


APPLIED MICROSOFT SQL SERVER 2012 ANALYSIS SERVICES

Tabular Modeling



Teo Lachev

WITH
VIDEO DEMOS

Applied Microsoft SQL Server 2012 Analysis Services

Tabular Modeling

Teo Lachev



Prologika Press

Applied Microsoft SQL Server 2012 Analysis Services Tabular Modeling

Published by:
Prologika Press
info@prologika.com
http://www.prologika.com

Copyright © 2012 Teo Lachev

All rights reserved. No part of this book may be reproduced, stored, or transmitted in any form or by any means, without the prior written permission of the publisher. Requests for permission should be sent to *info@prologika.com*.

Trademark names may appear in this publication. Rather than use a trademark symbol with every occurrence of a trademarked name, the names are used strictly in an editorial manner, with no intention of trademark infringement. The author has made all endeavors to adhere to trademark conventions for all companies and products that appear in this book, however, he does not guarantee the accuracy of this information.

The author has made every effort during the writing of this book to ensure accuracy of the material. However, this book only expresses the author's views and opinions. The information contained in this book is provided without warranty, either express or implied. The author, resellers or distributors, shall not be held liable for any damages caused or alleged to be caused either directly or indirectly by this book.

Printed in the United States of America
ISBN 13, print edition 978-0-9766353-5-2
ISBN 10, print edition 0-9766353-5-6
First printing 2012

Author:	Teo Lachev
Editor:	Edward Price
Technical Reviewers:	Burzin Mehernosh Daruwalla Greg Galloway Jen Underwood Norma O'Harra
Cover Designer:	Zamir Creations

contents

preface vii

acknowledgements viii

about the book ix

source code x

about the author xii

1 Introducing Business Intelligence Semantic Model 1

- 1.1 What is Business Intelligence Semantic Model? 1
 - Understanding Analysis Services History 3 •
 - Understanding the Microsoft BI Platform 4 • Understanding the BISM Components 9
- 1.2 Understanding Tabular 11
 - Why We Need Tabular 11 • Understanding Tabular Design Goals 12 • Understanding Tabular and VertiPaq Implementations 14
- 1.3 When to Use Tabular 15
 - Personal Business Intelligence 16 • Team Business Intelligence 18 •
 - Organizational Business Intelligence 20 • Comparing Tabular Features 21
- 1.4 Comparing Tabular with Other Models 22
 - Comparing Tabular and Multidimensional 22 •
 - Comparing Tabular and Report Builder Models 26
- 1.5 Summary 28

PART 1 PERSONAL BI INTELLIGENCE 29

2 Personal BI Basics 31

- 2.1 Introducing Personal Business Intelligence 31
 - About Organizational BI 32 • Understanding Personal BI 34
- 2.2 Understanding PowerPivot for Excel 35
 - Installing PowerPivot for Excel 35 • Getting to Know PowerPivot for Excel 38 •
 - Understanding PowerPivot Models 41 • Comparing Excel Reporting Options 43
- 2.3 Applied PowerPivot 45
 - Introducing Adventure Works Organizational BI 46 • Implementing a PowerPivot Model 46 • Analyzing Data 53
- 2.4 Summary 58

3 Importing Data 59

- 3.1 Data Modeling Concepts 59
 - Understanding Schema Types 59 • Introducing Table Relationships 61
- 3.2 Understanding Data Access Options 62
 - Introducing the Table Import Wizard 63 • Importing from Relational Databases 64 •
 - Importing from Multidimensional Databases 70 • Importing from Text Files 72 •
 - Importing from Data Feeds 75 • Other Import Options 79

3.3	Importing Adventure Works Data	81
	Importing from a Data Warehouse Database	81 • Importing from a Multidimensional Cube
	83 • Importing From Files	85 • Importing from Reports
	87 • Copying and Pasting Data	89
3.4	Summary	90
4	<i>Refining the Model</i>	91
4.1	Understanding Tables and Columns	91
	Understanding the Data View	92 • Exploring Data
	93 • Understanding Column Data Types	95 • Understanding Column Operations
	97 • Working with Tables and Columns	100
4.2	Managing Schema and Data Changes	103
	Managing Connections and Tables	103 • Managing Data Refresh
	105 • Performing Management Tasks	107
4.3	Relationships Revisited	108
	Relationship Rules and Limitations	108 • Auto-detecting Relationships
	109 • Creating Relationships Manually	111 • Understanding Diagram View
	114 • Working with Relationships	116
4.4	Summary	123
5	<i>Analyzing Data</i>	124
5.1	Understanding Excel Data Analysis	124
	Understanding Pivot Reports	125 • Understanding Report Layouts
	126 • Understanding PowerPivot Field List	129 • Understanding Integration with Excel
	132 • Understanding Slicers	134 • Understanding OLAP Pivot Features
	136	
5.2	Implementing Basic Pivot Reports	138
	Implementing a PivotTable Report	138 • Implementing a PivotChart Report
	141	
5.3	Implementing Dashboards	141
	Creating a Dashboard Report	142 • Working with Slicers
	146	
5.4	Implementing End-User Features	147
	Working with Dates	147 • Working with Perspectives
	150 • Working with Hierarchies	152
5.5	Summary	155
6	<i>Implementing Calculations</i>	156
6.1	Understanding Data Analysis Expressions	156
	Understanding DAX Calculations	157 • Understanding DAX Syntax
	160 • Understanding DAX Functions	161
6.2	Implementing Calculated Columns	170
	Creating Basic Calculated Columns	170 • Creating Advanced Calculated Columns
	174	
6.3	Implementing Measures	176
	Implementing Implicit Measures	176 • Implementing Explicit Measures
	179 • Implementing Key Performance Indicators	185
6.4	Implementing Advanced Relationships	187
	Implementing Role-Playing Relationships	188 • Implementing Parent-Child Relationships
	189 • Implementing Many-to-Many Relationships	191
6.5	Summary	193

PART 2 TEAM BUSINESS INTELLIGENCE 195

7 Team BI Basics 197

- 7.1 Understanding Team Business Intelligence 197
 - Introducing Microsoft SharePoint 198 • Understanding PowerPivot for SharePoint 200 • Understanding Runtime Interaction 201
- 7.2 Installing PowerPivot for SharePoint 203
 - Understanding Prerequisites 203 • Installing Active Directory Domain Services 203 • Installing SharePoint Server 206 • Configuring PowerPivot for SharePoint 209 • Upgrade Considerations 212
- 7.3 Verifying the SharePoint Setup 213
 - Publishing PowerPivot Models 213 • Troubleshooting Techniques 217
- 7.4 Summary 219

8 SharePoint Insights 220

- 8.1 Understanding SharePoint Insights 220
 - Understanding Integration Options 221 • Comparing Report Tools 222
- 8.2 Working with Excel Reports 223
 - Working with Excel Web Reports 223 • Working with Advanced Publishing Options 226 • Authoring Excel OLAP Reports 229
- 8.3 Working with Power View Reports 231
 - Understanding Power View 231 • Installing Reporting Services 232 • Understanding Reporting Properties 236 • Implementing the Sales by Country View 238 • Implementing the Sales Person Performance View 245
- 8.4 Authoring Operational Reports 248
 - Getting Started with Report Builder 248 • Working with Report Builder 252
- 8.5 Summary 255

9 Managing PowerPivot for SharePoint 256

- 9.1 Managing PowerPivot Workbooks 256
 - Managing Application Settings and Files 256 • Managing Data Sources 259 • Configuring Version History 262 • Approving Published Models 263
- 9.2 Managing Data Refresh 266
 - Understanding Data Refresh 266 • Configuring Data Refresh 267
- 9.3 Monitoring Server Utilization 272
 - Understanding PowerPivot Management Dashboard 272 • Analyzing Utilization and Activity 275
- 9.4 Managing Security 278
 - Understanding SharePoint Security 278 • Securing PowerPivot Models and Reports 281
- 9.5 Summary 285

PART 3 ORGANIZATIONAL BUSINESS INTELLIGENCE 287

10 Organizational BI Basics 289

- 10.1 Understanding Organizational Tabular Models 289
 - Understanding Professional Features 289 • Understanding the Professional Toolset 290 • Understanding the Development Lifecycle 292 • Installing Analysis Services in Tabular Mode 295
- 10.2 Designing Models 298
 - Getting Started with Project Development 298 • Importing from PowerPivot 301 • Implementing New Models 311 • Importing from Server 313
- 10.3 Authoring Reports 314
 - Implementing a Dashboard 314 • Introducing DAX Queries 318
- 10.4 Summary 324

11 Designing Storage and Security 325

- 11.1 Designing VertiPaq Storage 325
 - Understanding Partitions 326 • Understanding Processing 330 • Working with Partitions 333
- 11.2 Designing DirectQuery Storage 335
 - Understanding DirectQuery 336 • Implementing Low-Latency Models 340
- 11.3 Designing Data Security 343
 - Understanding Roles 343 • Understanding Row Filters 346 • Testing Database Roles 349 • Implementing Dynamic Data Security 350 • Externalizing Security Policies 352 • Advanced Security Scenarios 355
- 11.4 Summary 358

12 Managing Tabular Models 359

- 12.1 Management Fundamentals 359
 - Managing the Server 359 • Managing Memory 363 • Managing Databases 365
- 12.2 Managing Deployment 369
 - Working with the Deployment Wizard 369 • Synchronizing Databases 372
- 12.3 Managing Processing 374
 - Scripting Processing Jobs 374 • Automating Processing with Integration Services 376
- 12.4 Monitoring the Server 379
 - Tracing With SQL Server Profiler 379 • Monitoring Performance Counters 382 • Using Dynamic Management Views 384 • Error Logging and Reporting 386
- 12.5 Programming Management Tasks 386
 - Programming with AMO 386 • Scripting with PowerShell 390
- 12.6 Summary 392

Glossary of Terms 395

Index 397

preface

At SQL Pass 2010 in November 2010, Microsoft announced that in the forthcoming release of SQL Server, PowerPivot has been extended with features for organizational BI under a new name – Business Intelligence Semantic Model (BISM). This news took many BI professionals, including your humble correspondent, by surprise. Prior to PASS, Microsoft disclosed that organizational business intelligence (BI) would be a major focus in Denali, which I interpreted to mean that they'd add enhancements to multidimensional cubes, only to find out that all the buzz was about PowerPivot and BISM.

I have to admit that I was disappointed at first. I felt that Microsoft was shifting their focus in the wrong direction. Not surprisingly, there was a plenty of negativity from influencers and the community. Some pronounced UDM dead. Others sneered at BISM as a toy. To its credit, Microsoft was quick to cover the collateral damage and to clarify the BISM positioning. They made it clear that BISM doesn't intend to replace OLAP and multidimensional cubes. Later, at TechEd North America 2011 in Atlanta, Microsoft rebranded BISM as an umbrella name for both the multidimensional (Multidimensional) and tabular (Tabular) paths to further reinforce this message.

So, how did I convert from a Tabular critic to its aficionado and decided to write a book about it? During the year after PASS, I've exchanged plenty of e-mails with the product groups and Microsoft was patient enough to listen to me and address my concerns. Their answers helped me to "get it" and see the big picture through more optimistic lenses. Tabular is essentially a fresh start for Microsoft in the BI space. It has a great potential because it aims to simplify OLAP, and it makes BI even more accessible – "to the masses, by the masses", as the popular Microsoft slogan goes. Its elegance and versatility open new opportunities outside the core OLAP model, especially in the area of personal BI and team BI.

At the same time, however, I think that Microsoft got overexcited with Tabular and tilted too fast toward. They were missing the required balance at the expense of Multidimensional. Multidimensional has a very large install base, and it's still the preferred choice for data warehousing and OLAP. I personally believe that the Tabular journey will be evolutionary and that the market will decide the adoption pace.

This book attempts to restore that balance. I worked closely with Microsoft's Analysis Services team to provide an authoritative (yet independent) view of this technology to help you understand where and how to use it. Over the past two years, I've gathered plenty of real-life experience from BI solutions that used PowerPivot in one form or another. I've taught private and public classes and presented Tabular at conferences and local community chapters.

I decided to write this book to share with you the knowledge I harvested from my work with Tabular and to help you use it appropriately and efficiently. As its name suggests, the main objective of *Applied Microsoft SQL Server 2012 Analysis Services (Tabular Modeling)* is so to teach you the practical skills to implement Tabular-centric solutions for personal, team, and organizational BI.

Although this book is designed as a comprehensive guide to this technology, it's likely that you might have further questions or comments. As with my previous books, I am committed to help my readers with book-related questions via the book discussion forums on my company's web site, www.prologika.com. Happy modeling with Tabular!

Teo Lachev
Atlanta, GA

acknowledgements

Another SQL Server release and yet another book! Writing this one has been a lot of fun and a lot of work. It would not have been a reality without the help of many people to whom I am thankful. First and foremost, I'd like to thank my family for their ongoing support. They had to tolerate my long absence and put up with more than they had to. To my family I owe my greatest thanks.



Here is how my 8-year old son, Martin, and 11-year old daughter, Maya, portrait their absent father during his book writing project.

As an MVP, I've been privileged to enjoy close relationships with the SQL Server product groups. If it wasn't for their support, this book wouldn't have been a reality. Not only did they not mind my constant pestering and nagging, but they were even more eager to help me understand the "dark side" of Tabular! Akshai Mirchandani, Amir Netz, Ashvini Sharma, Cathy Dumas, Dave Wickert, Denny Lee, Edward Melomed, Jeffrey Wang, Julie Strauss, Kasper de Jonge, Marius Dumitru, Thomas Keiser, and T.K. Anand have contributed plenty of their time to educate me. The "insiders" lists rock! Special thanks to Akshai Mirchandani, Cathy Dumas, and Nickolai Medveditskov, for reviewing parts of the book and for helping ensure that it's technically accurate.

Writing a technical book while the product is still in development is like trying to hit a moving target because the product is constantly changing. Kudos to my technical reviewers, Burzin Mehernosh Daruwalla, Greg Galloway, Jen Underwood, and Norma O'Harra, for meticulously reviewing the manuscript and for proving feedback on how to improve it.

Thanks to my editor, Ed Price, from the SQL Server User Education team, for not losing faith that my incoherent readings could turn into something readable. Thank you for taking an extra mile and going beyond just checking the grammar. Your commitment and work was outstanding!

Finally, thank you for purchasing this book! I sincerely hope that you'll find it as enjoyable to read as it has been for me to write!

about the book

The book doesn't assume any prior experience with Microsoft Analysis Services or PowerPivot. It's designed as an easy-to-follow guide for navigating the personal-team-organizational BI continuum with Tabular. It starts by introducing you to the Microsoft BI platform and to the Business Intelligence Semantic Model (BISM). You need to know that each chapter builds upon the previous ones, to introduce new concepts and to extend the sample Adventure Works model that you'll implement in the exercises. Therefore, I'd recommend you read the chapters and do the exercises in the order they appear in the book.

Part 1, *Personal Business Intelligence*, teaches business users how to implement personal BI models with PowerPivot. Chapter 2, *Personal BI Basics*, lays out the foundation of personal BI, and it provides the necessary technical background to understand PowerPivot models. In Chapter 3, *Importing the Data*, you'll witness the impressive data acquisition capabilities of PowerPivot. In Chapter 4, *Refining the Model*, you'll explore and enhance the raw model to make it suitable for reporting. Chapter 5, *Analyzing Data*, shows you how to build interactive Excel pivot reports to gain data insights. And, in Chapter 6, *Implementing Calculations*, you'll further extend the model with useful business calculations.

Part 2, *Team Business Intelligence*, teaches business users and IT professionals how to take advantage of the SharePoint document management and collaboration features to share PowerPivot models with other team members. Chapter 7, *Team BI Basics*, introduces you to the SharePoint Products and Technologies and teaches you how to install and configure PowerPivot for SharePoint. Written for business users, Chapter 8, *SharePoint Insights*, walks you through the steps for authoring interactive and operational reports. In Chapter 9, *Managing PowerPivot for SharePoint*, IT administrators will learn how to manage published models, configure data refresh, secure user access, and monitor the server utilization and PowerPivot activity.

Part 3, *Organizational Business Intelligence*, shows BI pros how to bring tabular models to the enterprise. Chapter 10, *Organizational BI Basics*, introduces you to organizational BI and shows you different ways to implement tabular models. In Chapter 11, *Designing Storage and Security*, you'll learn how to extend your models with scalability, security, and low-latency features. In Chapter 12, *Managing Tabular Models*, you'll wear administrator hats and learn how to manage deployed models and how to perform common management tasks, including backing up databases, scripting objects, deploying changes, automating and programming management jobs.

source code

Applied Microsoft SQL Server 2012 Analysis Services demonstrates the full capabilities of Tabular for implementing personal, team, and organizational BI models. This requires installing and configuring various software products and technologies. **Table 1** lists the main software that you need for all the exercises in the book. Depending on your computer setup, you might need to download and install other components, as I explain throughout the book.

Table 1 The complete software requirements for practices and code samples in the book

Software	Required	Purpose	Chapters
Excel 2010 with PowerPivot for Excel	Yes	Implementing personal BI models.	2, 3, 4, 5, 6
SharePoint Server 2010 Enterprise with Reporting Services and PowerPivot for SharePoint	Yes	Implementing team BI models	7, 8, 9
Analysis Services Tabular with SQL Server Developer Tools Developer, Business Intelligence, or Enterprise edition	Yes	Implementing organizational BI models	10, 11, 12
SQL Server Database Engine Developer, Standard, or Enterprise 2008 R2 or later with the AdventureWorksDW2008R2 database	Yes	Importing and processing data	All
Analysis Services Multidimensional Developer, Standard, or Enterprise 2008 or later with the Adventure Works DW 2008 cube	No	Importing multidimensional data	3
Reporting Services Developer, Standard, or Enterprise 2008 R2 or later	No	Importing report feeds	3
Visual Studio 2010 Professional	No	Implement C# console application	12

Although the list is long, don't despair! The book includes step-by-step instructions for installing the first three software prerequisites, PowerPivot for Excel, PowerPivot for SharePoint, and Analysis Services in Tabular mode, in Chapters 2, 7, and 10 respectively. The book also provides alternative steps to complete the exercises if you don't install Analysis Services and Reporting Services.

You can download the book source code from the book page at <http://bit.ly/thebismbook>. After downloading the zip file, extract it to any folder of your hard drive. Once this is done, you'll see a folder for each chapter that contains the source code for that chapter. The source code in each folder includes the changes you need to make in the exercises in the corresponding chapter, plus any supporting files required for the exercises. For example, the Adventure Works.xlsx workbook in the Ch03 folder includes the changes that you'll make during the Chapter 3 practices and includes additional files for importing data. Save your files under different names or in different folders in order to avoid overwriting the files that are included in the source code.

The source code was tested with the Release Candidate 0 (RC0) build of SQL Server 2012 which was released in November 2011. Microsoft announced that this is a production-ready and feature-complete build.



NOTE The data source settings of the sample models in this book have connection strings to databases and text files. If you decide to test the finished samples, you need to update all the data sources to reflect your specific setup. To do so, open the Adventure Works workbook in Excel, click the PowerPivot ribbon tab (you must have PowerPivot for Excel installed first), and then click the PowerPivot Window button. In the PowerPivot window (Design tab), click the Existing Connections button. Click each connection, and then click Edit. In the Edit Connection dialog box, change the connection settings as needed.

Installing the Adventure Works databases

Most of the code samples use the AdventureWorksDW2008R2 database. This is the R2 build of a Microsoft-provided sample database that simulates a data warehouse. The Product Catalog 2008 report in Chapter 3 also uses the AdventureWorks2008R2 database. However, if you don't have the AdventureWorks2008R2 database or Reporting Services, you can still follow this book because I provide alternative steps to import the required data from the data warehouse database (AdventureWorksDW2008R2). Microsoft has bundled all the Adventure Works databases into a single download.



NOTE Microsoft is working on an updated version of the Adventure Works databases for SQL Server 2012. The SQL Server 2012 versions are expected to include incremental changes from the R2 releases. They haven't been tested with this book's code samples. If you decide to install the SQL Server 2012 version of the sample databases, you might find that the practice steps differ somewhat from your setup and that the reports might show different data.

Follow these steps to download the Adventure Works databases:

1. Open the Microsoft SQL Server Product Samples Database webpage on Codeplex (<http://www.codeplex.com/MSFTDBProdSamples>).
2. Click the SQL Server 2008R2 link. The link URL as of the time of this writing is <http://bit.ly/adventureworksR2>. Click the AdventureWorks2008R2_SR1.exe link.
3. When Internet Explorer prompts you, click Run to download and execute the setup program.
4. The setup program includes all the Adventure Works databases. In the Adventure Works Community Sample Databases step of the setup, make sure that the "AdventureWorks OLTP 2008R2" and "AdventureWorks Data Warehouse 2008R2" are selected. If you plan to deploy the Adventure Works cube, select also the "AdventureWorks OLAP Standard 2008R2" database if you have SQL Server 2008 (or above) Standard edition. Or, check the "AdventureWorks OLAP Enterprise 2008R2" database if you have SQL Server Developer, Enterprise, or Business Intelligence edition. The setup installs only the project files and it doesn't deploy the cube (the cube deployment steps are provided in the "Installing the Adventure Works cube" section).
5. To verify that the relational databases are installed, open SQL Server Management Studio (SSMS), and then connect to the SQL Server instance that hosts the Adventure Works databases.
6. Expand the Databases folder. You should see the AdventureWorks2008R2 and AdventureWorksDW2008R2 databases.

Installing the Adventure Works cube

In Chapter 3, you import data from the AdventureWorks Analysis Services cube, but I include alternative import steps if installing the cube is not an option. If you decide to import from the cube, install the Analysis Services AdventureWorks solution as follows:

1. Open the AdventureWorks Analysis Services solution (Adventure Works.sln) in Business Intelligence Development Studio (if you have SQL Server 2008 R2) or SQL Server Data Tools (if you have SQL Server 2012). Here is default folder where the Analysis Services solution is located:

|Program Files\Microsoft SQL Server\100\Tools\Samples\AdventureWorks 2008 Analysis Services Project

2. In the Solution Explorer, right-click the project node, and then click Properties.
3. In the Property Pages dialog box, click the Deployment tab. In the Server field, type the Analysis Services instance to which the cube will be deployed, such as *localhost* if you've installed Analysis Services on the default instance on your local server. Click OK.
4. In Solution Explorer, right-click the project node, and then click Deploy to deploy the project.
5. In SQL Server Management Studio, connect to the Analysis Services instance (see Step 3).
6. Expand the Databases folder. You should see the Adventure Works DW 2008 database.

Watching the video demos

Designing a tabular model is a click-intensive process. I captured video demos to help you stay on track when a picture is worth more than a thousand words (sometimes you just have to see it to understand). The video demos are bonus material to the book. The play symbol (▶) next to a section title indicates that there is a video demo for this section. Video demos are provided for a subset of the practical exercises and are not intended to exactly match the steps included in the book. The book web page (<http://bit.ly/thebismbook>) provides a link to the video demos that you can watch online.

Reporting errors

This book has no bugs! Well, that's the goal, anyway. We both know that this statement is over-ambitious to say the least. Please submit bug reports to the book discussion list on <http://prologika.com/cs/forums/default.aspx>. Confirmed bugs and inaccuracies will be published to the book errata document. A link to the errata document is provided in the book web page.

The book includes links to web resources for further study. Due to the transient nature of the Internet, some links might be no longer valid or might be broken. Searching for the document title is usually sufficient to recover the new link.

Your purchase of APPLIED MICROSOFT SQL SERVER 2012 ANALYSIS SERVICES (Tabular Modeling) includes free access to a web forum sponsored by the author, where you can make comments about the book, ask technical questions, and receive help from the author and the community. The author is not committed to a specific amount of participation or successful resolution of the question and his participation remains voluntary. You can subscribe to the forum from the author's personal website <http://prologika.com/cs/forums/default.aspx>.

about the author

Teo Lachev (MVP, MCSD, MCT, MCITP) is a consultant, author, and mentor, with a focus on Microsoft Business Intelligence. His Atlanta-based company "Prologika" helps organizations make sense of data by effectively applying Microsoft BI technologies. Teo has authored and co-authored several BI books, including *Applied Microsoft Analysis Services 2005* and *Applied SQL Server 2008 Reporting Services*, and he has co-authored the *Microsoft SQL Server Business Intelligence Implementation and Maintenance Training Kit*. He speaks frequently at industry conferences and leads the Microsoft Business Intelligence Group in Atlanta, GA.