

Contents

Acknowledgments	v
Introduction	xxiii
Chapter 1: Being Objective: Re-Examining Objects in SQL Server	1
So, What Exactly Do We Have Here?	1
An Overview of Database Objects	2
The Database Object	2
The Transaction Log	5
The Most Basic Database Object: Table	5
Schemas	6
Filegroups	7
Diagrams	7
Views	7
Stored Procedures	9
User-Defined Functions	9
Users and Roles	10
Rules	10
Defaults	10
User-Defined Data Types	10
Full-Text Catalogs	11
SQL Server Data Types	11
NULL Data	14
SQL Server Identifiers for Objects	15
What Gets Named?	15
Rules for Naming	15
Summary	16
Chapter 2: Tool Time	17
Books Online	18
The SQL Server Configuration Manager	19
Service Management	19
Network Configuration	20
The Protocols	21
On to the Client	23

Contents

The SQL Server Management Studio	24
Getting Started	25
Query Window	27
SQL Server Business Intelligence Development Studio	32
SQL Server Integration Services (SSIS)	32
Reporting Services	33
Bulk Copy Program (bcp)	33
SQL Server Profiler	33
sqlcmd	34
Summary	34
Chapter 3: Basic T-SQL	35
The Basic SELECT Statement	36
The SELECT Statement and FROM Clause	36
The JOIN Clause	38
The WHERE Clause	45
ORDER BY	49
Aggregating Data Using the GROUP BY Clause	52
Placing Conditions on Groups with the HAVING Clause	56
Outputting XML Using the FOR XML Clause	57
Making Use of Hints Using the OPTION Clause	57
DISTINCT	57
Adding Data with the INSERT Statement	59
The INSERT INTO . . . SELECT Statement	61
Changing What You've Got with the UPDATE Statement	62
The DELETE Statement	64
Exploring Alternative Syntax for Joins	66
An Alternative INNER JOIN	67
An Alternative OUTER JOIN	67
An Alternative CROSS JOIN	68
The UNION	69
Summary	72
Chapter 4: Creating and Altering Tables	73
Object Names in SQL Server	73
Schema Name (a.k.a. Ownership)	74
The Database Name	76
Naming by Server	76
The CREATE Statement	77
CREATE DATABASE	77
CREATE TABLE	82

The ALTER Statement	89
ALTER DATABASE	89
ALTER TABLE	92
The DROP Statement	95
Using the GUI Tool	96
Creating or Editing the Database	96
Creating and Editing Tables	97
Summary	99
Chapter 5: Reviewing Keys and Constraints	101
Types of Constraints	102
Domain Constraints	103
Entity Constraints	103
Referential Integrity Constraints	104
Constraint Naming	104
Key Constraints	105
PRIMARY KEY Constraints	106
FOREIGN KEY Constraints	108
UNIQUE Constraints	117
CHECK Constraints	118
DEFAULT Constraints	119
Defining a DEFAULT Constraint in Your CREATE TABLE Statement	120
Adding a DEFAULT Constraint to an Existing Table	121
Disabling Constraints	121
Ignoring Bad Data When You Create the Constraint	122
Temporarily Disabling an Existing Constraint	124
Rules and Defaults: Cousins of Constraints	126
Rules	126
Defaults	128
Determining Which Tables and Data Types Use a Given Rule or Default	129
Triggers for Data Integrity	129
Choosing What to Use	129
Summary	131
Chapter 6: Asking a Better Question: Advanced Queries	133
What Is a Subquery?	134
Building a Nested Subquery	135
Nested Queries Using Single Value SELECT Statements	136
Nested Queries Using Subqueries That Return Multiple Values	137
The ANY, SOME, and ALL Operators	138

Contents

Correlated Subqueries	139
How Correlated Subqueries Work	140
Correlated Subqueries in the WHERE Clause	140
Correlated Subqueries in the SELECT List	142
Derived Tables	144
The EXISTS Operator	146
Using EXISTS in Other Ways	147
Mixing Data Types: CAST and CONVERT	148
Using External Calls to Perform Complex Actions	150
Performance Considerations	151
JOINS vs. Subqueries vs. ?	152
Summary	153
Chapter 7: Daring to Design	155
Normalization 201	155
Where to Begin	156
Getting to Third Normal Form	157
Other Normal Forms	157
Relationships	158
Diagramming	159
A Couple of Relationship Types	160
The Entity Box	160
The Relationship Line	161
Terminators	162
Logical versus Physical Design	165
Purpose of a Logical Model	165
Parts of a Logical Model	166
Dealing with File-Based Information	168
Subcategories	171
Types of Subcategories	172
Keeping Track of What's What — Implementing Subcategories	173
Getting Physical — The Physical Implementation of Subcategories	175
Adding to Extensibility with Subcategories	176
Database Reuse	177
Candidates for Reusable Databases	177
How to Break Things Up	178
The High Price of Reusability	179
De-Normalization	179
Partitioning for Scalability	180

The SQL Server Diagramming Tools	181
Tables	183
Dealing with Constraints	185
Summary	187
Chapter 8: SQL Server — Storage and Index Structures	189
<hr/>	
SQL Server Storage	189
The Database	189
The File	190
The Extent	190
The Page	191
Rows	193
Full-Text Catalogs	194
Understanding Indexes	194
To “B,” or Not to “B”: B-Trees	195
How Data Is Accessed in SQL Server	199
Index Types and Index Navigation	200
Creating, Altering, and Dropping Indexes	208
The CREATE INDEX Statement	208
Creating XML Indexes	214
Implied Indexes Created with Constraints	215
ALTER INDEX	215
DROP INDEX	218
Choosing Wisely: Deciding What Index Goes Where and When	218
Selectivity	218
Watching Costs: When Less Is More	219
Choosing That Clustered Index	219
Column Order Matters	222
Dropping Indexes	222
Use the Database Engine Tuning Advisor	222
Maintaining Your Indexes	223
Fragmentation	223
Identifying Fragmentation	224
Summary	228
Chapter 9: Views	231
<hr/>	
Simple Views	232
More Complex Views	233
Using a View to Change Data — Before INSTEAD OF Triggers	236
Editing Views with T-SQL	237

Contents

Dropping Views	238
Auditing: Displaying Existing Code	238
Protecting Code: Encrypting Views	239
About Schema Binding	241
Making Your View Look Like a Table with VIEW_METADATA	241
Indexed (Materialized) Views	242
Partitioned Views	244
Summary	244
Chapter 10: Scripts and Batches	247
Script Basics	248
The USE Statement	248
Declaring Variables	249
Using @@IDENTITY	252
Using @@ROWCOUNT	252
Batches	253
Errors in Batches	255
When to Use Batches	256
SQLCMD	259
Dynamic SQL: Generating Your Code on the Fly with the EXEC Command	260
The Gotchas of EXEC	262
Control-of-Flow Statements	265
The IF . . . ELSE Statement	266
The CASE Statement	270
Looping with the WHILE Statement	275
The WAITFOR Statement	277
TRY/CATCH Blocks	277
Summary	280
Chapter 11: Getting Procedural: Stored Procedures and User-Defined Functions	281
Creating the Sproc: Basic Syntax	282
An Example of a Basic Sproc	282
Changing Stored Procedures with ALTER	283
Dropping Sprocs	283
Parameterization	284
Declaring Parameters	284
Creating Output Parameters	285
Confirming Success or Failure with Return Values	288
How to Use RETURN	288

Dealing with Errors	290
The Way We Were	291
Manually Raising Errors	296
Adding Your Own Custom Error Messages	299
What a Sproc Offers	301
Creating Callable Processes	301
Using Sprocs for Security	301
Sprocs and Performance	302
Extended Stored Procedures (XPs)	304
A Brief Look at Recursion	305
User-Defined Functions (UDFs)	307
What a UDF Is	308
UDFs Returning a Scalar Value	308
UDFs That Return a Table	310
Understanding Determinism	316
Debugging	317
Setting Up SQL Server for Debugging	318
Starting the Debugger	318
Parts of the Debugger	321
Using the Debugger after It's Started	322
Summary	327
Chapter 12: Transactions and Locks	329
Transactions	329
BEGIN TRAN	331
COMMIT TRAN	331
ROLLBACK TRAN	331
SAVE TRAN	331
How the SQL Server Log Works	336
Using the CHECKPOINT Command	337
CHECKPOINT on Recovery	337
At Normal Server Shutdown	338
At a Change of Database Options	338
When the Truncate on Checkpoint Option Is Active	338
When Recovery Time Would Exceed the Recovery Interval Option Setting	339
Failure and Recovery	339
Implicit Transactions	340
Locks and Concurrency	341
What Problems Can Be Prevented by Locks	342
Lockable Resources	346
Lock Escalation and Lock Effects on Performance	346

Contents

Lock Modes	347
Lock Compatibility	349
Specifying a Specific Lock Type — Optimizer Hints	349
Setting the Isolation Level	353
READ COMMITTED	353
READ UNCOMMITTED	354
REPEATABLE READ	354
SERIALIZABLE	355
Dealing with Deadlocks (a.k.a. “A 1205”)	355
How SQL Server Figures Out There’s a Deadlock	356
How Deadlock Victims Are Chosen	356
Avoiding Deadlocks	356
Summary	358
Chapter 13: Triggers	361
What Is a Trigger?	362
ON	363
WITH ENCRYPTION	363
The FOR AFTER versus the INSTEAD OF Clause	364
WITH APPEND	366
NOT FOR REPLICATION	367
AS	367
Using Triggers for Data Integrity Rules	367
Dealing with Requirements Sourced from Other Tables	367
Using Triggers to Check the Delta of an Update	369
Using Triggers for Custom Error Messages	371
Other Common Uses for Triggers	372
Updating Summary Information	372
Feeding Data into De-normalized Tables for Reporting	372
Setting Condition Flags	373
Other Trigger Issues	375
Triggers Can Be Nested	376
Triggers Can Be Recursive	376
Debugging Triggers	376
Triggers Don’t Get in the Way of Architecture Changes	377
Triggers Can Be Turned Off without Being Removed	377
Trigger Firing Order	378
INSTEAD OF Triggers	380
INSTEAD OF INSERT Triggers	381
INSTEAD OF UPDATE Triggers	384
INSTEAD OF DELETE Triggers	384

IF UPDATE() and COLUMNS_UPDATED	386
The UPDATE() Function	386
The COLUMNS_UPDATED() Function	386
Performance Considerations	388
Triggers Are Reactive Rather Than Proactive	388
Triggers Don't Have Concurrency Issues with the Process That Fires Them	389
Keep It Short and Sweet	389
Don't Forget Triggers When Choosing Indexes	389
Try Not to Roll Back within Triggers	389
Dropping Triggers	390
Debugging Triggers	390
Summary	392
Chapter 14: Nothing But NET!	393
Assemblies 101	394
Compiling an Assembly	394
Uploading Your Assembly to SQL Server	397
Creating Your Assembly-Based Stored Procedure	398
Creating Scalar User-Defined Functions from Assemblies	400
Creating Table-Valued Functions	403
Creating Aggregate Functions	407
Creating Triggers from Assemblies	412
Custom Data Types	417
Creating Your Data Type from Your Assembly	418
Accessing Your Complex Data Type	418
Dropping Data Types	419
Summary	419
Chapter 15: SQL Cursors	421
What Is a Cursor?	421
The Lifespan of a Cursor	422
Types of Cursors and Extended Declaration Syntax	427
Scope	428
Scrollability	432
Cursor Types	434
Concurrency Options	447
Detecting Conversion of Cursor Types: TYPE_WARNING	450
FOR <SELECT>	452
FOR UPDATE	452
Navigating the Cursor: The FETCH Statement	452

Contents

Altering Data within Your Cursor	453
Summary	456
Chapter 16: XML Integration	457
The XML Data Type	458
Defining a Column as Being of XML Type	458
XML Schema Collections	460
Creating, Altering, and Dropping XML Schema Collections	461
XML Data Type Methods	463
Enforcing Constraints beyond the Schema Collection	469
Retrieving Relational Data in XML Format	469
The FOR XML Clause	469
OPENXML	493
A Quick (Very Quick) Reminder of XML Indexes	497
HTTP Endpoints	497
Security	498
HTTP Endpoint Methods	499
Creating and Managing a HTTP Endpoint	499
Closing Thoughts	500
Summary	501
Chapter 17: Reporting for Duty, Sir!	503
Reporting Services 101	504
Building Simple Report Models	504
Data Source Views	507
Report Creation	512
Report Server Projects	517
Deploying the Report	522
Summary	523
Chapter 18: Buying in Bulk: the Bulk Copy Program (BCP) and Other Basic Bulk Operations	525
BCP Utility	526
BCP Syntax	526
BCP Import	531
BCP Export	534
Format Files	536
When Your Columns Don't Match	538
Using Format Files	540
Maximizing Import Performance	541

BULK INSERT	541
OPENROWSET (BULK)	542
ROWS_PER_BATCH	543
SINGLE_BLOB, SINGLE_CLOB, SINGLE_NCLOB	543
Summary	543
Chapter 19: Getting Integrated	545
Understanding the Problem	545
An Overview of Packages	546
Tasks	548
The Main Window	551
Solution Explorer	552
The Properties Window	552
Building a Simple Package	552
Executing Packages	560
Using the Execute Package Utility	560
Executing within Management Studio	562
Summary	563
Chapter 20: Replication	565
Replication Basics	566
Considerations When Planning for Replication	566
Replication Roles	568
Subscriptions	569
Types of Subscribers	570
Filtering Data	570
Replication Models	570
Snapshot Replication	571
Merge Replication	574
Transactional Replication	577
Immediate-Update Subscribers	580
Mixing Replication Types	581
Replication Topology	581
Simple Models	581
Mixed Models	584
Planning for Replication	587
Data Concerns	587
Mobile Devices	588
Setting Up Replication in Management Studio	588
Configuring the Server for Replication	588
Configuring a Publication	592

Contents

Setting Up Subscribers (via Management Studio)	598
Using Our Replicated Database	603
Replication Management Objects (RMO)	605
Summary	606
Chapter 21: Looking at Things in Full: Full-Text Search	607
Full-Text Search Architecture	608
Setting Up Full-Text Indexes and Catalogs	610
Enabling Full-Text for Your Database	610
Creating, Altering, Dropping, and Manipulating a Full-Text Catalog	611
Creating, Altering, Dropping, and Manipulating Full-Text Indexes	614
Creating Full-Text Catalogs Using the Old Syntax	619
Old Syntax for Indexes	621
More on Index Population	622
Full-Text Query Syntax	624
CONTAINS	624
FREETEXT	626
CONTAINSTABLE	626
FREETEXTTABLE	628
Dealing with Phrases	628
Booleans	629
Proximity	629
Weighting	630
Inflectional	631
Noise Words	631
Summary	632
Chapter 22: Security	633
Security Basics	634
One Person, One Login, One Password	634
Password Expiration	635
Password Length and Makeup	637
Number of Tries to Log In	637
Storage of User and Password Information	637
Security Options	638
SQL Server Security	639
Creating and Managing Logins	640
Windows Integrated Security	646

User Rights	647
Granting Access to a Specific Database	647
Granting Object Permissions within the Database	648
User Rights and Statement-Level Permissions	654
Server and Database Roles	655
Server Roles	657
Database Roles	658
Application Roles	661
Creating Application Roles	662
Adding Permissions to the Application Role	662
Using the Application Role	662
Getting Rid of Application Roles	663
More Advanced Security	664
What to Do about the guest Account	664
TCP/IP Port Settings	664
Don't Use the sa Account	665
Keep xp_cmdshell under Wraps	665
Don't Forget Views, Stored Procedures, and UDFs as Security Tools	665
Certificates and Asymmetric Keys	666
Certificates	667
Asymmetric Keys	667
Summary	667
Chapter 23: Playing a Good Tune: Performance Tuning	669
When to Tune	670
Index Choices	671
Check the Index Tuning Tool in the Database Tuning Advisor	672
Client- vs. Server-Side Processing	673
Strategic De-Normalization	674
Routine Maintenance	674
Organizing Your Sprocs Well	675
Keeping Transactions Short	675
Using the Least Restrictive Transaction Isolation Level Possible	675
Implementing Multiple Solutions if Necessary	675
Avoiding Cursors if Possible	676
Uses for Temporary Tables	677
Sometimes, It's the Little Things	677
Hardware Considerations	678
Exclusive Use of the Server	679
I/O vs. CPU Intensive	679
OLTP vs. OLAP	684

Contents

On-Site vs. Off-Site	684
The Risks of Being Down	685
Lost Data	685
Is Performance Everything?	685
Driver Support	686
The Ideal System	686
Troubleshooting	686
The Various Showplans and STATISTICS	687
The Database Consistency Checker (DBCC)	692
The Query Governor	692
The SQL Server Profiler	693
The Performance Monitor (Perfmon)	696
Summary	698
Chapter 24: Administrator	699
Scheduling Jobs	700
Creating an Operator	701
Creating Jobs and Tasks	704
Backup and Recovery	722
Creating a Backup — a.k.a. “A Dump”	722
Recovery Models	728
Recovery	729
Index Maintenance	733
ALTER INDEX	734
Index Name	734
Table or View Name	734
REBUILD	734
DISABLE	735
REORGANIZE	735
Archiving of Data	736
Summary	736
Chapter 25: SMO: SQL Management Objects	739
The History of SQL Server Management Object Models	740
SQL Distributed Management Objects	740
SQL Namespaces	740
Windows Management Instrumentation	741
SMO	741
The SMO Object Model	742

Walking through Some Examples	744
Getting Started	744
Creating a Database	745
Creating Tables	746
Dropping a Database	750
Backing Up a Database	750
Scripting	752
Pulling It All Together	753
Summary	758
Appendix A: System Functions	761
Appendix B: Connectivity	815
Appendix C: Getting Service	825
Index	843

