

MDX Solutions

Second Edition

**With Microsoft® SQL Server™ Analysis
Services 2005 and Hyperion® Essbase**

George Spofford
Sivakumar Harinath
Christopher Webb
Dylan Hai Huang
Francesco Civardi



WILEY

Wiley Publishing, Inc.

MDX Solutions

Second Edition

**With Microsoft® SQL Server™ Analysis
Services 2005 and Hyperion® Essbase**

George Spofford
Sivakumar Harinath
Christopher Webb
Dylan Hai Huang
Francesco Civardi



WILEY

Wiley Publishing, Inc.

MDX Solutions, Second Edition: With Microsoft® SQL Server™ Analysis Services 2005 and Hyperion® Essbase

Published by

Wiley Publishing, Inc.

10475 Crosspoint Boulevard

Indianapolis, IN 46256

www.wiley.com

Copyright © 2006 by Wiley Publishing, Inc., Indianapolis, Indiana

Published simultaneously in Canada

ISBN-13: 978-0-471-74808-3

ISBN-10: 0-471-74808-0

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

2MA/RX/QS/QW/IN

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 Legal Department, Wiley Publishing, Inc., 10475 Crosspoint Blvd., Indianapolis, IN 46256, (317) 572-3447, fax (317) 572-4355, 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 Website 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 or to obtain technical support, please contact our Customer Care Department within the U.S. at (800) 762-2974, outside the U.S. at (317) 572-3993 or fax (317) 572-4002.

Library of Congress Catalog Number: 2005032778

Trademarks: Wiley and the Wiley logo are 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 and SQL Server are trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. Hyperion is a registered trademark of Hyperion Solutions 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.

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

To my wife, Lisa, and to my parents, for their care and love with their children

–George Spofford

I dedicate this book in the grandest possible manner to my dear wife Shreepriya, who has been fully supportive and put up with me disappearing from home or spent several late nights when I worked on this book. It is also dedicated to my two-month-old twins Praveen and Divya, who do not know yet what a book is. I wish their cute photographs could be on the cover page. Finally, I would like to dedicate this book to the memory of my father Harinath Govindarajalu who passed away in 1999 who I am sure would have been proud of this great achievement and to my mother Sundara Bai.

–Sivakumar Harinath

For Helen and Natasha

–Chris Webb

To Alice, for your support and love.

–Dylan Hai Huang

I dedicate this book to my wife Teresa and my daughter Chiara, my main supporters in every challenge in my life. It is also dedicated to my cat Picci and its new friends Cartesio and Ipazia Hoppipolla for their particular attentions and their company during the writing of this book.

–Francesco Civardi



About the Authors

George Spofford is a Distinguished Engineer at Hyperion Solutions, Inc. He has been developing OLAP server and client software since 1998, when he led the development of FreeThink, one of the first desktop multidimensional modelling and analysis tools provided by Power Thinking tools, Inc. Upon its acquisition by Computer Corporation of America, he became product architect for their integrated OLAP/data warehouse server. Subsequently, he co-founded the consulting firms Dimensional systems and DSS Lab, where he provided technology consulting, benchmark development and auditing, and development services to vendors and users alike of decision-support technology, including organizations such as AC Nielsen, IMS Health, Intel, Oracle, and Microsoft. He has written numerous articles for trade journals, spoken at trade shows, co-authored the book *Microsoft OLAP Solutions*, and authored the first edition of *MDX Solutions*. george@dsslabs.com.

Sivakumar Harinath was born in Chennai, India. Siva has a Ph.D. in Computer Science from the University of Illinois at Chicago. His thesis title was: "Data Management Support for Distributed Data Mining of Large Datasets over High Speed Wide Area Networks." Siva has worked for Newgen Software Technologies (P) Ltd., IBM Toronto Labs, Canada, and has been at Microsoft since February of 2002. Siva started as a Software Design Engineer in Test (SDET) in the Analysis Services Performance Team and is currently an SDET Lead for Analysis Services 2005. Siva's other interests include high performance computing, distributed systems and high-speed networking. Siva is married to Shreepriya and had twins Praveen and Divya during the course of writing this book. His personal interests include travel, games/sports (in particular, Chess, Carrom, Racquet Ball, Board games) and Cooking. You can reach Siva at sivakumar.harinath@microsoft.com.

Christopher Webb (cwebb_olap@hotmail.com) has worked with Microsoft Analysis/OLAP Services since 1998 in a variety of roles, including three years spent with Microsoft Consulting Services. He is a regular contributor to the microsoft.public.sqlserver.olap newsgroup and his blog can be found at <http://spaces.msn.com/members/cwebbbi/>.

Dylan Hai Huang (Dylanh@microsoft.com) is currently working as a program manager in Microsoft Office Business Application Team. Before joining the current team, he had been working in Microsoft Analysis Service Performance Team. His main interests are business intelligence, high performance computation and data mining.

Francesco Civardi (fcivardi@daisylabs.com, fcivardi@hotmail.com) is Chief Scientist at DaisyLabs, a Business Intelligence Company. He has been working with Microsoft Analysis/OLAP Services since 1999, in the financial, industrial and retail sectors. His main interests are mathematical and financial modelling and Data Mining. He is Professor of Data Analysis Tools and Techniques at the Catholic University of Brescia and Cremona.



Executive Editor

Robert M. Elliott

Development Editor

Ed Connor

Technical Editor

Deepak Puri

Production Editor

Angela Smith

Copy Editor

Foxxe Editorial Services

Editorial Manager

Mary Beth Wakefield

Production Manager

Tim Tate

**Vice President and
Executive Group Publisher**

Richard Swadley

**Vice President and Executive
Publisher**

Joseph B. Wikert

Project Coordinator

Ryan Steffen

Graphics and Production Specialists

Denny Hager

Joyce Haughey

Jennifer Heleine

Stephanie D. Jumper

Alicia South

Quality Control Technicians

John Greenough

Charles Spencer

Brian Walls

Proofreading

TECHBOOKS Production Services

Indexing

Sherry Massey



Contents

Acknowledgments	xxi
Introduction	xxiii
Chapter 1 A First Introduction to MDX	1
What Is MDX?	1
Query Basics	2
Axis Framework: Names and Numbering	5
Case Sensitivity and Layout	6
Simple MDX Construction	7
Comma (,) and Colon (:)	7
.Members	9
Getting the Children of a Member with .Children	10
Getting the Descendants of a Member with Descendants()	11
Removing Empty Slices from Query Results	14
Comments in MDX	16
The MDX Data Model: Tuples and Sets	17
Tuples	18
Sets	20
Queries	21
Queries with Zero Axes	22
Axis-Only Queries	23
More Basic Vocabulary	23
CrossJoin()	23
Filter()	25
Order()	28
Querying for Member Properties	30
Querying Cell Properties	32
Client Result Data Layout	34
Summary	35

Chapter 2	Introduction to MDX Calculated Members and Named Sets	37
	Dimensional Calculations As Calculated Members	38
	Calculated Member Scopes	39
	Calculated Members and WITH Sections in Queries	39
	Formula Precedence (Solve Order)	42
	Basic Calculation Functions	48
	Arithmetic Operators	48
	Summary Statistical Operators	49
	Avg()	50
	Count(), .Count	50
	DistinctCount() (Microsoft extension)	51
	Sum()	52
	Max()	52
	Median()	52
	Min()	53
	NonEmptyCount() (Hyperion extension)	53
	Stdev(), Stddev()	54
	StdevP(), StddevP() (Microsoft Extension)	54
	Var(), Variance()	54
	VarP(), VarianceP() (Microsoft Extension)	55
	Additional Functions	55
	Introduction to Named Sets	57
	Named Set Scopes	58
	Summary	60
Chapter 3	Common Calculations and Selections in MDX	61
	Metadata Referencing Functions in MDX	64
	Many Kinds of Ratios, Averages, Percentages, and Allocations	65
	Percent Contribution (Simple Ratios between Levels in a Hierarchy)	65
	Percent Contribution to Total	66
	Using the .CurrentMember function	66
	Using the .Parent function	66
	Taking the Share-of-Parent Using .CurrentMember and .Parent	67
	Using the Ancestor() function	67
	Calculating the Share-of-Ancestor using .CurrentMember and Ancestor()	67
	Handling Division by Zero	69
	Basic Allocations	70
	Proportional Allocation of One Quantity Based on Ratios of Another	70
	Unweighted Allocations down the Hierarchy	71
	Averages	71
	Simple Averages	72
	Weighted Averages	73
	Time-Based References and Time-Series Calculations	74
	Period-to-Period References and Calculations	75
	Same-Period-Last-Year References and Calculations	76

Year-to-Date (Period-to-Date) Aggregations	76
Rolling Averages and 52-week High/Low	79
Using LastPeriods() to Select Time Ranges Based on a Target Member	81
Different Aggregations along Different Dimensions (Semi-Additive Measures Using MDX)	82
Mixing Aggregations: Sum across Non-Time, Average/Min/Max along Time	82
Mixing Aggregations: Sum across Non-time, Opening/Closing Balance along Time	83
Carryover of Balances for Slowly Changing Values and Reporting of Last Entered Balance	84
Finding the Last Child/Descendant with Data	87
Finding the Last Time Member for Which Any Data Has Been Entered	88
Using Member Properties in MDX Expressions (Calculations and Sorting)	89
Handling Boundary Conditions (Members out of Range, Division by Zero, and More)	92
Handling Insufficient Range Size	93
Handling Insufficient Hierarchical Depth	93
Handling a Wrong-Level Reference	94
Handling Division by Zero	95
Summary	95
Chapter 4 MDX Query Context and Execution	97
Cell Context and Resolution Order in Queries	98
The Execution Stages of a Query	99
The .DefaultMember Function	100
Default Context and Slicers	101
The Simplest Query: All Context, Nothing Else	101
The WHERE Clause: Default Context and Slicers	102
Adding Axes to a Query	103
Cell Context When Resolving Axes	104
Overriding Slicer Context	106
Cell Evaluation (For Any Cell)	107
Drilling in on Solve Order and Recursive Evaluation	108
Resolving NON EMPTY Axes	110
Resolving the HAVING Clause in AS2005	111
Looping Context and .CurrentMember	114
Interdependence of Members in AS2005: Strong Hierarchies, Autoexists, and Attribute Relationships	116
Strong Hierarchies	116
Autoexists	118
Modifying the Cube Context in AS2005	119
CREATE SUBCUBE Described	120
Subcube Restrictions and Attribute Relations	123
Further Details of Specifying a Subcube	125
Tuple Specifications for Subcubes	125

	Subcubes Based on Data Values	127
	Subcubes for Iterative Query Refinement	127
	Points to Consider When Using Subcubes	128
	Using SELECT in the FROM Clause in AS2005	128
	Infinite Recursion: A “Gotcha” Related to Calculation Context	131
	Product-Specific Solve Order Use	132
	Use of Solve Order between Global, Session, and Query Calculations in Analysis Services 2005	132
	Use of Solve Orders in Essbase	134
	Use of Solve Orders in Analysis Services 2000	135
	Nondata: Invalid Numbers, NULLs, and Invalid Members	135
	Invalid Calculations: Division by Zero and Numerical Errors	136
	Semantics of Empty Cells	136
	NULLs in Comparisons and Calculations	138
	Invalid Locations	140
	Precedence of Cell Properties in Calculations	143
	Precedence of Display Formatting	143
	Data Types from Calculated Cells	144
	Cube Context in Actions	146
	Cube Context in KPIs	146
	Visibility of Definitions between Global, Session, and Query-Specific Calculations in Analysis Services 2005	146
	Summary	148
Chapter 5	Named Sets and Set Aliases	149
	Named Sets: Scopes and Context	149
	Common Uses for Named Sets	150
	Set Aliases	152
	An Example of a Set Alias	153
	Set Aliases in More Detail	155
	When Set Aliases Are Required	157
	Summary	160
Chapter 6	Sorting and Ranking in MDX	161
	The Function Building Blocks	161
	Classic Top- <i>N</i> Selections	162
	Adding Ranking Numbers (Using the Rank() function)	165
	Handling Tied Ranks: Analysis Services	168
	Taking the Top- <i>N</i> Descendants or Other Related Members across a Set	169
	Getting the Fewest/Most Tuples to Reach a Threshold	172
	Retrieving the Top <i>N</i> Percent of Tuples	174
	Retrieving the Top <i>N</i> Percent of the Top <i>N</i> Percent	174
	Putting Members/Tuples in Dimension Order (Ancestors First or Last)	175
	Reversing a Set	176
	Summary	177

Chapter 7	Advanced MDX Application Topics	179
	Arranging Parents/Ancestors after Children, Not Before	181
	Returning the Subtree under a Member and the Ancestors of That Member Along with the Member	181
	Using Generate() to Turn Tuple Operations into Set Operations	182
	Calculating Dates/Date Arithmetic	183
	Defining Ratios against the Members Selected on Rows/Columns/Axes, Instead of against a Specific Dimension	187
	Report-Based Totals-to-Parent, Percentage Contribution to Report Totals	190
	Technique 1: Only Standard MDX Techniques	191
	Technique 2: Considering Using VisualTotals() in Analysis Services	197
	Using VisualTotals in Analysis Services 2000	197
	Using VisualTotals in AS2005	198
	Technique 3: Using AS2005 Subcubes	199
	Hierarchical Sorting That Skips Levels in the Hierarchy	200
	Sorting a Single Set on Multiple Criteria	202
	Multiple Layers or Dimensions of Sorting	202
	Sort Nested Dimensions with the Same Sorting Criterion for Each Dimension	203
	Sort Nested Dimensions by Different Criteria	204
	Pareto Analysis and Cumulative Sums	207
	Returning the Top-Selling Product (or Top-Selling Month or Other Most-Significant Name) As a Measure	211
	Most Recent Event for a Set of Selected Members	212
	How Long Did It Take to Accumulate This Many ? (Building a Set That Sums Backward or Forward in Time)	216
	Aggregating by Multiplication (Product Instead of Sum)	219
	One Member Formula Calculating Different Things in Different Places	220
	Including All Tuples with Tied Ranking in Sets	225
	Time Analysis Utility Dimensions	227
	A Sample Analysis	229
	Summary	237
Chapter 8	Using the Attribute Data Model of Microsoft Analysis Services	239
	The Unified Dimensional Model (UDM)	240
	Dimensions	242
	Attributes, Hierarchies, and Relationships	244
	Attributes	245
	Hierarchies and Levels	247
	Relationships	249
	Querying Dimensions	249
	Member Properties	252
	Parent-Child Hierarchies	254
	Time Dimension	257

Cubes	257
Dimension Relationships	260
Role-Playing Dimensions	264
Perspectives	265
Drill-Through	266
The Calculation Model in UDM	266
Defining Security on UDM	267
Summary	272
Chapter 9 Using Attribute Dimensions and Member Properties in Hyperion Essbase	273
UDAs and Attributes	273
Retrieving UDAs and Attribute Values on Query Axes	274
Predefined Attributes	275
Using UDA and Attribute Values in Calculations	275
Selecting Base Dimension Members Based on UDA and Attribute Values	276
Using Attribute() to Select Members Based on Shared Attribute Values	276
Using WithAttr() to Select Members Based on Attribute Values	278
Using UDA() to Select Members Sharing a UDA Value	279
Connecting Base Members to the Attribute Hierarchy with IN	280
Connecting Base Members to Their Actual Attribute Member	280
Connecting Attribute Members to Their Attribute Values	281
Summary	281
Chapter 10 Extending MDX through External Functions	283
Using Stored Procedures with MDX	285
.NET Stored Procedures	286
.NET Stored Procedure Parameters and Return Values	287
ADOMD Server objects	289
Expression	291
TupleBuilder	291
SetBuilder	292
MDX	292
Context	293
Server Metadata Objects	294
AMO .NET Management Stored Procedures	295
Performance Considerations of Static Functions and Nonstatic Functions	297
Debugging .NET Stored Procedures	299
Additional Programming Aspects NULL, ERROR(), and Exception	300
NULL Value As an Input Parameter	300
NULL Value As an Output Parameter	301
Exceptions during Execution	301
Error() Function	302
Using Stored Procedures for Dynamic Security	303
COM DLL Stored Procedures	305

Argument and Return-Type Details	306
Passing Arrays of Values to COM Stored Procedures	307
MDX Functions for Use with COM Stored Procedures	312
SetToStr(), TupleToStr()	312
Members(), StrToSet(), StrToTuple()	313
External Function Example: Time Span until Sum	315
Loading and Using Stored Procedures	316
Security of Stored Procedures	317
Stored Procedure Name Resolution	318
Invoke Stored Procedures in MDX	319
Additional Considerations for Stored Procedures	320
Summary	321
Chapter 11 Changing the Cube and Dimension Environment through MDX	323
Altering the Default Member for a Dimension in Your Session	324
Dimension Writeback Operations	325
Creating a New Member	325
Moving a Member within Its Dimension	326
Dropping a Member	327
Updating a Member's Definition	327
Refresh Cell Data and Dimension Members	328
Writing Data Back to the Cube	329
Standard Cell Writeback	329
Commit and Rollback	330
Using UPDATE CUBE	330
Summary	334
Chapter 12 The Many Ways to Calculate in Microsoft Analysis Services	335
Overview of Calculation Mechanisms	336
Intrinsic Aggregation for a Measure	336
Rollup by Unary Operator	338
Custom Member Formula	339
Calculated Member	341
Defining a Calculated Member	342
Dropping a Calculated Member	345
Cell Calculation	346
Defining a Cell Calculation	346
Dropping a Cell Calculation	350
Conditional Formatting	351
How Types of Calculations Interact	351
Interaction without Any Cell Calculations	352
Precedence of Custom Member Formulas on Multiple Dimensions	352
Precedence of Unary Operators on Multiple Dimensions	352
Cell Calculation Passes	353
Equation Solving and Financial Modeling	356
Using Solve Order to Determine the Formula in a Pass	358

Calculated Members Not Themselves Aggregated	360
Intrinsic Aggregation of Custom Rollups, Custom Members, and Calculated Cell Results	360
Tips on Using the Different Calculation Techniques	362
Summary	362
Chapter 13 MDX Scripting in Analysis Services 2005	365
MDX Scripting Basics	366
What Is an MDX Script?	366
The Calculate Statement	367
Subcubes	368
Assignments and Aggregation	371
Assignments and Calculated Members	376
Assignments and Named Sets	377
MDX Scripting and More Complex Cubes	379
Multiple Attribute Hierarchies	379
User Hierarchies	386
Parent/Child Attribute Hierarchies	387
Many-to-Many Dimensions	388
Fact Dimensions and Reference Dimensions	390
Semi-additive and Nonadditive Measures	390
Unary Operators and Custom Member Formulas	393
Advanced MDX Scripting	395
Defining Subcubes with SCOPE	395
Assignments That Are MDX Expressions	398
Assigning Error Values to Subcubes	402
Assigning Cell Property Values to Subcubes	402
Conditional Assignments	404
Real-World MDX Scripts	405
The Time Intelligence Wizard	405
Basic Allocations Revisited	408
Summary	410
Chapter 14 Enriching the Client Interaction	411
Using Drill-Through	412
Improvements and Changes in Microsoft Analysis Services 2005 for Drill-Through	413
MDX for Drill-Through I	413
Programmatic Aspects of Drill-Through	415
MDX for Drill-Through II	417
Drill-Through Security	418
Using Actions	419
What Can You Do with an Action?	419
Targets for Actions	424
Defining an Action	425
Programmatic Aspects of Actions	428
Dropping an Action	432

Using KPIs	432
Creating KPI	433
MDX KPI Function	436
Using KPI	437
Summary	439
Chapter 15 Client Programming Basics	441
ADOMD.NET Basics	442
Prerequisites	443
Making a Connection	443
Working with Metadata	444
Retrieving Schema Rowsets	445
Interoperability Considerations When Using Schema Rowsets	446
Working with the Metadata Object Model	446
Interoperability Considerations When Working with the Metadata Object Model	447
Dimension Particularities	448
Handling ADOMD.NET Metadata Caching	449
Executing a Query	450
Executing Commands	450
Parameterized Commands	451
Working with the CellSet Object	452
OlapInfo Holds Metadata	452
Axes Hold Axis Information	454
Cells Hold Cell Information	455
Further Details on Retrieving Information from a Query	457
Retrieving Member Property Information	457
Retrieving Additional Member Information	460
Further Details about Retrieving Cell Data	460
Retrieving Drill-Through Data As a Recordset	462
Key Performance Indicators	463
Executing Actions	464
Handling “Flattened” MDX Results	466
DataReader and Tabular Results for MDX Queries	466
Axis 0	467
Other Axes	468
Summary	470
Chapter 16 Optimizing MDX	471
Architecture Change from Analysis Services 2000 to 2005	472
Optimizing Set Operations	473
Sums along Cross-Joined Sets	474
Filtering across Cross-Joined Sets	475
Optimizing TopCount() and BottomCount()	477
NonEmpty function In Analysis Services 2005	478
Optimizing Sorting: Order()	480
UnOrder Function for a Query with a Large Dataset	480

Optimizing Summation	480
Designing Calculations into Your Database (Putting Member Properties into Measures and the New MDX function MemberValue)	482
MDX Script Optimization	483
Scope the Calculation in Detail	484
Avoid Leaf-Level Calculations	485
Cube Design to Avoid Leaf-Level Calculation	486
Measure Expression to Optimize Leaf-Level Calculation	487
MDX Script Optimization for Leaf-Level Calculation	488
Avoid Unnecessary Leaf-Level Calculation	489
Using NONEMPTY for Higher-Level Calculations	490
Using NonemptyBehavior to Provide a Hint for Server Calculations	491
Analysis Service 2005: Use Attribute Hierarchy Instead of Member Property	491
Analysis Service 2005: Use Scope Instead of IIF	492
Avoid Slow Functions in MDX Scripts	495
Change the Calculation Logic for Better Performance: Flow Calculation	495
Use Server Native Features Rather Than Scripts for Aggregation-Related Calculations	497
Summary	498
Chapter 17 Working with Local Cubes	501
Choosing Which Syntax to Use	502
Using the CREATE CUBE Statement	502
Overview of the Process	502
Anatomy of the CREATE CUBE Statement	503
Defining Dimensions	504
Overall Dimension	504
Named Hierarchies	505
Levels	505
Member Properties	508
FORMAT_NAME and FORMAT_KEY	510
Defining Measures	511
Adding Commands	512
ROLAP versus MOLAP	513
Anatomy of the INSERT INTO Statement	514
Cube Targets	515
Regular Dimension Levels	515
Parent-Child Dimensions	516
Member Properties	516
Custom Rollups	517
Measures	517
Column Placeholders in the Targets	517
Options for the INSERT INTO	517

The SELECT Clause	518
Select Statements That Are Not SQL	520
More Advanced Programming: Using Rowsets in Memory	520
Tips for Construction	521
Local Cubes from Server Cubes	521
Rollups and Custom Member Formulas	522
Using the CREATE GLOBAL CUBE Statement	524
Overview of the Process	524
Anatomy of the CREATE GLOBAL CUBE Statement	525
Defining Measures	525
Defining Dimensions	525
Defining Levels	526
Defining Members for Slicing	527
Things to Look Out For	527
Using Analysis Services Scripting Language	528
Overview of the Process	528
Anatomy of an ASSL Statement	528
Security	529
Summary	530
Appendix A MDX Function and Operator Reference	531
Appendix B Connection Parameters That Affect MDX	637
Appendix C Intrinsic Cell and Member Properties	661
Appendix D Format String Codes	675
Index	685

xxii Acknowledgments

Amynmohamed Rahan, Giuseppe Romanini, Marco Russo, Carlo Emilio Salaroglio, Lisa Santoro, Vaishnavi Sashikanth, Rick Sawa, John Shiangoli, Eric Smadja, Paolo Tenti, Alaric Thomas, Erik Thomsen, Richard Tkatchuk, Fabrizio Tosca, Piero Venturi, Massimo Villa, Mark Yang, Jeffrey Wang, Barry Ward, Robert Zare, and Tonglian Zheng. Ed Connor helped form this work into a coherent book, and Robert Elliott kept it on track.



Acknowledgments

This book would not be possible without the assistance of a large number of people providing several kinds of support, especially editing of prose, clarifying explanations, verifying the content of the material, hoisting of slack in workloads, and tolerating the long and persistent retreats into the home offices that enabled the authors to get the work done. George Chow and Richard Mao at Simba Technologies get special recognition for authoring chapter 15 (Client Programming Basics). Deepak Puri provided technical review during book development. The authors would like to thank the following individuals or organizations for their assistance during this effort:

Matt Abrams, Steve Allen, Luca Allieri, T.K. Anand, Teresa Andena, Dario Appini, Jon Axon, Graeme Axworthy, Nenshad Bardoliwalla, Marin Bezic, Veronica Biagi, Peter Bull, Matt Carroll, the Computer Café, Marcella Caviola, Saugata Chowdhury, Chiara Civardi, Claudio Civetta, Valerio Caloni, Marina Costa, Andrea Cozzoli, Thierry D’Hers, Marius Dumitru, Darryl Eckstein, Leah Etienne, Willfried Faerber, Hanying Feng, Curt Fish, Lon Fisher, Carlo Gerelli, Robert Gersten, Stefania Ghiretti, Alessandro Giancane, Sergei Gringauze, Colin Hardie, Alice Hu, Alessandro Huber, Graham Hunter, Andrea Improta, Gonsalvo Isaza, Thomas Ivarsson, Eric Jacobsen, Bruce Johnson, Yan Li, Ben van Lierop, Domiziana Luddi, John Lynn, Dermot MacCarthy, Aldo Manzetti, Al Marciante, Stefan Marin, Teodoro Marinucci, Frank McGuff, Krishna Meduri, Sreekumar Menon, Akshai Mirchandani, Michael Nader, Amir Netz, Ariel Netz, the Office Business Application team, Mimi Olivier, Marco Olivieri, Paul Olsen, Domenico Orofino, Massimo Pasqualetti, Mosha Pasumansky, Franco Perduca, Tim Peterson, Ernesto Pezzulla, Simona Poggi, John Poole,



Introduction

Dimensional applications are best and most easily built using a dimensional language. These dimensional applications are typified by the related notions of OLAP and dimensional data warehouses and marts. MDX (for *MultiDimensional eXpressions*) is the most widely accepted software language used for these applications. The book you read here, *MDX Solutions*, is the second edition of a guide to learning and using MDX. Since the first edition of *MDX Solutions*, the number of analytical applications that use MDX has grown very large, and several more servers and many third-party and homegrown client tools now allow you to use MDX to express the logic you use to calculate and to retrieve your information.

As a language, MDX is rather different in style and feel from SQL, and quite different from other programming languages like C++, C#, Lisp, Fortran, and so on. You can think of the formula language of a spreadsheet like Excel as another programming language, and while it is different than Excel, characterizing MDX as a sort of Excel-like SQL or a SQL-like Excel seems more apt than any other analogy. (If you're familiar with other OLAP query or calculation languages, it has more in common with them, but many readers will not be.)

In particular, compared to the first edition, this edition incorporates both a new version of a product and a product that has added support for MDX since then. Microsoft has released Microsoft® SQL Server 2005™ Analysis Services, which changes its use of MDX from the product's prior version. Hyperion Solutions has released Hyperion® System™ 9 BI+™ Analytic Services™, which builds on the Essbase functionality that introduced the term OLAP to the industry. (Since this book will refer to these products many times, we will refer to them with names that are informally shortened from the vendor's designations: *Analysis Services 2005*, *Analysis Services 2000*, and *Essbase 9*.)

Overview of the Book and Technology

The dimensional language works with a (multi-) dimensional data model. There are no formal or detailed standards of data model in the OLAP industry, and the number of details that would need to be worked out is quite large. However, there are a good number of aspects that are common by convention to different products. MDX has a standard syntax that handles the constructs and capabilities of many servers quite well. Vendors with additional capabilities extend MDX to provide access to them.

MDX originated as part of Microsoft's OLE DB for OLAP specification, and while the language was controlled by Microsoft, it sought input from several OLAP vendors and industry participants to help ensure that the language would be useful to multiple vendors. Microsoft has committed to transfer control of the specification to the XMLA Council (<http://www.xmla.org>), a consortium committed to coordinating and promoting a standard XML for Analysis specification. XML for Analysis is a web services API, spearheaded and supported by Microsoft, Hyperion and SAS Institute.

This book attempts to guide a pragmatic course between learning the language, and learning how to use it for the three product versions that we emphasize:

- Microsoft Analysis Services 2005
- Essbase 9
- Microsoft Analysis Services 2000

Microsoft has made a number of changes to its underlying data model with the 2005 release, and also a number of changes to how things built in MDX work with each other. This book devotes significant attention to using both of these kinds of changes. The Hyperion Essbase model has had a number of additional capabilities as well.

Other vendors or projects whose server products support MDX include: Applix, Microstrategy, MIS AG, Mondrian, SAP, and SAS Institute. Other companies, like Simba and Digital Aspects, provide tools and SDKs to assist construction of servers and clients that use MDX and related APIs. A large number of client tools use MDX to provide end users the ability to access sophisticated applications using MDX.

In order to fully wield MDX, you will need to really master how the server supports OLAP and how MDX behaves, since the two sides impact each other. Awesome MDX can make up for or extend the capabilities of a dimensional design; an awesome server design can eliminate the need to devote much MDX to solving the problem. This being a book on MDX, we won't spend much time

telling you how not to use MDX, but we will point out some cases where you're better off using some other aspect of your application environment.

How This Book Is Organized

If you are new to MDX, welcome! The chapters are sequenced to introduce you to the syntax, capabilities and use of MDX. Chapters 1, 2 and 3 introduce the basics of MDX and use. Chapter 4 discusses the logic of actually executing MDX in some depth, and although it is sort of a large chapter, there is a lot to know across the three product versions to actually understand in some detail how your MDX actually works. Chapters 5, 6, and 7 then build on using the details of how it works. Chapters 8 and on begin to cover product-specific aspects. So the first half of the book forms a background in capabilities and technique, and the second half introduces a range of possibilities specific to Microsoft and Hyperion.

Appendix A is a reference to all of the functions and operators of standard MDX and the extensions provided by the product versions covered. These functions are a big part of the language, and combining them in the right ways allows you to solve many different problems. They are a big part of the vocabulary for MDX, so the more of them you familiarize yourself with, the better. If the reference weren't so large, it would probably be a chapter in between Chapters 3 and 4, so that after learning the basics of MDX use you would immediately progress to the range of functions available. In several places within the chapters, you are specifically encouraged to look in Appendix A, and here in the introduction we will also encourage you to go through it. We won't be able to cover the use of every single function within the chapters otherwise, so you will pick up some useful nuggets by reviewing it on its own.

You will pick up a lot of technique and a number of tricks from the first seven chapters. These techniques will all be amplified by the contents of later chapters, or will be applicable to them. Unfortunately, we won't be able to give you every solution to every problem that you will have, but we hope to give you all you need to know about putting it together for yourself.

MDX is suited for a large number of applications. In order to simplify examples and explanations, this book includes only a few and really focuses on two. The Waremart 2005 database provides a common frame of reference between Hyperion and Microsoft products across many features and techniques, although it provides only a simple cube to work with. Chapters that deal largely with Analysis Services 2005 capabilities (including Chapters 8, 10, 13, and 14) will also refer to the Adventure Works database that ships with that product. Chapter 13 also includes a series of simple but incrementally more sophisticated databases.

What's Not in This Book

This book does not really cover the non-MDX aspects of building an analytical application. For developers targeting the Microsoft tool set, you may wish to refer to *Professional SQL Server Analysis Services 2005 with MDX* by Sivakumar Harinath, and Stephen R. Quinn, or *The Microsoft Data Warehouse Toolkit : With SQL Server 2005 and the Microsoft Business Intelligence Toolset* by Joy Mundy and Warren Thornthwaite.

Who Should Read This Book

This book is for any developer, consultant, or manager that needs to build and maintain a proficiency in MDX. MDX, dealing with calculations and selections, can be used for a large part of an overall application, so your concern may be as a front-end developer tasked with matching gestures or other specifications into MDX that retrieves a modified query. You may be an ASP or JSP developer, or developing reports in SQL Server's Reporting Services, and need to be able to translate simple and complex report requests. You may also be developing server-side calculation and modeling logic or security filters. Each of these may boil down to MDX that you write or MDX that you form in some helpful GUI. While the GUI may completely address your concerns, the code it generates will interact with other definitions and you are much better off having a good idea of what is going on and what the limits are.

Tools You Will Need

In order to run queries from the book, you will need a front end or API that you can send MDX in through and receive results back. Microsoft SQL Server 2005 ships with the SQL Server Management Studio, which can be used to run MDX queries. If you have upgraded from Analysis Services 2000, you may still use the MDX Sample application from that product edition against Analysis Services 2005. Hyperion Essbase includes the Analytic Administration Services console and the Essmsh command shell. Other tools are available as well; see the following on the related Web Site.

What's on the Web Site

The web site contains a collection of sample databases and code for use in Analysis Services and Essbase, and an MDX query interface for Essbase.

Summary

MDX is its own language, similar in some ways to languages you are familiar with and different in other ways. Even if you are familiar with OLAP concepts from one of the covered products, and especially if you are new to them, you may find it helpful to approach the language on its own terms. We don't assume that you know any of the language at the beginning, but if you do then we hope you will learn useful details that you didn't know at the outset. And we hope that you enjoy both learning and using it.

