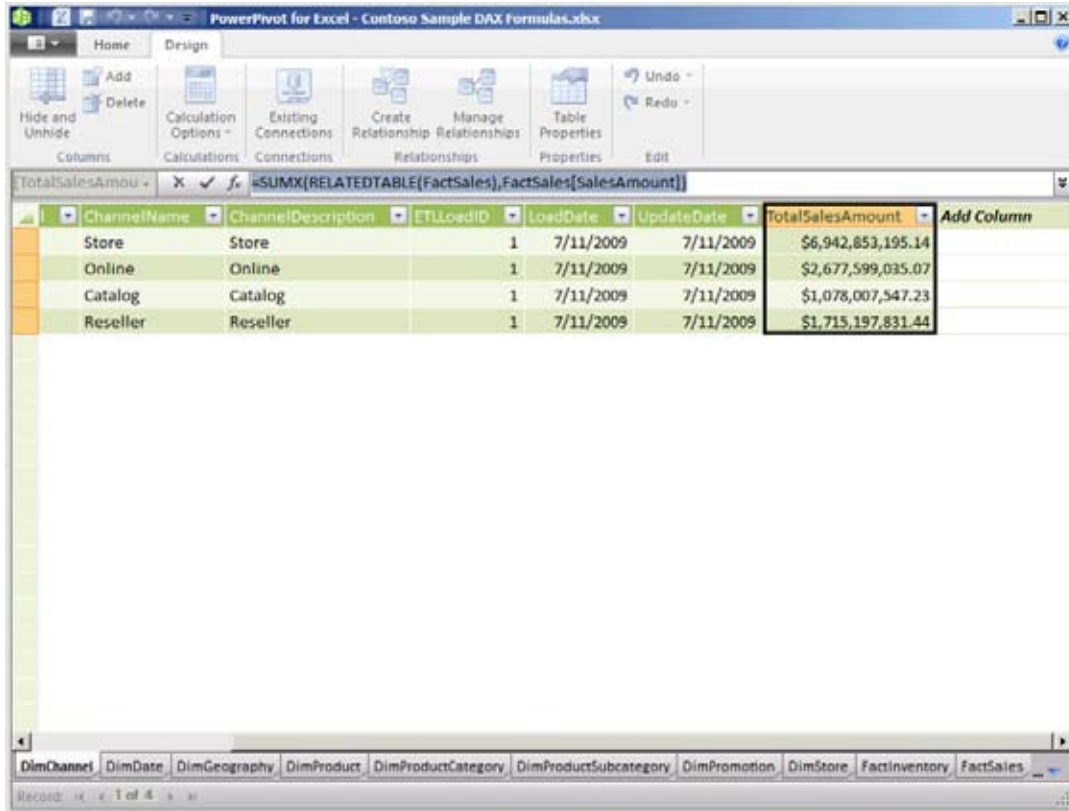


Figure 8: Advanced PowerPivot DAX – SUMX RELATEDTABLE Example 1

This advanced example shows aggregation of a “measure” TotalSalesAmount across a “dimension” DimChannel using the relationship between the DimChannel and the FactSales tables. SUMX is a new DAX function explicitly for this purpose. RELATEDTABLE generates a table of aggregation based on the stated measure. Measures and Dimensions are standard objects defined in a relational Data Warehouse or a Multi-dimensional “Cube.” PowerPivot can be used with any Data Source but in this case, the source data comes from a relational Data Warehouse supported by SQL Server 2008 R2. The power here is that the same DAX expression can be made on any table with a relationship with the FactSales table giving aggregation by all the rows in that dimension table. Below the same expression is made against the ProductCategory table giving context sensitive results.

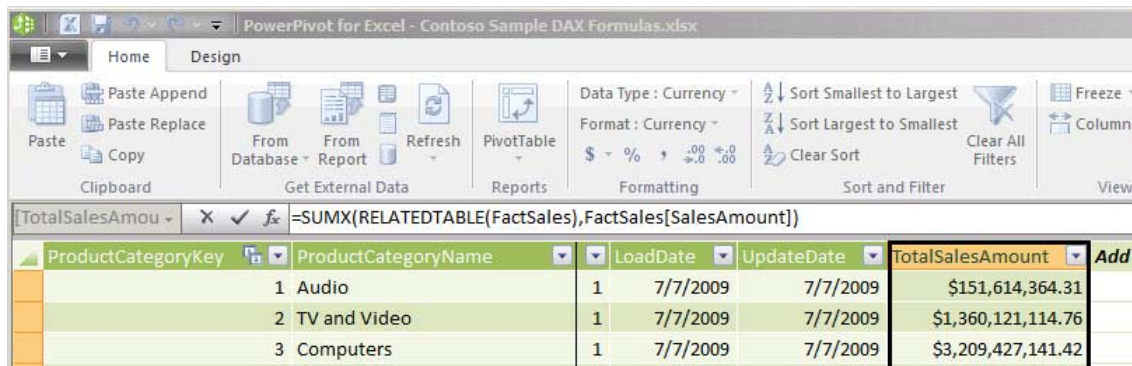


Figure 9: Advanced PowerPivot DAX – SUMX RELATEDTABLE Example 2

Once the PowerPivot model is defined, PivotTables, and PivotCharts can be made using Excel worksheets. Note that the PowerPivot model could be predefined and saved as an .xlsx file, if necessary, allowing the design work to be performed ahead of time.

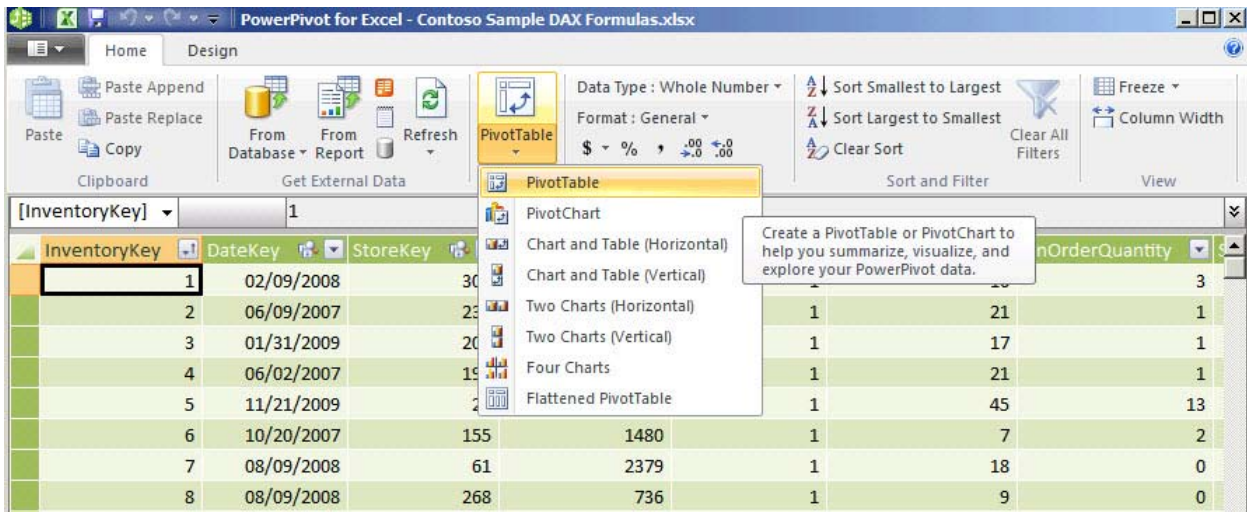


Figure 10: Create PivotTable

A new worksheet can be created to host the PivotTable.




Figure 11: New Worksheet for PivotTable

The PivotTable is displayed in the new worksheet.

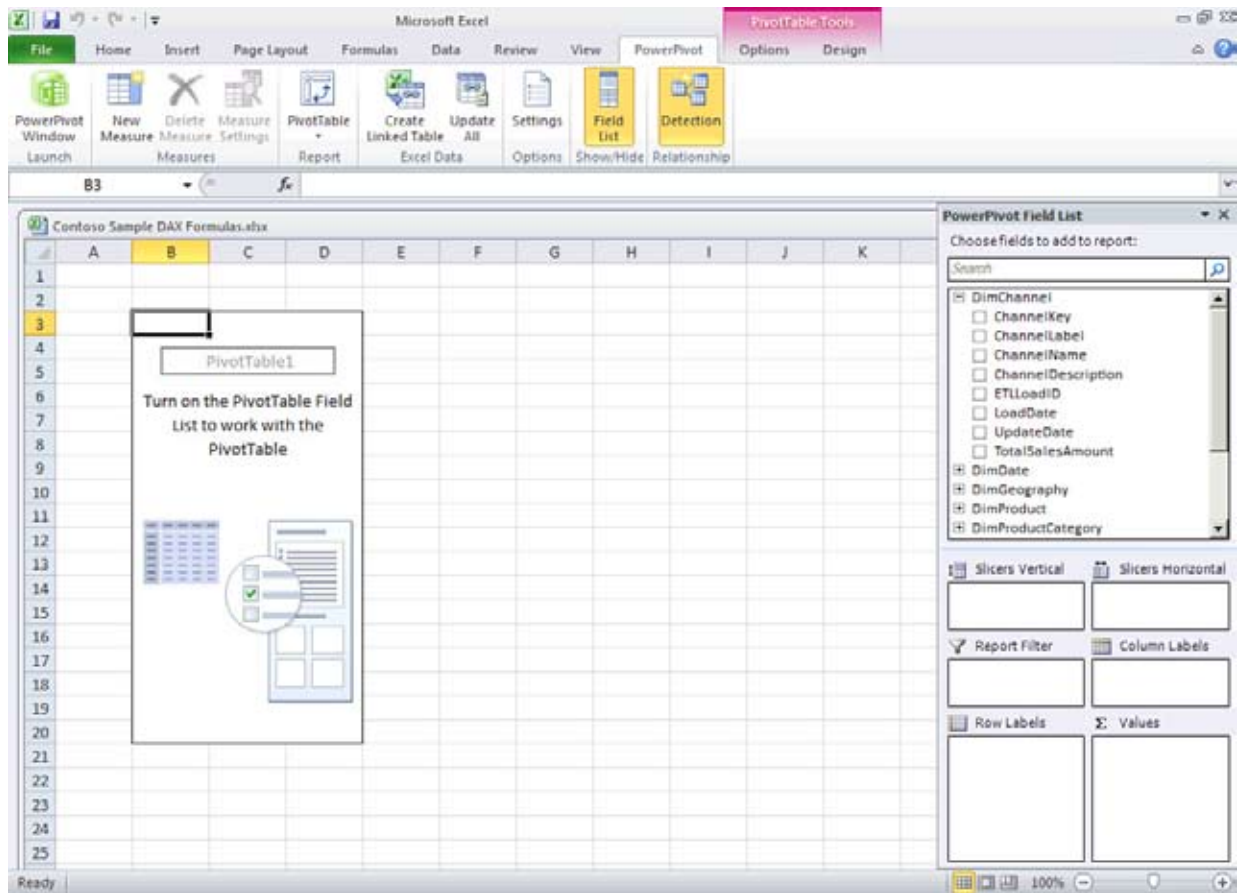


Figure 12: PivotTable with Field List

The PivotTable user interface is essentially standard Excel functionality with the Field List coming from the Data Sources in the PowerPivot Window rather than from a local spreadsheet. In this way, the analyst can use their existing skills to develop results while PowerPivot accesses the remote data complexities “under the hood.”

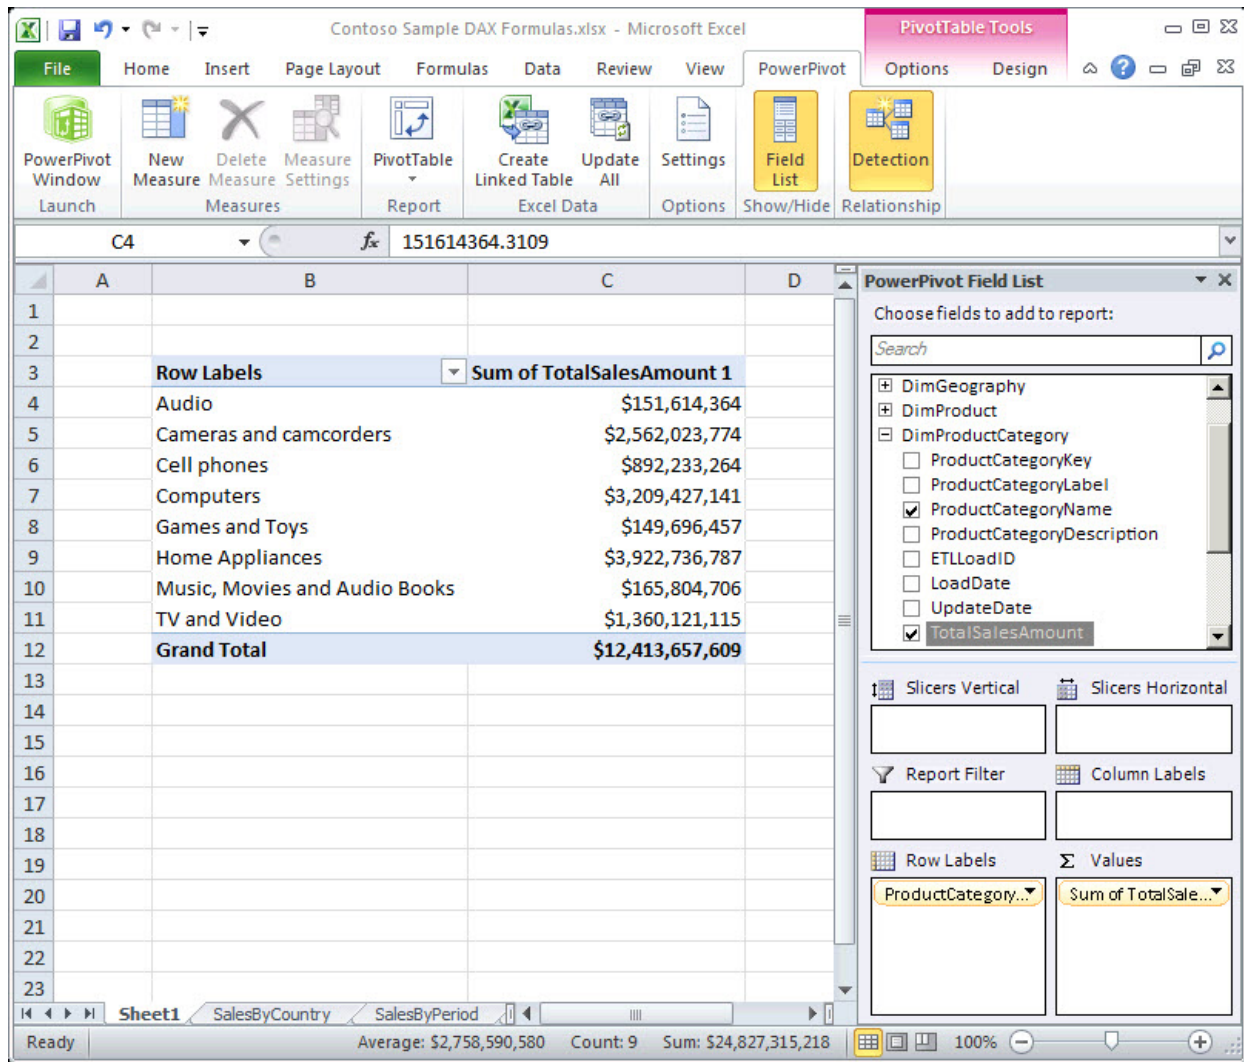


Figure 13: Pivot Table

Here, we are generating a simple PivotTable to show the total sales by Product Category. Filters and Slicers can also be defined to generate powerful results.

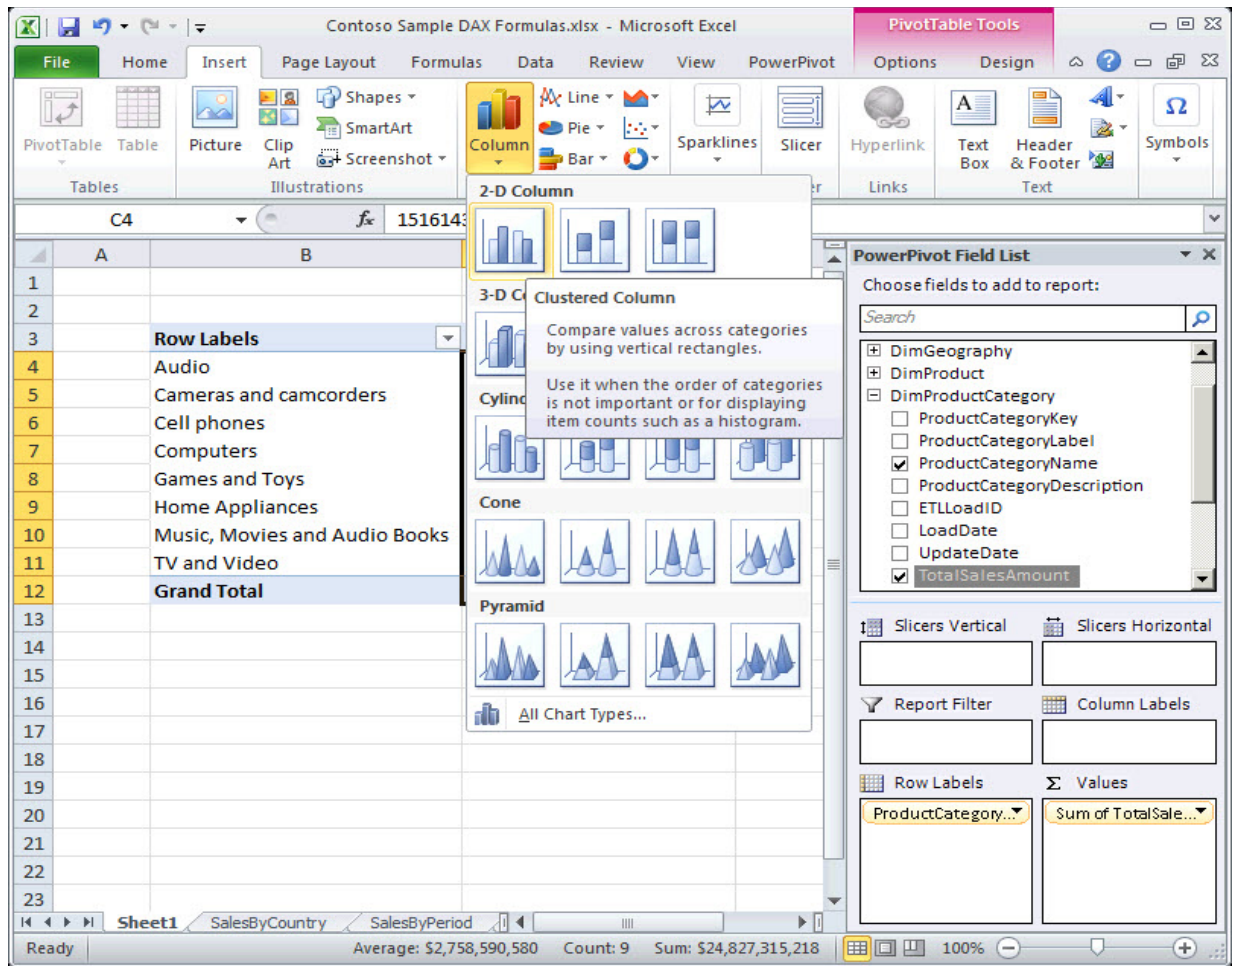


Figure 14: PivotChart

PivotCharts can be added to the workbook at any time, again, with the familiar Excel functionality and user interface. Here we are using the Insert ribbon within the worksheet. A new capability supported in Office 2010 is the sparkline feature that allows complex charts to be simplified by generating "mini-graphs" within cells to show individual trends.

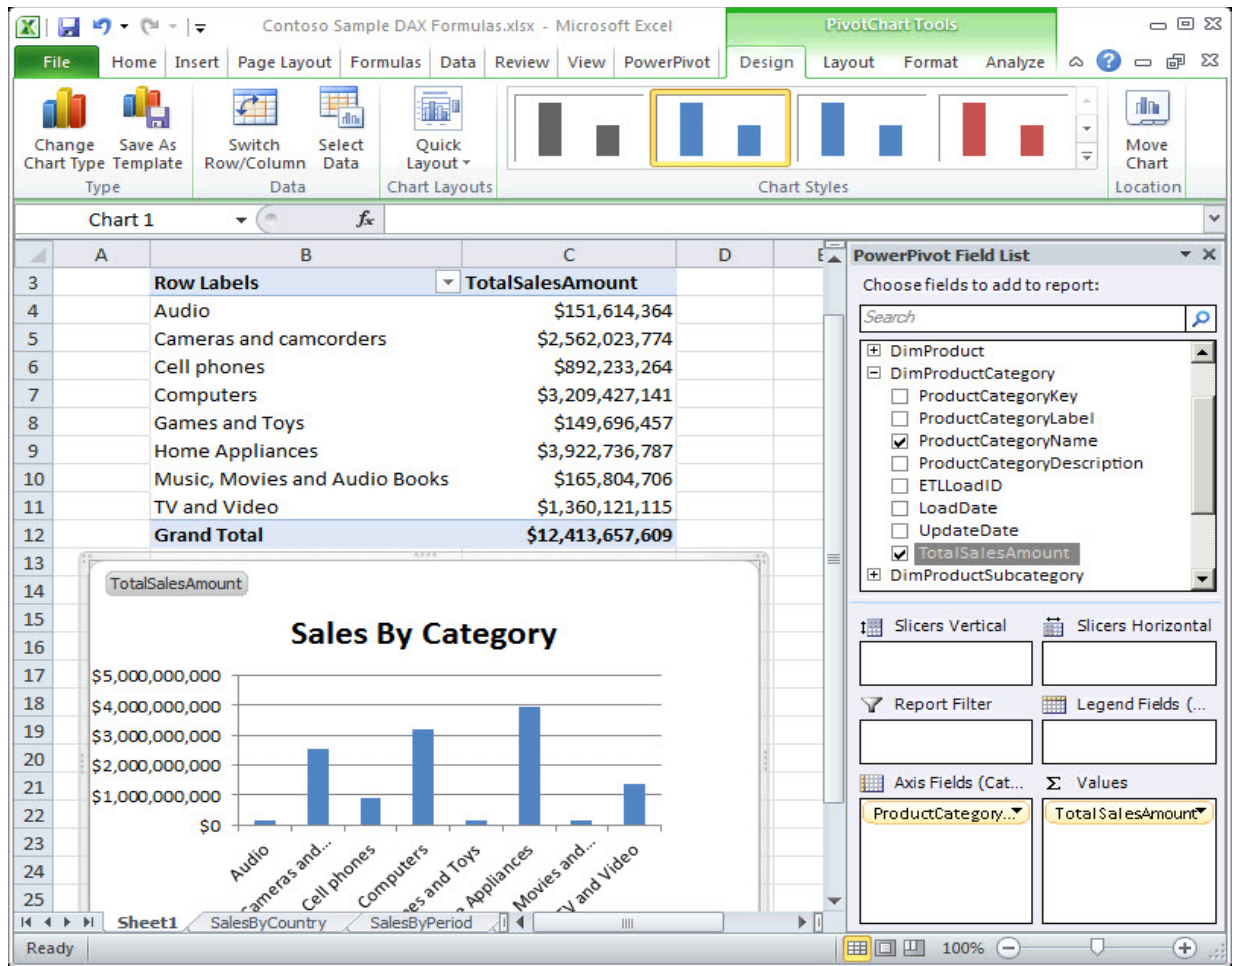


Figure 15: Pivot Table and Chart

The PivotTable and Chart can be saved and published together using Excel Services and PowerPivot for Sharepoint.

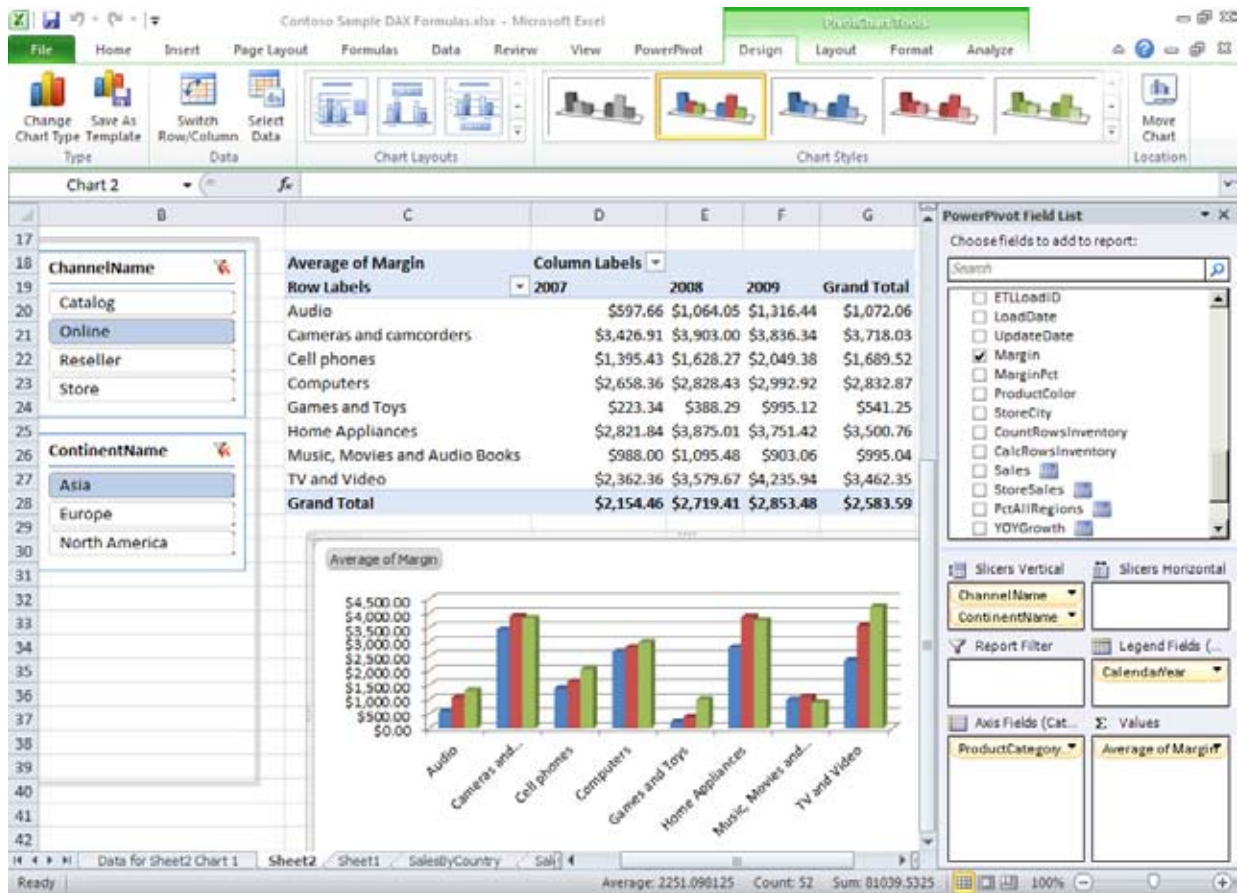


Figure 16: Example PowerPivot results

Slicers can be added to provide interactive functionality for the end user. In this case, we are analyzing average Margins Year over Year using Channel and Continent as slicers. These controls would be available to the Sharepoint users when published using PowerPivot for Sharepoint 2010.

Conclusion

PowerPivot is Microsoft's new technology that leverages Microsoft SQL Server 2008 R2, Office 2010, and Sharepoint 2010 to provide an exciting new Self Service Business Intelligence platform. The ease of use is impressive for the end user with the power and complexity being handled dynamically by the underlying integration of the software.

Learn More

Learn more about how you can improve productivity, enhance efficiency, and sharpen your competitive edge. Check out the following Global Knowledge course(s):

[SQL Server 2008 for Administration \(M6231, M6232\)](#)

[Implementing and Maintaining SQL Server 2008 Analysis Services and Business Intelligence Solutions \(M6234\)](#)

For more information or to register, visit www.globalknowledge.com or call **1-800-COURSES** to speak with a sales representative.

Our courses and enhanced, hands-on labs and exercises offer practical skills and tips that you can immediately put to use. Our expert instructors draw upon their experiences to help you understand key concepts and how to apply them to your specific work situation. Choose from our more than 1,200 courses, delivered through Classrooms, e-Learning, and On-site sessions, to meet your IT and business training needs.

About the Author

Brian Egler is a Global Knowledge instructor and Course Director specializing in Microsoft SQL Server technologies. He currently resides in Holly Springs, North Carolina.