




The Microsoft® Data Warehouse Toolkit

Second Edition


With SQL Server®
2008 R2 and the
Microsoft Business
Intelligence Toolset

Joy Mundy and
Warren Thornthwaite
with Ralph Kimball



The Microsoft® Data Warehouse Toolkit

Second Edition



The Microsoft® Data Warehouse Toolkit

**With SQL Server 2008 R2
and the Microsoft® Business
Intelligence Toolset**

Second Edition

Joy Mundy and
Warren Thornthwaite
with Ralph Kimball



WILEY

Wiley Publishing, Inc.

**The Microsoft® Data Warehouse Toolkit: With SQL Server 2008 R2
and the Microsoft® Business Intelligence Toolset, Second Edition**

Published by
Wiley Publishing, Inc.
10475 Crosspoint Boulevard
Indianapolis, IN 46256
www.wiley.com

Copyright © 2011 by Joy Mundy and Warren Thornthwaite with Ralph Kimball
Published by Wiley Publishing, Inc., Indianapolis, Indiana

Published simultaneously in Canada

ISBN: 978-0-470-64038-8
ISBN: 978-1-118-06793-2 (ebk)
ISBN: 978-1-118-06795-6 (ebk)
ISBN: 978-1-118-06794-9 (ebk)

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 646-8600. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

Limit of Liability/Disclaimer of Warranty: The publisher and the author make no representations or warranties with respect to the accuracy or completeness of the contents of this work and specifically disclaim all warranties, including without limitation warranties of fitness for a particular purpose. No warranty may be created or extended by sales or promotional materials. The advice and strategies contained herein may not be suitable for every situation. This work is sold with the understanding that the publisher is not engaged in rendering legal, accounting, or other professional services. If professional assistance is required, the services of a competent professional person should be sought. Neither the publisher nor the author shall be liable for damages arising herefrom. The fact that an organization or Web site is referred to in this work as a citation and/or a potential source of further information does not mean that the author or the publisher endorses the information the organization or website may provide or recommendations it may make. Further, readers should be aware that Internet websites listed in this work may have changed or disappeared between when this work was written and when it is read.

For general information on our other products and services please contact our Customer Care Department within the United States at (877) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.

Library of Congress Control Number: 2011920894

Trademarks: Wiley and the Wiley logo are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates, in the United States and other countries, and may not be used without written permission. Microsoft is a registered trademark of Microsoft Corporation. All other trademarks are the property of their respective owners. Wiley Publishing, Inc. is not associated with any product or vendor mentioned in this book.



About the Authors

Joy Mundy has focused on DW/BI systems since 1992 with stints at Stanford, WebTV, and Microsoft's SQL Server product development organization. Joy graduated from Tufts University with a BA in Economics, and from Stanford University with an MS in Engineering Economic Systems.

Warren Thornthwaite began his DW/BI career in 1980. After managing Metaphor's consulting organization, he worked for Stanford University and WebTV. Warren holds a BA in Communications Studies from the University of Michigan and an MBA from the University of Pennsylvania's Wharton School.

Ralph Kimball founded the Kimball Group. Since the mid 1980s, he has been the DW/BI industry's thought leader on the dimensional approach and has trained more than 10,000 IT professionals. Prior to working at Metaphor and founding Red Brick Systems, Ralph co-invented the Star workstation at Xerox's Palo Alto Research Center (PARC). Ralph has a Ph.D. in Electrical Engineering from Stanford University.



Credits

Executive Editor

Robert Elliott

Project Editors

Sara Shlaer

Ginny Munroe

Technical Editor

Ralph Kimball

Senior Production Editor

Debra Banninger

Copy Editor

Kim Cofer

Editorial Director

Robyn B. Siesky

Editorial Manager

Mary Beth Wakefield

Freelancer Editorial Manager

Rosemarie Graham

Marketing Manager

Ashley Zurcher

Production Manager

Tim Tate

Vice President and**Executive Group Publisher**

Richard Swadley

Vice President and**Executive Publisher**

Barry Pruett

Associate Publisher

Jim Minatel

Project Coordinator, Cover

Katie Crocker

Compositor

Craig Johnson,

Happenstance Type-O-Rama

Proofreader

Jen Larsen, Word One

Indexer

Robert Swanson

Cover Image

© Getty Images

Cover Designer

Ryan Sneed



Acknowledgments

First, we want to thank the thousands of you who have read the Kimball Group's *Toolkit* books, attended our courses, and engaged us in consulting projects. We always learn from you, and you've had a profound impact on our thinking and the business intelligence industry.

This book would not have been written without the assistance of many people on the SQL Server product development team. Dave Wickert reviewed the PowerPivot and SharePoint chapters and provided many excellent suggestions for improving the content. Bryan Smith was kind enough to read the chapters on Integration Services and Analysis Services, and those chapters are the better for his assistance. Carolyn Chau reviewed the Reporting Services chapter, Eric Hanson read the relational database chapter, Pej Javaheri commented on the SharePoint chapter, and Raman Iyer read the data mining chapter. Our sincere gratitude to all of them.

Other members of the SQL Server team provided significant assistance reviewing the SQL Server 2005 version of this book, and we were too embarrassed to impose on them a second time. These include Bill Baker, Stuart Ozer, Grant Dickinson, Donald Farmer, Siva Harinath, Jamie MacLennan, John Miller, Ashvini Sharma, Stephen Quinn, and Rob Zare.

Our colleagues at the Kimball Group were invaluable. Their encouragement kept us going while we were writing the book, and their reviews helped us polish and prune material. Ralph Kimball, of course, had a huge impact on the book, not just from his writing and thinking in the business intelligence arena but more directly by helping us improve the book's overall structure and flow.

Sara Shlaer, Ginny Munroe, and Bob Elliott, our editors at Wiley, have been very helpful and encouraging. It's been a pleasure to work with them.

To our life partners, thanks for being there when we needed you, for giving us the time we needed, and for occasionally reminding us that it was time to take a break. Tony Navarrete and Elizabeth Wright, the book wouldn't exist without you.



Contents at a Glance

Foreword		xxvii
Introduction		xxix
Part 1	Requirements, Realities, and Architecture	1
Chapter 1	Defining Business Requirements	3
Chapter 2	Designing the Business Process Dimensional Model	29
Chapter 3	The Toolset	79
Chapter 4	System Setup	99
Part 2	Building and Populating the Databases	133
Chapter 5	Creating the Relational Data Warehouse	135
Chapter 6	Master Data Management	165
Chapter 7	Designing and Developing the ETL System	187
Chapter 8	The Core Analysis Services OLAP Database	245
Chapter 9	Design Requirements for Real-Time BI	305
Part 3	Developing the BI Applications	321
Chapter 10	Building BI Applications in Reporting Services	323
Chapter 11	PowerPivot and Excel	375
Chapter 12	The BI Portal and SharePoint	403
Chapter 13	Incorporating Data Mining	429

Part 4	Deploying and Managing the DW/BI System	491
Chapter 14	Designing and Implementing Security	493
Chapter 15	Metadata Plan	523
Chapter 16	Deployment	545
Chapter 17	Operations and Maintenance	583
Chapter 18	Present Imperatives and Future Outlook	613
Index		625



Contents

Foreword	xxvii
Introduction	xxix
Part 1 Requirements, Realities, and Architecture	1
Chapter 1 Defining Business Requirements	3
The Most Important Determinant of Long-Term Success	5
Adventure Works Cycles Introduction	6
Uncovering Business Value	6
Obtaining Sponsorship	7
Defining Enterprise-Level Business Requirements	8
Establishing Initial Enterprise Requirements Project Scope	9
Gathering and Documenting Enterprise-Level Business Requirements	10
Data Auditing/Data Profiling	16
Creating the Program Requirements Findings Document	18
Prioritizing the Business Requirements	22
Revisiting the Project Planning	25
Gathering Project-Level Requirements	26
Summary	28
Chapter 2 Designing the Business Process Dimensional Model	29
Dimensional Modeling Concepts and Terminology	30
Facts	31
Just the Facts	32
The Grain	33
Dimensions	33
Bringing Facts and Dimensions Together	34
The Bus Matrix, Conformed Dimensions, and Drill Across	36

Additional Design Concepts and Techniques	38
Surrogate Keys	38
Slowly Changing Dimensions	39
Dates	42
Degenerate Dimensions	43
Snowflaking	43
Many-to-Many or Multivalued Dimensions	44
The Many-to-Many Relationship Between a Fact and Dimension	45
Many-to-Many Between Dimensions	46
Hierarchies	47
Variable-Depth Hierarchies	48
Frequently Changing Hierarchies	49
Aggregate Dimensions	49
Junk Dimensions	51
The Three Fact Table Types	52
Aggregates	53
The Dimensional Modeling Process	54
Preparation	55
Identify Roles and Participants	55
Revisit the Requirements	56
Understand the Data Architecture Strategy	58
Set Up the Modeling Environment	58
Establish Naming Conventions	59
Data Profiling and Research	60
Data Profiling and Source System Exploration	60
Building Dimensional Models	63
High-Level Dimensional Model Design Session	63
Identifying Dimension Attributes and Fact Measures	65
Developing the Detailed Dimensional Model	66
Testing and Refining the Model	68
Reviewing and Validating the Model	68
Case Study: The Adventure Works Cycles	
Orders Dimensional Model	69
The Orders Fact Table	69
The Dimensions	69
Employee	69
Customer and Reseller	70
Currency	71
Identifying Dimension Attributes and Facts for the Orders Business Process	72
The Final Draft of the Initial Orders Model	74
Detailed Orders Dimensional Model Development	75
Identifying SCD Change Types	75
Reviewing the Issues	76
Identifying the Facts	76
Final Dimensional Model	77
Summary	77

Chapter 3	The Toolset	79
	The Microsoft DW/BI Toolset	80
	Why Use the Microsoft Toolset?	82
	Architecture of a Microsoft DW/BI System	83
	Why Analysis Services?	84
	Why a Relational Store?	86
	ETL Is Not Optional	86
	The Role of Master Data Services	88
	Delivering BI Applications	88
	Overview of the Microsoft Tools	89
	Which Products Do You Need?	90
	SQL Server Development and Management Tools	92
	SQL Server Management Studio	92
	Business Intelligence Development Studio (BIDS)	95
	Summary	97
Chapter 4	System Setup	99
	System Sizing Considerations	100
	Calculating Data Volumes	101
	Determining Usage Complexity	102
	Simple or Controlled Access	102
	Moderate Complexity	103
	Highly Demanding Use	103
	Estimating Simultaneous Users	104
	Assessing System Availability Requirements	105
	How Big Will It Be?	105
	System Configuration Considerations	105
	Memory	106
	Monolithic or Distributed?	106
	Storage System Considerations	110
	Balancing the Data Pipeline	110
	Disk Performance	111
	Fault Tolerance and RAID	112
	Storage Area Networks	112
	Processors	113
	Setting Up for High Availability	114
	Software Installation and Configuration	115
	Development Environment Software Requirements	116
	Development Database Server	117
	Database Designer	117
	Development Database Administrator	118
	ETL System Developer	118
	Report Designer	119
	Reporting Portal Developer	119
	Data Mining Model Developer	120
	Analytic Application Developer	120

Test and Production Software Requirements	120
Database Server	121
Database Administrator	121
Operating Systems	122
SQL Server Relational Database Setup	122
Security Options During Installation	123
Files, Filegroups, and RAID	123
Database Recovery Model	125
Initial Database Size	126
Analysis Services Setup	126
Analysis Services File Locations and Storage Requirements	127
Analysis Services and Memory	128
Integration Services Setup	129
Integration Services File Locations and Storage Requirements	130
Reporting Services Setup	130
Summary	131
Part 2 Building and Populating the Databases	133
Chapter 5 Creating the Relational Data Warehouse	135
Getting Started	136
Complete the Physical Design	137
Surrogate Keys	138
String Columns	138
To Null, or Not to Null?	140
Housekeeping Columns	140
Table and Column Extended Properties	142
Define Storage and Create Constraints and Supporting Objects	142
Create Files and Filegroups	142
Data Compression	144
Entity and Referential Integrity Constraints	145
Initial Indexing and Database Statistics	147
Dimension Indexes	147
Fact Table Indexes	149
Statistics	149
Aggregate Tables	150
Create Table Views	151
Insert an Unknown Member Row	152
Example CREATE TABLE Statement	152
Partitioned Tables	153
How Does Table Partitioning Work?	154
Managing Partitioned Tables	158
Using Partitioned Tables in the Data Warehouse Database	162
Finishing Up	163
Staging Tables	163
Metadata Setup	163
Summary	164

Chapter 6	Master Data Management	165
	Managing Master Reference Data	166
	Incomplete Attributes	167
	Data Integration	168
	Systems Integration	170
	Master Data Management Systems and the Data Warehouse	171
	Introducing SQL Server Master Data Services	171
	Model Definition Features	172
	User Interface: Defining the Master Data Model	172
	Creating Database Structures	173
	Security	173
	Programmability of Model Definition	174
	Full Versioning of Models	174
	Data Management Features	174
	User Interface: Exploring and Managing the Master Data	174
	Importing and Updating Data	176
	Exporting Data	177
	Full Versioning of All Attributes	179
	Creating a Simple Application	179
	The Business Scenario	179
	Keep It Simple	180
	Create the MDS Model	180
	Load the Subcategory Members	183
	Polish the Model	185
	Export to the Data Warehouse	185
	Summary	186
Chapter 7	Designing and Developing the ETL System	187
	Round Up the Requirements	188
	Develop the ETL Plan	191
	Introducing SQL Server Integration Services	192
	Control Flow and Data Flow	194
	Data Flow	195
	Error Flows	196
	SSIS Package Architecture	197
	The Major Subsystems of ETL	198
	Extracting Data	199
	Subsystem 1: Data Profiling	199
	Subsystem 2: Change Data Capture System	200
	Subsystem 3: Extract System	202
	Loading Data	204
	Cleaning and Conforming Data	206
	Subsystem 4: Data Cleaning System	206
	Cleaning Data in the Data Flow	207
	Halting Package Execution	211

Subsystem 5: Error Event Schema	214
Subsystem 6: Audit Dimension Assembler	215
Subsystem 7: Deduplication System	216
Subsystem 8: Conforming System	217
Delivering Data for Presentation	218
Subsystem 9: Slowly Changing Dimension Manager	218
Standard Handling for Slowly Changing Dimensions	219
Custom Handling for Slowly Changing Dimensions	221
Alternatives to the Slowly Changing Dimension Transform	221
Subsystem 10: Surrogate Key Generator	223
Subsystem 11: Hierarchy Manager	223
Subsystem 12: Special Dimensions Manager	224
Subsystem 13: Fact Table Builders	225
Transaction Grain Fact Table Loader	226
Periodic Snapshot Fact Table Loader	228
Accumulating Snapshot Fact Table Loader	229
Subsystem 14: Surrogate Key Pipeline	229
Surrogate Key Pipeline Technique #1: Cascading Lookups	231
Surrogate Key Pipeline Technique #2: Database Joins	233
Which Technique Is Best?	234
Subsystem 15: Multi-Valued Dimension Bridge Table Builder	235
Subsystem 16: Late Arriving Data Handler	235
Late Arriving Dimension Members	235
Late Arriving Dimension Updates	236
Late Arriving Facts	236
Subsystem 17: Dimension Manager	238
Subsystem 18: Fact Provider System	238
Subsystem 19: Aggregate Builder	239
Subsystem 20: OLAP Cube Builder	239
Subsystem 21: Data Propagation Manager	240
Managing the ETL Environment	240
Summary	243
Chapter 8 The Core Analysis Services OLAP Database	245
Overview of Analysis Services OLAP	247
Why Use Analysis Services?	247
Why Not Analysis Services?	249
Designing the OLAP Structure	250
Planning	251
Cube Content	251
Cube Granularity	252
Cube Usage	252
Getting Started	253
Setup	253
Create Relational Views	253
Populate the Data Warehouse Database	254

Create a Project and a Data Source View	255
Dimension Designs	257
Standard Dimensions	258
Variable Depth or Parent-Child Hierarchies	259
Multivalued or Many-to-Many Dimensions	259
Role-Playing Dimensions	260
Creating and Editing Dimensions	261
Editing Dimension Properties	265
Editing Attribute Properties	266
Time and Account Dimensions	267
Creating Hierarchies	268
Set Up Attribute Relationships	270
Browsing Dimension Data	273
Creating and Editing the Cube	274
Edit the Cube Structure	276
Edit Dimension Usage	279
Build, Deploy, and Process the Project	281
Create Calculations	282
Define Key Performance Indicators	287
Create Actions	288
Partitions and Aggregations	289
Maintain Perspectives	290
Translations	290
Physical Design Considerations	291
Understanding Storage Modes	293
Developing the Partitioning Plan	294
Designing Performance Aggregations	296
Planning for Deployment	298
Processing the Full Cube	299
Developing the Incremental Processing Plan	299
Scheduled Processing	300
Planning for Updates to Dimensions	302
Planning for Fact Updates and Deletes	303
Summary	304
Chapter 9 Design Requirements for Real-Time BI	305
Real-Time Triage	306
What Does Real-Time Mean?	306
Who Needs Real Time?	307
Real-Time Tradeoffs	308
Scenarios and Solutions	311
Executing Reports in Real Time	313
Serving Reports from a Cache	313
Creating an ODS with Mirrors and Snapshots	314
Creating an ODS with Replication	314
Building a BizTalk Application	315

Building a Real-Time Relational Partition	315
Querying Real-Time Data in the Relational Database	317
Using Analysis Services to Query Real-Time Data	318
Summary	319
Part 3 Developing the BI Applications	321
Chapter 10 Building BI Applications in Reporting Services	323
A Brief Overview of BI Applications	324
Types of BI Applications	325
The Value of Business Intelligence Applications	326
A High-Level Architecture for Reporting	328
Reviewing Business Requirements for Reporting	328
Examining the Reporting Services Architecture	330
Using Reporting Services as a Standard Reporting Tool	332
Creating Reports	333
Finding Reports	334
Viewing Reports	336
Receiving Results	337
Changing Reports	337
Solid, Reliable System	339
Reporting Services Assessment	339
The Reporting System Design and Development Process	340
Reporting System Design	341
Determining the Initial Report Set	341
Creating the Report Template	343
Creating Report Specifications and Documentation	345
Designing the Navigation Framework	347
Conducting the User Review	348
Reporting System Development	348
Prepare	349
Build	349
Test	350
Deploy	350
Building and Delivering Reports	351
Planning and Preparation	351
Setting Up the Development Environment	351
Creating Standard Templates	352
Creating Reports	354
Revisit the Report Specifications	354
The Report Creation Process	355
Creating Your First Standard Report	356
Reporting Operations	368
Ad Hoc Reporting Options	369
The Report Model	370
Shared Datasets	371
Report Parts	371
Summary	372

Chapter 11 PowerPivot and Excel	375
Using Excel for Analysis and Reporting	376
The PowerPivot Architecture: Excel on Steroids	378
Creating and Using PowerPivot Databases	380
Getting Started	381
PowerPivot Table Design	381
Loading the Data	382
Creating the Relationships	383
Creating Analytics with PowerPivot	385
Creating a PowerPivot PivotTable	385
Adding New Measures to the PivotTable	386
The CALCULATE() Function	387
Adding a Computed Column to the PivotTable Database	390
Observations and Guidelines on PowerPivot for Excel	392
PowerPivot for SharePoint	394
The PowerPivot SharePoint User Experience	394
PowerPivot Publishing	394
PowerPivot Viewing	396
PowerPivot as a Data Source	396
Server-Level Resources	397
PowerPivot Monitoring and Management	397
PowerPivot Monitoring	398
PowerPivot Workbook Publishing Process	400
PowerPivot's Role in a Managed DW/BI Environment	400
Summary	401
 Chapter 12 The BI Portal and SharePoint	 403
The BI Portal	404
Planning the BI Portal	405
Impact on Design	406
Business Process Categories	407
Additional Functions	408
Building the BI Portal	409
Using SharePoint as the BI Portal	411
Architecture and Concepts	412
SharePoint's Three-Tier Architecture and Topology	412
SharePoint Terminology	414
Setting Up SharePoint	417
The Installation Process	417
Plan SharePoint for BI	419
Installing the Test System	421
Completing the BI Portal	424
The Added Functionality of the BIPortal Site Template	425
Exploring SharePoint	426
Summary	426

Chapter 13	Incorporating Data Mining	429
	Defining Data Mining	430
	Basic Data Mining Terminology	432
	Business Uses of Data Mining	433
	Classification	434
	Estimation (Regression)	434
	Prediction	435
	Association or Affinity Grouping	436
	Clustering (Segmentation)	436
	Anomaly Detection	438
	Description and Profiling	438
	Business Task Summary	438
	Roles and Responsibilities	440
	SQL Server Data Mining Architecture Overview	440
	The Data Mining Design Environment	442
	Build, Deploy, and Process	442
	Accessing the Mining Models	443
	Integration Services and Data Mining	443
	Additional Features	444
	Architecture Summary	445
	Microsoft Data Mining Algorithms	445
	Decision Trees	446
	Naïve Bayes	447
	Clustering	448
	Sequence Clustering	448
	Time Series	449
	Association	449
	Neural Network	449
	The Data Mining Process	450
	The Business Phase	451
	Identifying Business Opportunities	451
	Understanding the Data	452
	Describing the Data Mining Opportunity	452
	The Data Mining Phase	453
	Data Preparation	453
	Model Development	456
	Model Validation (Evaluation)	457
	The Operations Phase	460
	Implementation	460
	Assess Impact	461
	Maintain the Model	461
	Metadata	462

Data Mining Examples	463
Case Study: Categorizing Cities	463
Categorizing Cities: Business Opportunity	464
Categorizing Cities: Data Understanding	464
Categorizing Cities: Data Preparation	465
Categorizing Cities: Model Development	466
Categorizing Cities: Model Validation	470
Categorizing Cities: Implementation	470
Categorization Cities: Maintenance and Assessment	471
Case Study: Product Recommendations	472
Product Recommendations: The Business Phase	472
Product Recommendations: The Data Mining Phase	474
Product Recommendations: The Operations Phase	487
Summary	488
Part 4 Deploying and Managing the DW/BI System	491
Chapter 14 Designing and Implementing Security	493
Identifying the Security Manager	494
Securing the Hardware and Operating System	495
Securing the Operating System	495
Using Windows Integrated Security	496
Securing the Development Environment	497
Securing the Data	498
Providing Open Access for Internal Users	498
Itemizing Sensitive Data	500
Securing Various Types of Data Access	500
Securing the Components of the DW/BI System	502
Reporting Services Security	502
Administrative Roles for Reporting Services	502
User Roles for Reporting Services	504
Reporting Services in SharePoint Integrated Mode	504
Analysis Services Security	505
Administrative Roles for Analysis Services	505
User Roles for Analysis Services	506
Dynamic Security	513
PowerPivot Security	514
Relational DW Security	514
Administrative Roles for the Relational Database	515
User Roles for the Relational Database	515
Integration Services Security	520
Usage Monitoring	521
Summary	521

Chapter 15 Metadata Plan	523
Metadata Basics	524
The Purpose of Metadata	524
Metadata Categories	525
The Metadata Repository	526
Metadata Standards	526
SQL Server 2008 R2 Metadata	527
Cross-Tool Components	528
Relational Engine Metadata	532
Analysis Services	532
Integration Services	533
Reporting Services	533
Master Data Services	534
SharePoint	534
External Metadata Sources	534
System Monitor	534
Active Directory	535
Looking to the Future	535
A Practical Metadata Approach	535
Creating the Metadata Strategy	536
Business Metadata Reporting	538
Analysis Services as Primary Query Platform	538
Relational Engine Extended Properties	539
Business Metadata Schema	540
Process Metadata Reporting	541
Technical Metadata Reporting	542
Ongoing Metadata Management	543
Summary	543
Chapter 16 Deployment	545
Setting Up the Environments	546
Testing	550
Development Testing	551
System Testing	555
Data Quality Assurance Testing	557
Performance Testing	559
Service Level Confirmation	559
Processing Performance: Getting Data In	560
Query Performance: Getting Data Out	561
Usability Testing	562
Testing Summary	563
Deploying to Production	564
Relational Database Deployment	565
Integration Services Package Deployment	567
Analysis Services Database Deployment	568
Reporting Services Report Deployment	571
Master Data Services Deployment	572

Data Warehouse and BI Documentation	573
Core Descriptions	573
Business Process Dimensional Model Descriptions	573
Table and Column Descriptions	574
Report Descriptions	575
Additional Documentation	575
User Training	576
User Support	579
Desktop Readiness and Configuration	580
Summary	581
Chapter 17 Operations and Maintenance	583
Providing User Support	584
Maintaining the BI Portal	585
Extending the BI Applications	586
System Management	587
Governing the DW/BI System	588
Identifying and Terminating User Sessions	588
Resource Governance	591
Performance Monitoring	593
Relational Database	594
Analysis Services Database	597
Reporting Services	598
Integration Services	598
PowerPivot	600
Usage Monitoring	600
Managing Disk Space	602
Service and Availability Management	603
Performance Tuning the DW/BI System	604
Backup and Recovery	606
Relational Databases	606
Integration Services	608
Analysis Services	609
Reporting Services	610
Recovery	610
Executing the ETL Packages	611
Summary	611
Chapter 18 Present Imperatives and Future Outlook	613
Growing the DW/BI System	613
Lifecycle Review with Common Problems	615
Phase I — Requirements, Realities, Plans, and Designs	616
Phase II — Developing the Databases	616
Phase III — Developing the BI Applications and	
Portal Environment	617
Phase IV — Deploying and Managing the DW/BI System	618
Iteration and Growth	618

What We Like in the Microsoft BI Toolset	619
Future Directions: Room for Improvement	620
Query Tools	620
Metadata	621
Relational Database Engine	621
Analysis Services	622
Master Data Services	622
Integration	623
Customer Focus	623
Conclusion	623

Index	625
--------------	------------



Foreword

In the five years since the first edition was published, Microsoft has made impressive progress in building out its data warehousing and business intelligence tools suite. It is gratifying to those of us who work in this space to see the steady commitment that Microsoft has made to provide usable, professional quality tools. During these five years, Warren and Joy have consulted with dozens of clients, taught scores of classes, answered hundreds if not thousands of questions, had many “schema lunches” where the schema diagrams competed with the food, and have pounded on every module in Microsoft’s DW/BI toolset. This current edition remains a unique reference, combining overall perspectives on what the tools do with accurate assessments of how well they do it. This book teaches judgment, not button clicks!

—Ralph Kimball

