

Contents

Introduction	xxxv
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	6
The Most Basic Database Object: Table	6
Schemas	7
Filegroups	7
Diagrams	8
Views	8
Stored Procedures	10
User-Defined Functions	10
Users and Roles	10
Rules	11
Defaults	11
User-Defined Data Types	11
Full-Text Catalogs	11
SQL Server Data Types	12
NULL Data	17
SQL Server Identifiers for Objects	17
What Gets Named?	17
Rules for Naming	17
Summary	18
Chapter 2: Tool Time	19
Books Online	20
The SQL Server Configuration Manager	20
Service Management	21
Network Configuration	21
The Protocols	22
On to the Client	24
The SQL Server Management Studio	26

Contents

Getting Started	27
Query Editor	28
SQL Server Business Intelligence Development Studio	33
SQL Server Integration Services (SSIS)	33
Reporting Services	33
Bulk Copy Program (bcp)	34
SQL Server Profiler	34
sqlcmd	34
Summary	35
Chapter 3: Asking a Better Question: Advanced Queries	37
A Quick Review of Subqueries	38
Building a Nested Subquery	39
Nested Queries Using Single Value SELECT Statements	39
Nested Queries Using Subqueries That Return Multiple Values	40
The ANY, SOME, and ALL Operators	41
Correlated Subqueries	42
How Correlated Subqueries Work	42
Correlated Subqueries in the WHERE Clause	42
Correlated Subqueries in the SELECT List	44
Derived Tables	45
The EXISTS Operator	47
Using EXISTS in Other Ways	48
The INTERSECT and EXCEPT Operators	49
EXCEPT	50
INTERSECT	50
Comparing EXCEPT and INTERSECT with Their EXISTS and NOT EXISTS Equivalents	51
Common Table Expressions (CTEs)	54
Recursive Queries	55
MERGE	59
Using External Calls to Perform Complex Actions	62
Performance Considerations	63
JOINS vs. Subqueries vs. ?	63
Summary	65
Chapter 4: XML Integration	67
The XML Data Type	68
Defining a Column as Being of XML Type	69
XML Schema Collections	70
Creating, Altering, and Dropping XML Schema Collections	71

XML Data Type Methods	73
Enforcing Constraints beyond the Schema Collection	79
Retrieving Relational Data in XML Format	80
The FOR XML Clause	80
OPENXML	106
A Quick Heads Up Regarding XML Indexes	110
A Brief Word on Hierarchical Data	110
Summary	111
Chapter 5: Daring to Design	113
Normalization 201	113
Where to Begin	114
Getting to Third Normal Form	115
Other Normal Forms	116
Relationships	116
Diagramming	117
A Couple of Relationship Types	118
The Entity Box	119
The Relationship Line	120
Terminators	120
Logical versus Physical Design	123
Purpose of a Logical Model	123
Parts of a Logical Model	125
Dealing with File-Based Information Via Classic BLOBs	126
Subcategories	129
Types of Subcategories	131
Keeping Track of What's What — Implementing Subcategories	132
Getting Physical — The Physical Implementation of Subcategories	134
Adding to Extensibility with Subcategories	135
Database Reuse	135
Candidates for Reusable Databases	136
How to Break Things Up	136
The High Price of Reusability	137
De-Normalization	137
Partitioning for Scalability	138
The SQL Server Diagramming Tools	139
Tables	141
Dealing with Constraints	143
Regarding Date Columns	146
Summary	146

Contents

Chapter 6: Core Storage and Index Structure	147
SQL Server Storage	148
The Database	148
The File	148
The Extent	149
The Page	149
Rows	152
Full-Text Catalogs	152
File Streams	152
Understanding Indexes	153
To “B,” or Not to “B”: B-Trees	154
How Data Is Accessed in SQL Server	157
Index Types and Index Navigation	158
Creating, Altering, and Dropping Indexes	166
The CREATE INDEX Statement	166
Implied Indexes Created with Constraints	171
ALTER INDEX	172
DROP INDEX	174
Choosing Wisely: Deciding What Index Goes Where and When	174
Selectivity	175
Watching Costs: When Less Is More	175
Choosing That Clustered Index	176
Column Order Matters	178
Dropping Indexes	179
Use the Database Engine Tuning Advisor	179
Maintaining Your Indexes	179
Fragmentation	180
Identifying Fragmentation	180
Summary	187
Chapter 7: More Advanced Index Structures	189
XML Indexes	189
The Primary XML Index	190
Secondary XML Indexes	191
Creating XML Indexes	194
User-Defined Data Types	195
Classic UDTs	195
.NET UDTs	196
Tabular UDTs	197
Dropping a User-Defined Type	199
Hierarchical Data	200

Contents

Understanding Depth Versus Fanout	200
HierarchyID Type Structure	201
Working with HierarchyID Values — HierarchyID Methods	202
Indexing Hierarchy Data	213
Performance Considerations	215
Spatial Data	215
Spatial Concepts	216
Implementing Planar Data Representations — The GEOMETRY Data Type	221
Implementing Geodetic Representations — The GEOGRAPHY Type	227
Filestreams	229
Enabling Filestreaming	231
Enabling a Database for Filestreams	232
Creating a Filestream-Enabled Table	232
Using T-SQL with Filestreams	233
Using Filestreams with .NET	234
Table Compression	235
Enabling Compression	235
Summary	236
Chapter 8: Views	237
Reviewing View Syntax	237
More Complex Views	239
Using a View to Change Data — Before INSTEAD OF Triggers	242
Editing Views with T-SQL	243
Dropping Views	243
Auditing: Displaying Existing Code	244
Protecting Code: Encrypting Views	246
About Schema Binding	248
Making Your View Look Like a Table with VIEW_METADATA	248
Indexed (Materialized) Views	248
Partitioned Views	250
Summary	253
Chapter 9: Scripts and Batches	255
Script Basics	256
Batches	256
Errors in Batches	258
When to Use Batches	259
SQLCMD	262
Dynamic SQL: Generating Your Code on the Fly with the EXEC Command	263
The Gotchas of EXEC	265

Contents

Control-of-Flow Statements	268
The IF . . . ELSE Statement	269
The CASE Statement	273
Looping with the WHILE Statement	276
The WAITFOR Statement	278
TRY/CATCH Blocks	278
Summary	281
Chapter 10: Advanced Programmability	283
A More Advanced Look At Stored Procedures	283
Output Parameters	284
Dealing with Errors	287
Table-Valued Parameters (TVPs)	298
Debugging	301
Starting the Debugger	301
Parts of the Debugger	302
Using the Debugger Once It's Started	306
Understanding the SQLCLR and .NET Programming in SQL Server	307
Assemblies 101	308
Compiling an Assembly	308
Uploading Your Assembly to SQL Server	311
Creating Your Assembly-Based Stored Procedure	312
Creating Scalar User-Defined Functions from Assemblies	314
Creating Table-Valued Functions	317
Creating Aggregate Functions	320
Creating Triggers from Assemblies	325
Custom Data Types	330
Creating Your Data Type from Your Assembly	331
Accessing Your Complex Data Type	331
Dropping Data Types	332
Summary	333
Chapter 11: Transactions and Locks	335
Transactions	335
BEGIN TRAN	337
COMMIT TRAN	337
ROLLBACK TRAN	338
SAVE TRAN	338
How the SQL Server Log Works	343
Using the CHECKPOINT Command	343
At Normal Server Shutdown	344

At a Change of Database Options	345
When the Truncate on Checkpoint Option Is Active	345
When Recovery Time Would Exceed the Recovery Interval Option Setting	345
Failure and Recovery	346
Implicit Transactions	347
Locks and Concurrency	348
What Problems Can Be Prevented by Locks	349
Lockable Resources	352
Lock Escalation and Lock Effects on Performance	352
Lock Modes	353
Lock Compatibility	355
Specifying a Specific Lock Type — Optimizer Hints	356
Setting the Isolation Level	359
READ COMMITTED	360
READ UNCOMMITTED	360
REPEATABLE READ	361
SERIALIZABLE	361
SNAPSHOT	362
Dealing with Deadlocks (a.k.a. “A 1205”)	363
How SQL Server Figures Out There’s a Deadlock	363
How Deadlock Victims Are Chosen	363
Avoiding Deadlocks	363
Summary	366
Chapter 12: Triggers	367
What Is a Trigger?	368
ON	369
WITH ENCRYPTION	369
The FOR AFTER versus the INSTEAD OF Clause	370
WITH APPEND	372
NOT FOR REPLICATION	373
AS	373
Using Triggers for Data Integrity Rules	373
Dealing with Requirements Sourced from Other Tables	374
Using Triggers to Check the Delta of an Update	376
Using Triggers for Custom Error Messages	378
Other Common Uses for Triggers	379
Updating Summary Information	379
Feeding Data into De-normalized Tables for Reporting	379
Setting Condition Flags	380
Other Trigger Issues	383

Contents

Triggers Can Be Nested	383
Triggers Can Be Recursive	383
Debugging Triggers	384
Triggers Don't Get in the Way of Architecture Changes	384
Triggers Can Be Turned Off without Being Removed	385
Trigger Firing Order	385
INSTEAD OF Triggers	387
INSTEAD OF INSERT Triggers	388
INSTEAD OF UPDATE Triggers	391
INSTEAD OF DELETE Triggers	392
IF UPDATE() and COLUMNS_UPDATED()	393
The UPDATE() Function	393
The COLUMNS_UPDATED() Function	394
Performance Considerations	396
Triggers Are Reactive Rather Than Proactive	396
Triggers Don't Have Concurrency Issues with the Process That Fires Them	396
Keep It Short and Sweet	397
Don't Forget Triggers When Choosing Indexes	397
Try Not to Roll Back within Triggers	397
Dropping Triggers	397
Summary	398
Chapter 13: SQL Cursors	399
What Is a Cursor?	399
The Lifespan of a Cursor	400
Types of Cursors and Extended Declaration Syntax	405
Scope	406
Scrollability	411
Cursor Types	413
Concurrency Options	425
Detecting Conversion of Cursor Types: TYPE_WARNING	428
FOR <SELECT>	430
FOR UPDATE	430
Navigating the Cursor: The FETCH Statement	430
Altering Data within Your Cursor	431
Summary	434
Chapter 14: Reporting Services	435
A Quick Look at Reports as a Concept	435
Reporting Services 101	436
Tools Used with Reporting Services	437

Other Means of Accessing Reporting Services	437
Report Server Projects	438
Data Sources	439
Using the Report Wizard	442
Editing Reports	446
Parameterizing Reports	452
Providing/Controlling Parameter Values and How They Are Used	456
Adding Charts	465
Linking Reports	466
Deploying the Report	468
A Brief Note on RDL	469
Summary	471
Chapter 15: Buying in Bulk: The Bulk Copy Program (BCP) and Other Basic Bulk Operations	473
bcp Utility	474
bcp Syntax	474
bcp Import	479
bcp Export	483
Format Files	484
When Your Columns Don't Match	486
Using Format Files	489
Maximizing Import Performance	489
BULK INSERT	490
OPENROWSET (BULK)	491
ROWS_PER_BATCH	492
SINGLE_BLOB, SINGLE_CLOB, SINGLE_NCLOB	492
Summary	492
Chapter 16: Getting Integrated	493
Understanding the Problem	493
An Overview of Packages	494
Tasks	496
The Main Window	499
Solution Explorer	500
The Properties Window	500
Building a Simple Package	500
Executing Packages	508
Using the Execute Package Utility	508
Executing within Management Studio	511
Summary	511

Contents

Chapter 17: Replication	513
Replication Basics	514
Considerations When Planning for Replication	514
Replication Roles	516
Subscriptions	517
Types of Subscribers	518
Filtering Data	518
Replication Models	518
Snapshot Replication	519
Merge Replication	522
Transactional Replication	525
Immediate-Update Subscribers	528
Mixing Replication Types	529
Replication Topology	529
Simple Models	530
Mixed Models	531
Planning for Replication	534
Data Concerns	535
Mobile Devices	535
Setting Up Replication in Management Studio	536
Configuring the Server for Replication	536
Configuring a Publication	540
Setting Up Subscribers (via Management Studio)	546
Using Our Replicated Database	550
Replication Management Objects (RMO)	552
Summary	553
Chapter 18: Looking at Things in Full: Full-Text Search	555
Full-Text Search Architecture	556
Setting Up Full-Text Indexes and Catalogs	558
Enabling Full-Text for Your Database	558
Creating, Altering, Dropping, and Manipulating a Full-Text Catalog	558
Creating, Altering, Dropping, and Manipulating Full-Text Indexes	561
A Note Regarding the Older Syntax	566
More on Index Population	566
Full-Text Query Syntax	568
CONTAINS	568
FREETEXT	570
CONTAINSTABLE	570
FREETEXTTABLE	572
Dealing with Phrases	572

Booleans	573
Proximity	573
Weighting	574
Inflectional	575
Stop Words	575
Summary	576
Chapter 19: Feeling Secure	579
Security Basics	580
One Person, One Login, One Password	580
Password Expiration	582
Password Length and Makeup	583
Number of Tries to Log In	584
Storage of User and Password Information	584
Security Options	585
SQL Server Security	585
Creating and Managing Logins	586
Windows Authentication	593
User Permissions	593
Granting Access to a Specific Database	594
Granting Object Permissions within the Database	595
User Rights and Statement-Level Permissions	601
Server and Database Roles	602
Server Roles	603
Database Roles	604
Application Roles	608
Creating Application Roles	609
Adding Permissions to the Application Role	610
Using the Application Role	610
Getting Rid of Application Roles	611
More Advanced Security	612
What to Do About the Guest Account	612
TCP/IP Port Settings	612
Don't Use the sa Account	613
Keep xp_cmdshell under Wraps	613
Don't Forget Views, Stored Procedures, and UDFs as Security Tools	613
Certificates and Asymmetric Keys	614
Certificates	615
Asymmetric Keys	615
Database Encryption	615
Summary	615

Contents

Chapter 20: A Grand Performance: Designing a Database That Performs Well	617
When to Tune	618
Index Choices	619
Check the Index Tuning Tool in the Database Engine Tuning Advisor	620
Client vs. Server-Side Processing	621
Strategic De-Normalization	622
Organizing Your Sprocs Well	623
Keeping Transactions Short	623
Using the Least Restrictive Transaction Isolation Level Possible	623
Implementing Multiple Solutions if Necessary	623
Avoiding Cursors if Possible	624
Uses for Temporary Tables	625
Using Temp Tables to Break Apart Complex Problems	625
Using Temp Tables to Allow Indexing on Working Data	625
Update Your Code In a Timely Fashion	626
Sometimes, It's the Little Things	626
Hardware Considerations	627
Exclusive Use of the Server	628
I/O vs. CPU Intensive	628
OLTP vs. OLAP	632
On-Site vs. Off-Site	633
The Risks of Being Down	634
Lost Data	635
Is Performance Everything?	635
Driver Support	635
The Ideal System	636
Summary	636
Chapter 21: What Comes After: Forensic Performance Tuning	637
When to Tune (Mark Two)	638
Routine Maintenance	638
Troubleshooting	639
The Data Collector	639
The Various Showplans and STATISTICS	640
Database Console Commands (DBCC)	645
Dynamic Management Views	645
The Activity Monitor	647
The SQL Server Profiler	649
The Performance Monitor (PerfMon)	652
Summary	654

Chapter 22: Administration	655
Scheduling Jobs	656
Creating an Operator	657
Creating Jobs and Tasks	661
Backup and Recovery	681
Creating a Backup — a.k.a. “A Dump”	681
Recovery Models	687
Recovery	688
Index Maintenance	692
ALTER INDEX	693
Index Name	693
Table or View Name	693
REBUILD	693
DISABLE	694
REORGANIZE	694
Archiving of Data	695
PowerShell	695
Trying Out PowerShell	696
Navigating in PowerShell	701
A Final Word on PowerShell	703
Policy-Based Management	703
Summary	703
Chapter 23: SMO: SQL Management Objects	705
The History of SQL Server Management Object Models	706
SQL Distributed Management Objects	706
SQL Namespaces	706
Windows Management Instrumentation	707
SMO	707
The SMO Object Model	708
Walking through Some Examples	709
Getting Started	710
Creating a Database	711
Creating Tables	712
Dropping a Database	716
Backing Up a Database	716
Scripting	718
Pulling It All Together	719
Summary	724

Contents

Chapter 24: Data Warehousing	727
Considering Differing Requirements	728
Online Transaction Processing (OLTP)	728
Online Analytical Processing (OLAP)	729
A Brief Word on Data Mining	730
OLTP or OLAP?	730
Dimensional Databases	731
The Fact Table	731
The Dimension Tables	733
The Star and Snowflake Schemas	733
Data Cubes	733
Data Warehouse Concepts	734
Data Warehouse Characteristics	735
Data Marts	736
SQL Server Integration Services	736
Data Validation	737
Data Scrubbing	737
Creating an Analysis Services Solution	737
Accessing a Cube	746
Summary	749
Chapter 25: Being Well Connected	751
Appendix A: System Functions	753
Legacy System Functions (a.k.a. Global Variables)	754
@@CONNECTIONS	754
@@CPU_BUSY	754
@@IDLE	755
@@IO_BUSY	755
@@PACK_RECEIVED and @@PACK_SENT	755
@@PACKET_ERRORS	755
@@TIMETICKS	755
@@TOTAL_ERRORS	755
@@TOTAL_READ and @@TOTAL_WRITE	756
@@TRANCOUNT	756
Aggregate Functions	756
AVG	757
CHECKSUM	757
CHECKSUM_AGG	758
COUNT	758
COUNT_BIG	758

Contents

GROUPING	759
MAX	759
MIN	759
STDEV	759
STDEVP	760
SUM	760
VAR	760
VARP	760
Configuration Functions	760
@@DATEFIRST	761
@@DBTS	761
@@LANGID and @@LANGUAGE	761
@@LOCK_TIMEOUT	762
@@MAX_CONNECTIONS	762
@@MAX_PRECISION	762
@@NESTLEVEL	762
@@OPTIONS	762
@@REMSERVER	764
@@SERVERNAME	764
@@SERVICENAME	764
@@SPID	764
@@TEXTSIZE	764
@@VERSION	765
Cryptographic Functions	765
AsymKey_ID	765
Cert_ID	766
CertProperty	766
DecryptByAsmKey	766
DecryptByCert	766
DecryptByKey	767
DecryptByPassPhrase	767
EncryptByAsmKey	767
EncryptByCert	767
EncryptByKey	768
EncryptByPassPhrase	768
Key_GUID	768
Key_ID	768
SignByAsmKey	768
SignByCert	769
VerifySignedByAsmKey	769
VerifySignedByCert	769
Cursor Functions	769

Contents

@@CURSOR_ROWS	769
@@FETCH_STATUS	770
CURSOR_STATUS	770
Date and Time Functions	771
CURRENT_TIMESTAMP	772
DATEADD	772
DATEDIFF	772
DATENAME	772
DATEPART	772
DAY	773
GETDATE	773
GETUTCDATE	773
ISDATE	773
MONTH	773
SYSDATETIME	773
SYSDATETIMEOFFSET	774
SYSUTCDATETIME	774
SWITCHOFFSET	774
TODATETIMEOFFSET	775
YEAR	775
Mathematical Functions	775
ABS	776
ACOS	776
ASIN	776
ATAN	776
ATN2	777
CEILING	777
COS	777
COT	777
DEGREES	777
EXP	777
FLOOR	778
LOG	778
LOG10	778
PI	778
POWER	778
RADIANS	778
RAND	778
ROUND	779
SIGN	779
SIN	779
SQRT	779

Contents

SQUARE	780
TAN	780
Basic Metadata Functions	780
COL_LENGTH	781
COL_NAME	781
COLUMNPROPERTY	781
DATABASEPROPERTY	782
DATABASEPROPERTYEX	783
DB_ID	784
DB_NAME	784
FILE_ID	784
FILE_NAME	784
FILEGROUP_ID	784
FILEGROUP_NAME	785
FILEGROUPPROPERTY	785
FILEPROPERTY	785
FULLTEXTCATALOGPROPERTY	786
FULLTEXTSERVICEPROPERTY	786
INDEX_COL	786
INDEXKEY_PROPERTY	787
INDEXPROPERTY	787
OBJECT_ID	787
OBJECT_NAME	788
OBJECTPROPERTY	788
OBJECTPROPERTYEX	790
@@PROCID	791
SCHEMA_ID	791
SCHEMA_NAME	791
SQL_VARIANT_PROPERTY	791
TYPEPROPERTY	792
Rowset Functions	793
CHANGETABLE	793
CONTAINSTABLE	793
FREETEXTTABLE	793
OPENDATASOURCE	793
OPENQUERY	794
OPENROWSET	794
OPENXML	794
Security Functions	795
HAS_DBACCESS	795
IS_MEMBER	795
IS_SRVROLEMEMBER	796

Contents

SUSER_ID	796
SUSER_NAME	796
SUSER_SID	797
SUSER_SNAME	797
USER	797
USER_ID	797
USER_NAME	797
String Functions	798
ASCII	798
CHAR	799
CHARINDEX	799
DIFFERENCE	799
LEFT	799
LEN	799
LOWER	800
LTRIM	800
NCHAR	800
PATINDEX	800
QUOTENAME	800
REPLACE	801
REPLICATE	801
REVERSE	801
RIGHT	801
RTRIM	801
SOUNDEX	802
SPACE	802
STR	802
STUFF	802
SUBSTRING	802
UNICODE	803
UPPER	803
System Functions	803
APP_NAME	804
CASE	804
CAST and CONVERT	805
COALESCE	805
COLLATIONPROPERTY	805
CURRENT_USER	806
DATALength	806
@@ERROR	806
FORMATMESSAGE	806
GETANSINULL	807

Contents

HOST_ID	807
HOST_NAME	807
IDENT_CURRENT	807
IDENT_INCR	807
IDENT_SEED	808
@@IDENTITY	808
IDENTITY	808
ISNULL	808
ISNUMERIC	809
NEWID	809
NULLIF	809
PARSENAME	809
PERMISSIONS	809
@@ROWCOUNT	809
ROWCOUNT_BIG	810
SCOPE_IDENTITY	810
SERVERPROPERTY	810
SESSION_USER	812
SESSIONPROPERTY	812
STATS_DATE	812
SYSTEM_USER	812
USER_NAME	813
Text and Image Functions	813
TEXTPTR	813
TEXTVALID	813
Appendix B: Going Meta: Getting Data About Your Data	815
System Views	815
sys.assemblies	816
sys.columns	817
sys.databases	818
sys.database_files	819
sys.identity_columns	821
sys.indexes	822
sys.index_columns	824
sys.objects	825
sys.partitions	826
sys.partition_functions	827
sys.schemas	828
sys.servers	828
sys.spatial_indexes	829
sys.synonyms	830

Contents

sys.user_token	830
sys.xml_indexes	831
Dynamic Management Views	832
Index-Related dm Views and Functions	832
Appendix C: The Basics	841
Everything to Do with Queries	841
The Basic SELECT Statement	842
Beyond Inner Joins	846
The INSERT Statement	848
The INSERT INTO . . . SELECT Statement	850
The UPDATE Statement	851
The DELETE Statement	852
Exploring Alternative Syntax for Joins	853
An Alternative Inner Join	853
An Alternative Outer Join	854
An Alternative CROSS JOIN	855
UNION	855
Subqueries and Derived Tables	858
Summary	860
Index	861