



Beginning Microsoft[®] SQL Server[®] 2012 Programming

Paul Atkinson, Robert Vieira

BEGINNING MICROSOFT® SQL SERVER® 2012 PROGRAMMING

INTRODUCTIO	N xxv
CHAPTER 1	RDBMS Basics: What Makes Up a SQL Server Database?1
CHAPTER 2	Learning the Tools of the Trade
CHAPTER 3	The Foundation Statements of T-SQL
CHAPTER 4	JOINs
CHAPTER 5	Creating and Altering Tables 129
CHAPTER 6	Keys and Constraints 175
CHAPTER 7	Adding More to Your Queries 213
CHAPTER 8	Being Normal: Normalization and Other Basic Design Issues257
CHAPTER 9	SQL Server Storage and Index Structures
CHAPTER 10	Views
CHAPTER 11	Writing Scripts and Batches
CHAPTER 12	Stored Procedures
CHAPTER 13	User-Defined Functions
CHAPTER 14	Transactions and Locks
CHAPTER 15	Triggers
CHAPTER 16	A Brief XML Primer
CHAPTER 17	Business Intelligence Fundamentals
CHAPTER 18	BI Storage and Reporting Resources
CHAPTER 19	Reporting for Duty, Sir! A Look at Reporting Services
CHAPTER 20	Getting Integrated with Integration Services
CHAPTER 21	Playing Administrator
APPENDIX	Answers to Exercises
INDEX	
► ONLINE A	APPENDICES
BONUS APPE	NDIX 1: System Functions1
BONUS APPE	NDIX 2: Very Simple Connectivity Examples

BEGINNING

Microsoft[®] SQL Server[®] 2012 Programming

BEGINNING Microsoft[®] SQL Server[®] 2012 Programming

Paul Atkinson Robert Vieira



John Wiley & Sons, Inc.

Beginning Microsoft® SQL Server® 2012 Programming

Published by John Wiley & Sons, Inc. 10475 Crosspoint Boulevard Indianapolis, IN 46256 www.wiley.com

Copyright © 2012 by John Wiley & Sons, Inc., Indianapolis, Indiana

Published simultaneously in Canada

ISBN: 978-1-118-10228-2 ISBN: 978-1-118-22388-8 (ebk) ISBN: 978-1-118-23621-5 (ebk) ISBN: 978-1-118-26216-0 (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 Web site may provide or recommendations it may make. Further, readers should be aware that Internet Web sites 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 publishes in a variety of print and electronic formats and by print-on-demand. Some material included with standard print versions of this book may not be included in e-books or in print-on-demand. If this book refers to media such as a CD or DVD that is not included in the version you purchased, you may download this material at http://booksupport.wiley.com. For more information about Wiley products, visit www.wiley.com.

Library of Congress Control Number: 2012933622

Trademarks: Wiley, the Wiley logo, Wrox, the Wrox logo, Wrox Programmer to Programmer, and related trade dress 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 and SQL Server are registered trademarks of Microsoft Corporation. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc., is not associated with any product or vendor mentioned in this book.

This book is dedicated to my parents, and to the unique perspective they provide from the union of the often professorial Atkinsons and the utterly hilarious Walshes. If any fraction of their wit has leaked through into its pages, this book will have succeeded.

CREDITS

EXECUTIVE EDITOR Robert Elliott

PROJECT EDITOR Maureen Spears

TECHNICAL EDITOR Eric Aldinger

PRODUCTION EDITOR Daniel Scribner

COPY EDITOR Gayle Johnson

EDITORIAL MANAGER Mary Beth Wakefield

FREELANCER EDITORIAL MANAGER Rosemarie Graham

ASSOCIATE DIRECTOR OF MARKETING David Mayhew

MARKETING MANAGER Ashley Zurcher

BUSINESS MANAGER Amy Knies PRODUCTION MANAGER Tim Tate

VICE PRESIDENT AND EXECUTIVE GROUP PUBLISHER Richard Swadley

VICE PRESIDENT AND EXECUTIVE PUBLISHER Neil Edde

ASSOCIATE PUBLISHER Jim Minatel

PROJECT COORDINATOR, COVER Katie Crocker

PROOFREADER Louise Watson, Word One

INDEXER Ron Strauss

COVER DESIGNER Ryan Sneed

COVER IMAGE © Arthur S. Aubry / Getty Images

ABOUT THE AUTHORS

PAUL ATKINSON has been working with database technology since his days at Microsoft in the early 1990s, where he learned Microsoft Access first as a side project and later as a support rep. From there he went on to get a degree in Computer Science from Portland State University, rode the dot-com wave as a DBA, and came out of the dot-com bust as a contractor primarily in the health care industry. Currently, Paul works for Huron Healthcare as a BI Architect and Team Lead developing both traditional and real-time BI solutions. His training classes in high-performance T-SQL programming are among the most popular course offerings available at Huron. Paul, his wife Kristin, and their daughter Maureen love living in Portland, Oregon. When he's not working or writing Paul plays on multiple soccer teams, and otherwise keeps busy as a mountain biker, snowboarder, hasher, bicycle commuter, chairman of the school site council, proud member of the Timbers Army, and semi-professional chocolatier. As Robert Heinlein wrote, specialization is for insects.

ROBERT VIEIRA is a Software Architect with Huron Consulting Group and is considered one of the leading authorities on Microsoft SQL Server. He speaks at conferences nationally and is well known for his unique tutorial approach in his teaching and writing.

ABOUT THE TECHNICAL EDITOR

ERIC ALDINGER is a quality assurance engineer working out of Portland, Oregon. He spends much of his time working on thorny data integration and data warehousing issues on green field development projects. Editing this text proves the improbable axiom that a University of California, Riverside, literature degree is good for something outside of academia, given enough time. Eric has many hobbies, none of which are of interest to readers of this book.

ACKNOWLEDGMENTS

NOBODY COMPLETES A PROJECT like this alone. Crafting a technical book that's readable despite the fairly dry material found therein takes an enormous amount of time and dedication, and that's not from the author — it's from those who support, guide, and encourage him. While there are many who helped with this, my first foray, there are a few without whom I'd never have gotten through it.

Rob Vieira, whose efforts early on helped create the genre of the developer-centric SQL Server book, has probably contributed more to this title than anyone — even me. Besides laying down a frame-work, Rob has also been a mentor and a friend, and without him, I'd never have started on this effort at all. That said, I'm sure he'll be pleased when I've finished, since he's probably bored of my escaping work before 6:00 to come home and write.

Bob Elliott, for giving me the chance based on nothing but a recommendation (albeit a good one).

Kristin Atkinson, my wife of more than 15 years now; you have carried the house on your shoulders since I picked up this project. Each night I come home and sit down to write, and you've taken care of every other thing. Not only have you cleared the path, but you've made an excellent sounding board. This book is as much yours as it is anyone's, and by now you probably know more about SQL Server than any other artisan chocolatier I know.

Maureen Spears, you have phenomenal tolerance for my first-timer mistakes and questions, not to mention delays. Your help, attitude, and feedback have kept my enthusiasm up long after it could have flagged.

Eric Aldinger, by stepping up and enthusiastically accepting the role of technical editor, you've ensured the errata section on this title will be ridiculously clean. Once we're done here we need to spend less time on technical books and more on bikes.

CONTENTS

INTRODUCTION	XXV
CHAPTER 1: RDBMS BASICS: WHAT MAKES UP A SQL SERVER DATABASE?	1
An Overview of Database Objects	2
The Database Object	2
The Transaction Log	8
The Most Basic Database Object: Table	8
Filegroups	9
Diagrams	9
Views	10
Stored Procedures	11
User-Defined Functions	11
Sequences	12
Users and Roles	12
Rules	12
Defaults	12
User-Defined Data Types	13
Full-Text Catalogs	13
SQL Server Data Types	13
NULL Dala SOL Server Identifiers for Objects	19 19
What Cote Named?	20
Pulos for Naming	20
Summary	20 21
Summary	21
CHAPTER 2: LEARNING THE TOOLS OF THE TRADE	23
Getting Help with Books Online	24
SQL Server Configuration Manager	25
Service Management	26
Network Configuration	26
The Protocols	28
On to the Client	30
SQL Server Management Studio	31
Getting Started with the Management Studio	32
Interacting Through the Query Window	36
SQL Server Data Tools (formerly BIDS)	43

SQL Server Integration Services (SSIS) SQL Server Reporting Services (SSRS) SQL Server Analysis Services (SSAS) Bulk Copy Program (BCP) SQL Server Profiler sqlcmd PowerShell Summary	44 45 46 46 46 47 47
CHAPTER 3: THE FOUNDATION STATEMENTS OF T-SQL	49
Getting Started with a Basic SELECT Statement	50
The SELECT Statement and FROM Clause	51
The WHERE Clause	55
ORDER BY	58
Aggregating Data Using the GROUP BY Clause	61
Placing Conditions on Groups with the HAVING Clause	71
Outputting XML Using the FOR XML Clause	72
Making Use of Hints Using the OPTION Clause	73
Ine DISTINCT and ALL Predicates	/3
Adding Data with the INSERT Statement	/6
Multirow Inserts	٥١ دە
Changing What You've Get with the LIPDATE Statement	02 9 /
The DELETE Statement	87
Summary	88
	91
	51
Combining Table Data with JOINs	92
Selecting Matching Rows with INNER JOIN	93
How an INNER JOIN IS LIKE a WHERE Clause	99
The Simple OUTER JOIN	102
Dealing with More Complex OUTED JOINs	103
Seeing Both Sides with ELUL IOINs	109
Understanding CROSS IOINs	117
Exploring Alternative Syntax for Joins	118
An Alternative INNER JOIN	118
An Alternative OUTER JOIN	119
An Alternative CROSS JOIN	120
Stacking Results with UNION	121
Summary	125

CHAPTER 5: CREATING AND ALTERING TABLES	
Object Names in SQL Server	130
Schema Name (aka Ownership)	130
The Database Name	133
Naming by Server	133
Reviewing the Defaults	134
The CREATE Statement	134
CREATE DATABASE	134
Building a Database	139
CREATE TABLE	141
The ALTER Statement	155
ALTER DATABASE	156
ALTER TABLE	160
The DROP Statement	164
Using the GUI Tool	165
Creating a Database Using the Management Studio	165
Backing into the Code: Creating Scripts with the Management Studio	170
Summary	171
CHAPTER 6: KEYS AND CONSTRAINTS	175
Types of Constraints	176
Domain Constraints	177
Entity Constraints	177
Referential Integrity Constraints	178
Constraint Naming	178
Key Constraints	179
Primary Key Constraints	180
Foreign Key Constraints	183
UNIQUE Constraints	196
CHECK Constraints	197
DEFAULI Constraints	199
Defining a DEFAULT Constraint in Your CREATE	200
TABLE Statement	200
Adding a DEFAULT Constraint to an Existing Table	200
Disabiling Constraints	201
Ignoring Bad Data when You Create the Constraint	202
Pulse and Defaulte Cousing of Constraint	204
Rules and Detaults — Cousins of Constraints	206
Rules	206
Dropping Rules	208

Defaults	208
Dropping Defaults	209
Determining Which Tables and Data Types Use a Given Rule or Default	209
Triggers for Data Integrity	210
Choosing What to Use	210
Summary	211
CHAPTER 7: ADDING MORE TO YOUR QUERIES	213
What Is a Subquery?	214
Building a Nested Subquery	215
Building Correlated Subqueries	218
How Correlated Subqueries Work	219
Dealing with NULL Data — the ISNULL Function	223
Derived Tables	225
Using Common Table Expressions (CTEs)	228
Using the WITH Keyword	228
Using Multiple CTEs	231
Recursive CTEs	231
Using the EXISTS Operator	232
Filtering with EXISTS	232
Using EXISTS in Other Ways	234
Mixing Data Types: CAST and CONVERT	236
Synchronizing Data with the MERGE Command	239
The Merge Command in Action	240
A Brief Look at BY TARGET versus BY SOURCE	243
Gathering Affected Rows with the OUTPUT Clause	243
Through the Looking Glass: Windowing Functions	246
ROW_NUMBER	246
RANK, DENSE_RANK, and NTILE	249
One Chunk at a Time: Ad Hoc Query Paging	251
Performance Considerations	252
Measuring Performance	252
JOINs versus Subqueries versus CTEs versus	253
Summary	255
CHAPTER 8: BEING NORMAL: NORMALIZATION AND	
OTHER BASIC DESIGN ISSUES	257
Understanding Tables	258
Keeping Your Data "Normal"	258
Before the Beginning	260
The First Normal Form	261

The Second Normal Form	265
The Third Normal Form	267
Other Normal Forms	270
Understanding Relationships	270
One-to-One	271
Zero or One-to-One	272
One-to-One or Many	272
One-to-Zero, One, or Many	273
Many-to-Many	275
Diagramming Databases	279
Tables	282
Creating Relationships in Diagrams	288
Denormalization	291
Beyond Normalization	291
Keep It Simple	292
Choosing Data Types	292
Err on the Side of Storing Things	293
Drawing Up a Quick Example	293
Creating the Database	293
Adding the Diagram and the Initial Tables	294
Adding the Relationships	299
Adding Some Constraints	301
5	
Summary	302
Summary CHAPTER 9: SQL SERVER STORAGE AND INDEX STRUCTURES	302 305
Summary CHAPTER 9: SQL SERVER STORAGE AND INDEX STRUCTURES SQL Server Storage	302 305 306
Summary CHAPTER 9: SQL SERVER STORAGE AND INDEX STRUCTURES SQL Server Storage The Database	302 305 306 306
Summary CHAPTER 9: SQL SERVER STORAGE AND INDEX STRUCTURES SQL Server Storage The Database The Extent	302 305 306 306 306
Summary CHAPTER 9: SQL SERVER STORAGE AND INDEX STRUCTURES SQL Server Storage The Database The Extent The Page	302 305 306 306 306 307
Summary CHAPTER 9: SQL SERVER STORAGE AND INDEX STRUCTURES SQL Server Storage The Database The Extent The Page Rows	302 305 306 306 306 307 307
Summary CHAPTER 9: SQL SERVER STORAGE AND INDEX STRUCTURES SQL Server Storage The Database The Extent The Page Rows Sparse Columns	302 305 306 306 306 307 307 307
Summary CHAPTER 9: SQL SERVER STORAGE AND INDEX STRUCTURES SQL Server Storage The Database The Extent The Page Rows Sparse Columns Understanding Indexes	302 305 306 306 306 307 307 308 308 308
Summary CHAPTER 9: SQL SERVER STORAGE AND INDEX STRUCTURES SQL Server Storage The Database The Extent The Page Rows Sparse Columns Understanding Indexes B-Trees	302 305 306 306 306 307 307 307 308 308 308 310
Summary CHAPTER 9: SQL SERVER STORAGE AND INDEX STRUCTURES SQL Server Storage The Database The Extent The Page Rows Sparse Columns Understanding Indexes B-Trees How Data Is Accessed in SQL Server	302 305 306 306 306 307 307 308 308 308 308 310 313
Summary CHAPTER 9: SQL SERVER STORAGE AND INDEX STRUCTURES SQL Server Storage The Database The Extent The Page Rows Sparse Columns Understanding Indexes B-Trees How Data Is Accessed in SQL Server Creating, Altering, and Dropping Indexes	302 305 306 306 306 307 307 308 308 308 310 313 323
Summary CHAPTER 9: SQL SERVER STORAGE AND INDEX STRUCTURES SQL Server Storage The Database The Extent The Page Rows Sparse Columns Understanding Indexes B-Trees How Data Is Accessed in SQL Server Creating, Altering, and Dropping Indexes The CREATE INDEX Statement	302 305 306 306 306 307 307 307 308 308 308 310 313 323 323
Summary CHAPTER 9: SQL SERVER STORAGE AND INDEX STRUCTURES SQL Server Storage The Database The Extent The Page Rows Sparse Columns Understanding Indexes B-Trees How Data Is Accessed in SQL Server Creating, Altering, and Dropping Indexes The CREATE INDEX Statement Creating XML Indexes	302 305 306 306 307 307 308 308 308 310 313 323 323 330
Summary CHAPTER 9: SQL SERVER STORAGE AND INDEX STRUCTURES SQL Server Storage The Database The Database The Extent The Page Rows Sparse Columns Understanding Indexes B-Trees How Data Is Accessed in SQL Server Creating, Altering, and Dropping Indexes The CREATE INDEX Statement Creating XML Indexes Implied Indexes Created with Constraints	302 305 306 306 306 307 307 308 308 310 313 323 323 320 330 330
Summary CHAPTER 9: SQL SERVER STORAGE AND INDEX STRUCTURES SQL Server Storage The Database The Extent The Page Rows Sparse Columns Understanding Indexes B-Trees How Data Is Accessed in SQL Server Creating, Altering, and Dropping Indexes The CREATE INDEX Statement Creating XML Indexes Implied Indexes or Sparse and Geospatial Columns	302 305 306 306 306 307 307 308 308 308 310 313 323 323 330 330 330 331
Summary CHAPTER 9: SQL SERVER STORAGE AND INDEX STRUCTURES SQL Server Storage The Database The Extent The Page Rows Sparse Columns Understanding Indexes B-Trees How Data Is Accessed in SQL Server Creating, Altering, and Dropping Indexes The CREATE INDEX Statement Creating XML Indexes Implied Indexes Created with Constraints Creating Indexes on Sparse and Geospatial Columns Choosing Wisely: Deciding Which Index Goes	302 305 306 306 307 307 308 308 308 310 313 323 323 330 330 331
Summary CHAPTER 9: SQL SERVER STORAGE AND INDEX STRUCTURES SQL Server Storage The Database The Database The Extent The Page Rows Sparse Columns Understanding Indexes B-Trees How Data Is Accessed in SQL Server Creating, Altering, and Dropping Indexes The CREATE INDEX Statement Creating XML Indexes Implied Indexes Created with Constraints Creating Indexes on Sparse and Geospatial Columns Choosing Wisely: Deciding Which Index Goes Where and When	302 305 306 306 307 307 308 308 308 310 313 323 323 323 330 330 331 331

Watching Costs: When Less Is More Choosing That Clustered Index Column Order Matters Covering Indexes Filtered Indexes ALTER INDEX Dropping Indexes Taking a Hint from the Query Plan Why Isn't Your Index Being Used? Use the Database Engine Tuning Advisor Maintaining Your Indexes Fragmentation Identifying Fragmentation versus Likelihood of Page Splits	335 336 338 339 340 340 343 343 343 344 345 345 345 346 350
CHAPTER 10: VIEWS	353
Creating Simple Views Views as Filters Editing Views with T-SQL Dropping Views Creating and Editing Views in the Management Studio Creating Views in Management Studio Editing Views in the Management Studio Auditing: Displaying Existing Code Protecting Code: Encrypting Views About Schema Binding Making Your View Look Like a Table with VIEW_METADATA Indexed (Materialized) Views Indexing an Aggregate View Summary	354 357 367 367 368 371 373 371 373 374 375 375 378 381
CHAPTER 11: WRITING SCRIPTS AND BATCHES	383
Understanding Script Basics Selecting a Database Context with the USE Statement Declaring Variables Setting the Value in Your Variables Reviewing System Functions Retrieving IDENTITY Values Generating SEQUENCES Using @@ROWCOUNT	384 385 385 386 388 390 395 398

Grouping Statements into Batches	399
A Line to Itself	400
Each Batch Is Sent to the Server Separately	400
GO Is Not a T-SQL Command	401
Errors in Batches	401
When to Use Batches	402
Running from the Command Prompt: sqlcmd	405
Dynamic SQL: Using the EXEC Command	409
Generating Your Code on the Fly	410
Understanding the Dangers of Dynamic SQL	411
Using Control-of-Flow Statements	415
The IF ELSE Statement	416
The CASE Statement	420
Looping with the WHILE Statement	426
The WAITFOR Statement	428
Dealing with Errors with TRY/CATCH Blocks	428
Summary	431
CHAPTER 12: STORED PROCEDURES	435
Creating the Sproc: Basic Syntax	436
An Example of a Basic Sproc	436
Changing Stored Procedures with ALTER	438
Dropping Sprocs	438
Parameterizing Sprocs	438
Declaring Parameters	439
Confirming Success or Failure with Return Values	445
How to Use RETURN	445
More on Dealing with Errors	448
Handling Inline Errors	449
Making Use of @@ERROR	450
Using @@ERROR in a Sproc	452
Handling Errors Before They Happen	455
Manually Raising Errors	461
Re-throwing Errors	465
Adding Your Own Custom Error Messages	467
What a Sproc Offers	468
Creating Callable Processes	469
Using Sprocs for Security	469
Sprocs and Performance	469
Extended Stored Procedures (XPs)	472
A Brief Look at Recursion	472

Debugging	475
Starting the Debugger	475
Parts of the Debugger	476
Taking a Break: Breakpoints	479
Using the Debugger Once It's Started	481
Understanding .NET Assemblies	484
When to Use Stored Procedures	485
Summary	486
CHAPTER 13: USER-DEFINED FUNCTIONS	489
What a UDF Is	489
UDFs Returning a Scalar Value	491
UDFs That Return a Table	496
Inline UDFs	496
Understanding Determinism	503
Debugging User-Defined Functions	506
Using .NET in a Database World	507
Summary	507
CHAPTER 14: TRANSACTIONS AND LOCKS	509
Understanding Transactions	509
ACID Transactions	510
Introducing Transaction Operations	510
Using BEGIN TRAN	511
Using COMMIT TRAN	511
Using ROLLBACK TRAN	512
Using SAVE TRAN	512
How the SQL Server Log Works	514
Using the Log for Failure and Recovery	516
Activating Implicit Transactions	517
Understanding Locks and Concurrency	518
Clarifying the Purpose of Locks	518
Defining Lockable Resources	522
Lock Escalation and Lock Effects on Performance	522
Understanding Lock Modes	523
	010
Clarifying Lock Compatibility	525
Clarifying Lock Compatibility Specifying a Specific Lock Type — Optimizer Hints	525 526
Clarifying Lock Compatibility Specifying a Specific Lock Type — Optimizer Hints Setting the Isolation Level	525 526 527
Clarifying Lock Compatibility Specifying a Specific Lock Type — Optimizer Hints Setting the Isolation Level Setting READ COMMITTED	525 526 527 528

Setting REPEATABLE READ	529
Setting SERIALIZABLE	529
Setting SNAPSHOT	530
Dealing with Deadlocks (aka "a 1205")	531
How SQL Server Figures Out There's a Deadlock	531
How Deadlock Victims Are Chosen	531
Avoiding Deadlocks	532
Summary	534
CHAPTER 15: TRIGGERS	537
What Is a Trigger?	538
ON	540
WITH ENCRYPTION	540
FORIAFTER	540
The FORIAFTER versus the INSTEAD OF Clause	541
NOT FOR REPLICATION	543
AS	543
Using Triggers for Data Integrity Rules	543
Dealing with Requirements Sourced from Other Tables	544
Using Triggers to Check the Delta of an Update	545
Using Triggers for Custom Error Messages	547
Other Common Uses for Triggers	548
Other Trigger Issues	548
Triggers Can Be Nested	548
Triggers Can Be Recursive	549
Triggers Don't Prevent Architectural Changes	549
Triggers Can Be Turned Off without Being Removed	550
Trigger Firing Order	550
INSTEAD OF Triggers	552
Performance Considerations	552
Triggers Are Reactive Rather Than Proactive	552
Triggers Don't Have Concurrency Issues with the Process That Fires Them	553
Using IF UPDATE() and COLUMNS_UPDATED	553
Keep It Short and Sweet	556
Don't Forget Triggers When Choosing Indexes	556
Try Not to Roll Back within Triggers	556
Dropping Triggers	556
Debugging Triggers	557
Summary	558

CHAPTER 16: A BRIEF XML PRIMER	561
XML Basics	562
Parts of an XML Document	563
Namespaces	574
Element Content	576
Being Valid versus Being Well Formed — Schemas and DTDs	576
What SQL Server Brings to the Party	577
Defining a Column as Being of XML Type	578
XML Schema Collections	579
Creating, Altering, and Dropping XML Schema Collections	581
XML Data Type Methods	583
Enforcing Constraints Beyond the Schema Collection	591
Retrieving Relational Data in XML Format	592
RAW	593
AUTO	596
EXPLICIT	597
PATH	599
OPENXML	605
A Brief Word on XSLT	610
Summary	613
CHAPTER 17: BUSINESS INTELLIGENCE FUNDAMENTALS	615
What Is Business Intelligence?	616
BI Is Better Than the Alternative	617
Making Decisions Based on Data	617
Five Experts, Five Opinions	618
Past Performance Predicts Future Behavior	618
Who Uses BI?	619
Those Who Forget History: The Data Warehouse	619
The Purpose of the Data Warehouse	620
OLTP versus OLAP	620
Kimball versus Inmon	621
Dimensional Modeling: Why Be Normal?	624
Measures and Metrics	625
Facts	626
1 4615	020
Dimensions	629
Dimensions ETLs	629 637
Dimensions ETLs Sourcing Your Data	629 637 637
Dimensions ETLs Sourcing Your Data Enforcing Conformity	629 637 637 638
Dimensions ETLs Sourcing Your Data Enforcing Conformity Resolving Conflicts by Defining Authoritative Data	629 637 637 638 639

Dealing with Concurrency During ETLs	640
Caching Your Extract — The Operational Data Store (ODS)	641
Making Your Data Actionable: BI Reporting Techniques	642
Goals	642
Key Performance Indicators	643
Dashboards	644
Scorecards	645
Detail Reports	645
Drill-Down	645
Drill-Through	646
Ad Hoc Reporting	646
Summary	647
CHAPTER 18: BI STORAGE AND REPORTING RESOURCES	649
BI in T-SQL	650
Introducing the AdventureWorks Data Warehouse	650
Using T-SQL Analysis Functions	651
Columnstore Indexes	660
SQL Server Analysis Services, or How I Learned to Stop Worrying	
and Love the Cube	661
Understanding SSAS Modeling	662
Building Your First Cube	663
Data Sources	669
Data Source Views	670
Cubes	671
Measure Groups	671
Dimensions	672
The Rest of the Cube	674
Self-Service BI: User Tools	675
Report Builder	675
PowerPivot: Lightweight BI with SharePoint and Excel	675
Power View	677
Summary	678
CHAPTER 19: REPORTING FOR DUTY, SIR! A LOOK AT	
REPORTING SERVICES	681
Reporting Services 101	682
Understanding the SSRS Report Lifecycle	683
Developing Reports in SSRS	683
Managing Your Reports	683
Delivering Reports to Users	684

Understanding the Reporting Services Architecture	684
SSDT: Report Designer	684
SSRS Windows Service	684
Report Manager	685
Report Builder	685
Report Viewer	685
Building Simple Report Models	686
Creating a Project	686
Adding to Your Data Model	695
Building and Deploying Your Data Model	701
Report Creation	702
Report Server Projects	711
Building a Report Server Project	711
Deploying the Report	716
Summary	717
CHAPTER 20: GETTING INTEGRATED WITH	
INTEGRATION SERVICES	719
Understanding the Problem	720
Using the Import/Export Wizard to Generate Basic Packages	720
Starting the Wizard	721
Choosing Your Data Sources	721
Completing the Wizard	725
Examining Package Basics	727
Connection Managers	728
The Package Editor Pane	730
Executing Packages	738
Executing a Package inside the Development Studio	739
Executing within SSMS	739
Using the Execute Package Utility	742
Executing Using the SQL Server Agent	744
Executing a Package from within a Program	744
A Final Word on Packages	745
Summary	745
CHAPTER 21: PLAYING ADMINISTRATOR	747
Scheduling Jobs	748
	/40
Creating an Operator	749
Creating an Operator Creating Jobs and Steps	749 751
Creating an Operator Creating Jobs and Steps Logins and Users	749 751 760
Creating an Operator Creating Jobs and Steps Logins and Users Backup and Recovery	749 751 760 764

APPENDIX: ANSWERS TO EXERCISES	783
Summary	779
Deployment and Execution	778
Sqlps	776
Cmdlets and Modules	776
Installing and Enabling PowerShell	775
What Is PowerShell?	775
Automating Administration Tasks with PowerShell	775
Policy Based Management	774
ALTER INDEX	771
Index Maintenance	770
Recovery	769
Recovery Models	767

INDEX

807

ONLINE APPENDICES

BONUS APPENDIX 1: SYSTEM FUNCTIONS	1
System Statistical Functions	2
Aggregate Functions	4
Analytic Functions	8
Configuration Functions	10
Conversion Functions	15
Cryptographic Functions	16
Cursor Functions	20
Date and Time Functions	22
Hierarchy Functions	27
Mathematical Functions	28
Basic Metadata Functions	33
Rowset Functions	47
Security Functions	49
String Functions	52
System Functions	57
Text and Image Functions	67
BONUS APPENDIX 2: VERY SIMPLE CONNECTIVITY EXAMPLES	69
Some General Concepts	70
Connectivity Examples	71
A Brief Word on LINQ	76

INTRODUCTION

OUT OF EVERY ENDING COMES the beginning of something new. This title has been Rob Vieira's for many years, and now he's wrapping up that chapter of his life while I begin a new chapter of my own — and the first chapter of this text. Likewise, you, as a reader, are also entering something of a transition; you know something about programming, probably, but you're about to begin a completely new experience as a programmer of relational databases in general and Microsoft SQL Server 2012 in particular.

Database programming is a pretty big shift from most other kinds of software engineering. If you're at all like me, something brought you from your early programming experience (6502-based machines in BASIC for me) to databases. For me it happened at Microsoft in the early '90s, supporting a beta product that flew under the radar; I was asked to track certain aspects of the work I was doing in Access (2.0 at that time). Once I learned to do that, I was hooked, and a couple years on the Access support queue (learning SQL Server 4.2 and beyond as well) sealed the deal.

Working with SQL Server is a whole different animal from ordinary procedural programming. You get to think in mathematical and set-based terms, and learn how to ask carefully for what you want without spelling out how to actually accomplish the work. Transitioning from procedural programming to this kind of thinking without help is like trying to make a paradigm shift without a clutch. And yet this language, SQL, has a certain simplicity to it sometimes that makes it a pleasure to work with, once you learn how to think like it thinks.

I learned from immersion, from Microsoft's internal trainers, and from peers; what I wished for at the time was a book I could read that would give me the concepts and the functional knowledge to understand what I was seeing and know what was out there that I didn't know about yet. This book is the book I wanted, which means if you're in that early learning phase with T-SQL, it's probably the book you need as well.

This is a step-by-step tutorial, providing you the concepts you need in bite-sized pieces presented in an orderly way, each building on the last. The whole reason it exists is that you'd likely have a terrible time picking up such a completely new set of concepts by choosing topics out of an online help database. Books are still, in 2012, a great way to learn new ideas.

My hope is that, in this book, you find something that covers all of the core elements of SQL Server with the same success that we had in the original *Professional SQL Server Programming* titles. When you're done, you should be set to be a highly functional SQL Server 2012 programmer and, with any luck, you'll enjoy the unique challenges of database programming for years to come.

WHO THIS BOOK IS FOR

It is almost sad that the word "beginner" is in the title of this book. Don't get me wrong; if you are a beginner, this title is for you. But it is designed to last you well beyond your beginning days. What is covered in this book is necessary for the beginner, but there is simply too much information

for you to remember all of it all the time, and so it is laid out in a fashion that should make a solid review and reference book even for the more intermediate, and, yes, even advanced user.

The beginning user will want to start right at the beginning. Things are designed such that just about everything in this book is a genuine "need to know" sort of thing. With the possible exception of the chapters on XML, Reporting Services, and Integration Services, every item in this book is fundamental to you having the breadth of understanding you need to make well-informed choices on how you approach your SQL Server problems. Even these three topics are increasingly fundamental to being a serious SQL Server developer.

For the intermediate user, you can probably skip perhaps as far as Chapter 7 or 8 to start. Although I still recommend scanning the prior chapters for holes in your skills or general review, you can probably skip ahead with little harm done and get to something that might be a bit more challenging for you.

Advanced users, in addition to utilizing this as an excellent reference resource, will probably want to focus on Chapter 12 and beyond. Virtually everything from that point forward should be of some interest (the new debugging, transactions, XML, BI, Reporting Services, Integration Services, and more!).

WHAT THIS BOOK COVERS

Well, if you've read the title, you're probably not shocked to hear that this book covers SQL Server 2012 with a definite bent toward the developer's perspective.

SQL Server 2012 is the latest incarnation of a database management system that has now been around for more than two decades. It builds on the base redesign that was done to the product in version 7.0 — this time adding some brilliant analysis and reporting functionality, among other things. This book focuses on core development needs of every developer, regardless of skill level. The focus is highly oriented to just the 2012 version of the product, but there is regular mention of backward-compatibility issues, as they may affect your design and coding choices.

HOW THIS BOOK IS STRUCTURED

The book is designed to become increasingly more advanced as you progress through it, but, from the very beginning, I'm assuming that you are already an experienced developer — just not necessarily with databases. In order to make it through this book, you do need to already have understanding of programming basics such as variables, data types, and procedural programming. You do not have to have seen a query before in your life (though I suspect you have).

The focus of the book is highly developer-oriented. This means that it will, for the sake of both brevity and sanity, sometimes gloss over or totally ignore items that are more the purview of the database administrator than the developer. You will, however, be reminded of administration issues either as they affect the developer or as they need to be thought of during the development process — you'll also take a brief look at several administration-related issues in Chapter 21.

The book makes a very concerted effort to be language independent in terms of your client-side development. VB, C#, C++, Java, and other languages are generally ignored (it focuses on the server side of the equation) and treated equally where addressed.

In terms of learning order, you'll start by learning the foundation objects of SQL, and then move on to basic queries and joins. From there, you can begin adding objects to your database and discuss items that are important to the physical design — then it is on to the more robust code aspects of SQL Server scripting, stored procedures, user-defined functions, and triggers. After a short tutorial on business intelligence, you'll have a look at a few of the relatively peripheral features of SQL Server. Last but not least, you can wrap things up with a bit of important material on administration meant to help you keep the databases you develop nice and healthy.

NOTE This book is a tutorial, but there's reference material online you might want to refer to even after you're done. This bonus material, including a full reference containing system functions and a document of very simple connectivity examples, is available at www.wrox.com. Once at the site, search for the book by its ISBN, which is 978-1-118-102282.

WHAT YOU NEED TO USE THIS BOOK

In order to make any real, viable use of this book, you will need an installation of SQL Server. The book makes extensive use of the actual SQL Server 2012 management tools, so I highly recommend that you have a version that contains the full product, rather than just using SQL Server Express. That said, the book is focused on the kind of scripting required for developers, so even SQL Server Express users should be able to get the lion's share of learning out of most of the chapters. You will also need the AdvenureWorks sample database, the AdventureWorks database for BI and reporting, and a few custom databases installed. Instructions for accessing these databases can be found in the ReadMe file on this book's website (www.wrox.com).

A copy of Visual Studio is handy for working with this book, but most of the Visual Studio features needed are included in the Business Intelligence Studio that comes along with the SQL Server product.

CONVENTIONS

To help you get the most from the text and keep track of what's happening, I've used a number of conventions throughout the book.

TRY IT OUT

The Try It Out is an exercise you should work through, following the text in the book.

1. The exercises usually consist of a set of steps.

- **2.** Each step has a number.
- **3.** Follow the steps through with your copy of the database.

How It Works

After each Try It Out, the code you've typed will be explained in detail.

WARNING Boxes with a warning icon like this one hold important, not-to-beforgotten information that is directly relevant to the surrounding text.

NOTE The pencil icon indicates notes, tips, hints, tricks, and asides to the current discussion.

- > We *italicize* new terms and important words when we introduce them.
- ▶ We show keyboard strokes like this: Ctrl+A.
- > We show filenames, URLs, and code within the text like so: persistence.properties.
- > We present code in two different ways:

We use a monofont type with no highlighting for most code examples.

We use bold to emphasize code that's particularly important in the present context.

SOURCE CODE

As you work through the examples in this book, you may choose either to type in all the code manually or to use the source code files that accompany the book. All the source code used in this book is available for download at www.wrox.com. When at the site, simply locate the book's title (use the Search box or one of the title lists) and click the Download Code link on the book's detail page to obtain all the source code for the book. Code that is included on the website is highlighted by the following icon:



Listings include the filename in the title. If it is just a code snippet, you'll find the filename in a code note such as this:

Code snippet filename