

preface

Analysis Services and the Microsoft Business Intelligence Platform share the common goal of making business data usable and accessible to as many people as possible and delivering on the Microsoft promise – “*BI for the Masses*”. Most businesses today have OLTP-based systems to capture raw data. This is usually the first IT requirement an organization has to tackle. The IT technology has matured to a point where implementing highly-transactional, highly-available applications is not an issue. The real predicament that most organizations face in the information age is analyzing the mountains of data that are piled up every day.

Think of the last time you had to look up a person’s contact. How long did it take you? Did you start with a stack of scattered name cards in your drawer, then checked the Outlook Contacts folder, and finally gave up after checking the contacts stored in your cell phone? Wouldn’t it be much faster and simpler if all these pieces of data were stored in one repository? Now, give this data predicament an enterprise “dimension” and you are getting closer to what a corporate business analyst has to deal with on a regular basis. Not to mention that the numbers often don’t tally from one reporting medium to another and there are different versions of the “truth” depending of what reporting tool is used.

Don’t get me wrong. Enterprise reporting plays an important role as a business intelligence vehicle. However, once you’ve successfully addressed the basic reporting requirements, you may realize that it is time to “graduate” to the next level – efficient, consistent and rich data analytics. This is where Microsoft Analysis Services and Microsoft Business Intelligence Platform come in. By allowing you to organize your data assets and build an efficient presentation layer on top of them, they allow you to get *out* what was put *in*. As a result, you spend less time in searching for information and have more time to act upon it and make decisions.

Microsoft Analysis Services has a very strategic role in the Microsoft Business Intelligence initiative. At the same time, however, Analysis Services is not a trivial technology. An Analysis Services-based solution is often the focal point of several IT processes, including OLAP, data mining, ETL, data warehousing, standard and ad hoc reporting, performance management, etc. Understanding how these technologies fit together is not an easy task. I wrote this book to help you in this endeavor.

As its name suggests, the main objective of *Applied Microsoft Analysis 2005* is to teach you the practical skills you need to implement Analysis Services-centric business intelligence solutions. This book doesn’t assume any prior experience with OLAP and Microsoft Analysis Services. It is designed as an easy-to-follow guide where each chapter builds upon the previous to implement the components of the innovative Unified Dimensional Model (UDM) in chronological order. New concepts are introduced with step-by-step instructions and hands-on demos.

Artists say that after Picasso’s Cubism, the world of art never looked the same. I hope that you will find Analysis Services 2005, code-named *Picasso*, as influential for transforming your company’s BI vision to the next level. Happy data analyzing!

Teo Lachev
Atlanta, GA

acknowledgements

Writing this book 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.

As with my previous book, first and foremost, I would like to thank my family for their ongoing support. They had to tolerate my total dedication to the book for almost a year! My wife, Elena, contributed directly to the book by converting my incoherent writings into a readable manuscript. To my family I owe my greatest thanks.

I would like to thank Ariel Netz, Group Program Manager with Microsoft Analysis Services for letting me harass his team with naïve questions. Special thanks to Akshai Mirchandani, Alex Bocharov, Amir Netz, Andrew Garbuzov, Dave Wickert, Jamie MacLennan, Marin Bezic, Mosha Pasumansky, Paul Sanders, Raman Iyer, Richard Tkachuk, Rob Zare, Sasha Berger, Siva Harinath, T.K. Anand, Thierry D'Hers, ZhaoHui Tang, and other Microsoft engineers who not only didn't mind my constant pestering, but were always willing to help me understand the "dark side" of OLAP and data mining.

Donald Farmer from the Integration Services team helped me tremendously with the Integration Services code samples. As with my first book, Brian Welcker and the Reporting Services team have been extremely accommodating and supportive. Special thanks to Jason Carlson and Carolyn Chao for addressing tough questions about ad hoc reporting and Rajeev Karunakaran for helping me with the report viewer controls.

Kudos to my technical reviewer, Alexander Nepomnjashiy (Alex), for meticulously reviewing the manuscript consistency and code accuracy. As with my previous book, Alexander helped me tremendously in transitioning the book from the yearly builds of Analysis Services to its final release – a long, painstaking and sometimes frustrating process.

Many, many thanks to my peer reviewers for supporting me with valuable feedback and encouragement. Marco Russo, Dejan Sarka (MVP – SQL Server), and Deepak Puri (MVP – SQL Server) and have been phenomenal in sharing their extensive OLAP experience and suggestions for improvement. Stacy Parish, Brendon Schwartz, and Sorin Belean helped me improve the book readability and technical accuracy.

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

about the book

Chapter 1 provides a panoramic overview of Microsoft Analysis Services 2005 and its architecture. You will understand what UDM is and how it bridges the “great divide” between the relational and dimensional reporting models. We will finish with a hands-on lab that demonstrates how you can produce an Excel-based report that sources data from an Analysis Services cube.

Part 1, *The Unified Dimensional Model*, is meant for database modelers and developers that will implement the core dimensional model of UDM. Chapter 2 teaches you how to work with data and set up data sources and data source views. Chapter 3 introduces you to the dimensional model and shows you how to use the Cube Wizard and Dimension Wizard to build quickly the raw UDM. In chapter 4, you will refine the UDM and get more insight into the UDM building blocks – dimensions and measures. In chapter 5, we will look at some innovative features that debut in Analysis Services 2005, such as dimension relationships, semi-additive measures, linked cubes and measure groups, and error configurations.

Part 2, *Data Warehousing and Mining*, discusses the two most common usage scenarios for using Analysis Services. Chapter 6 shows you how to build a cube top-down by starting with the cube schema and implementing an ETL package to load the cube with data. Chapter 7 introduces you to the powerful data mining features in UDM and shows you how to build mining models on top of relational data sources. Chapter 8 expands the data mining concepts by teaching you how to use UDM as a data source for the mining models.

Part 3, *Programming Analysis Services*, targets developers who will be using MDX to extend the cube capabilities. Chapter 9 introduces you to MDX and walks you through many examples that demonstrate how you can navigate the dimensional model horizontally and vertically. In chapter 10, you will learn how to work with sets and subcube expressions. Chapter 11 teaches you how to extend UDM with custom code in the form of VBA functions and Analysis Services stored procedures. In chapter 11, you will build a rich end-user model by using KPIs, actions, perspectives, and translations.

Part 4, *Managing Analysis Services*, is for administrators who are responsible for managing Analysis Services servers. Chapter 13 introduces you to the management features of the SQL Server Management Studio and shows you how to perform routine tasks, such as backing up, synchronizing, deploying, and scripting databases. In chapter 14, you will learn how to work with the cube storage settings and improve the cube performance with aggregations. Chapter 15 discusses the Analysis Services processing architecture by walking you through different operational scenarios. In chapter 16, you will understand how to secure Analysis Services by using role-based security policies.

Part 5, *Building Business Intelligence Solutions*, gives developers and power users the necessary background to build Analysis Services-centric solutions. Chapter 17 teaches you how to leverage MDX and DMX to build custom business intelligence applications. In chapter 18, you will build standard and ad hoc reports with Microsoft Reporting Services. Finally, in Chapter 19, you will learn how to use Microsoft Office analytics tools to implement interactive reporting solutions and the Microsoft Office Business Scorecard Manager 2005 to build dashboards for corporate performance management.

source code

Applied Microsoft Analysis 2005 introduces the new concepts with hands-on labs where each chapter builds upon the previous to implement a complete UDM solution. Table 1 lists the software requirements to run all the code samples included in the book.

Table 1 Software requirements for running the book source code.

Software	Purpose
SQL Server 2005 Developer or Enterprise Edition	These two editions provide all Analysis Services features. Use the Developer Edition on Windows XP and Developer or Enterprise Edition if you are running Windows 2003.
Microsoft Visual Studio 2005	To edit the Visual Basic code samples and SSAS stored procedures.
Microsoft Office 2003	To use the analytics tools for building reports (Excel and OWC) in Chapter 19.
Microsoft Office Business Scorecard Manager 2005 and Windows Server 2003 with SharePoint Services	To build a performance scorecard in Chapter 19.

The code samples can be downloaded from the book web page at <http://www.prologika.com>. After downloading the zip file, extract it to any folder of your hard drive. Once this is done, you will see a folder for each chapter. The UDM code samples use the AdventureWorks and AdventureWorksDW sample databases that come with SQL Server 2005. Follow the instructions in Appendix A to install these databases. For the hands-on lab in chapter 1, you will also need the sample Adventure Works DW UDM which is included in Analysis Services 2005 samples. Chapter 1 provides instructions about how to deploy the Adventure Works DW solution.

about the author

Teo Lachev has more than a decade of experience in designing and developing Microsoft-centric solutions. He is a technical architect for a leading financial institution where he designs and implements Business Intelligence applications for the banking industry. Teo is a Microsoft Most Valuable Professional (*MVP*) for SQL Server, Microsoft Certified Solution Developer (*MCSD*), and Microsoft Certified Trainer (*MCT*). He is the author of *Microsoft Reporting Services in Action* (Manning Publications, 2004). Teo lives with his family in Atlanta, GA.

Your purchase of *Applied Microsoft Analysis Services 2005* 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 www.prologika.com.