

Microsoft®

Microsoft®

SQL Server® 2008

Mike Hotek




eBook + exercises



Step by Step



How to access your CD files



The print edition of this book includes a CD. To access the CD files, go to <http://aka.ms/626041/files>, and look for the Downloads tab.

Note: Use a desktop web browser, as files may not be accessible from all ereader devices.

Questions? Please contact: mspinput@microsoft.com

Microsoft Press

PUBLISHED BY

Microsoft Press

A Division of Microsoft Corporation

One Microsoft Way

Redmond, Washington 98052-6399

Copyright © 2009 by Mike Hotek

All rights reserved. No part of the contents of this book may be reproduced or transmitted in any form or by any means without the written permission of the publisher.

Library of Congress Control Number: 2008935428

Printed and bound in the United States of America.

ISBN: 978-0-7356-2604-1

5 6 7 8 9 10 11 12 13 QGT 7 6 5 4 3 2 1

Distributed in Canada by H.B. Fenn and Company Ltd.

A CIP catalogue record for this book is available from the British Library.

Microsoft Press books are available through booksellers and distributors worldwide. For further information about international editions, contact your local Microsoft Corporation office or contact Microsoft Press International directly at fax (425) 936-7329. Visit our Web site at www.microsoft.com/mspress. Send comments to mspinput@microsoft.com.

Microsoft, Microsoft Press, Active Directory, ActiveX, Excel, Internet Explorer, MSDN, PowerPoint, SQL Server, Visio, Visual Basic, Visual C#, Visual Studio, Windows, Windows Server, and Windows Vista are either registered trademarks or trademarks of the Microsoft group of companies. Other product and company names mentioned herein may be the trademarks of their respective owners.

The example companies, organizations, products, domain names, e-mail addresses, logos, people, places, and events depicted herein are fictitious. No association with any real company, organization, product, domain name, e-mail address, logo, person, place, or event is intended or should be inferred.

This book expresses the author's views and opinions. The information contained in this book is provided without any express, statutory, or implied warranties. Neither the authors, Microsoft Corporation, nor its resellers, or distributors will be held liable for any damages caused or alleged to be caused either directly or indirectly by this book.

Acquisitions Editor: Ken Jones

Developmental Editor: Sally Stickney

Project Editor: Denise Bankaitis

Editorial Production: S4Carlisle Publishing Services

Technical Reviewer: Randall Galloway; Technical Review services provided by Content Master,
a member of CM Group, Ltd.

Cover: Tom Draper Design

To Genilyn,

You make it all worthwhile.

Contents at a Glance

Part I	Getting Started with Microsoft SQL Server 2008	
1	Overview of Microsoft SQL Server	3
2	Installing and Configuring SQL Server 2008	15
3	Using the Tools in SQL Server 2008	31
Part II	Designing Databases	
4	Creating Databases	51
5	Designing Tables	63
6	Indexes	87
7	Partitioning	103
Part III	Retrieving and Manipulating Data	
8	Data Retrieval	121
9	Advanced Data Retrieval	137
10	Data Manipulation	153
Part IV	Designing Advanced Database Objects	
11	Views	179
12	Stored Procedures	185
13	Functions	199
14	Triggers	213
15	Database Snapshots	219
16	Service Broker	225
17	Full-Text Indexing	247
Part V	Database Management	
18	Security	259
19	Policy-Based Management	289
20	Data Recovery	301

21	SQL Server Agent	319
22	Dynamic Management Views	333
Part VI	High Availability Overview	
23	High Availability	345
Part VII	Business Intelligence	
24	SQL Server Integration Services	377
25	SQL Server Reporting Services	413
26	SQL Server Analysis Services	451

Table of Contents

Acknowledgments	xvii
Introduction	xix

Part I **Getting Started with Microsoft SQL Server 2008**

1 Overview of Microsoft SQL Server	3
Database Engine	3
Storage Engine	4
Security Subsystem	4
Programming Interfaces	5
Service Broker	6
SQL Server Agent	6
Replication	6
High Availability	7
The Relational Engine in SQL Server 2008	8
Business Intelligence	9
Integration Services	10
Reporting Services	12
Analysis Services	13
Chapter 1 Quick Reference	14
2 Installing and Configuring SQL Server 2008	15
Editions of SQL Server 2008	15
Infrastructure Requirements	17
Service Accounts	17
Collation Sequences	19
Authentication Modes	19
SQL Server Instances	20
Upgrading to SQL Server 2008	20
In-Place Upgrade	20
Side-by-Side Upgrade	21

 **What do you think of this book? We want to hear from you!**

Microsoft is interested in hearing your feedback so we can continually improve our books and learning resources for you. To participate in a brief online survey, please visit:

www.microsoft.com/learning/booksurvey/

Installing SQL Server 2008	22
Install Sample Databases	28
Chapter 2 Quick Reference.....	29
3 Using the Tools in SQL Server 2008	31
SQL Server Documentation	31
Management Tools in SQL Server 2008	34
SQL Server Configuration Manager	35
SQL Server Management Studio	38
Database Mail	42
Performance Management Tools	44
Profiler	44
Database Engine Tuning Advisor	44
Performance Studio	45
Business Intelligence Tools	46
Business Intelligence Development Studio	46
Chapter 3 Quick Reference.....	47

Part II **Designing Databases**

4 Creating Databases.....	51
SQL Server System Databases	51
SQL Server Database Structure	53
Database Files	53
Filegroups	56
Creating a Database.....	57
Moving Databases	59
Detaching a Database.....	59
Attaching a Database	60
Chapter 4 Quick Reference.....	62
5 Designing Tables.....	63
Naming Objects.....	64
Schemas.....	64
Data Types.....	67
Numeric Data	67
Character Data	69
Date and Time Data.....	71
Binary Data	72
XML	72

FILESTREAM Data.....	74
Spatial Data Type.....	74
HierarchyID Data Type.....	75
Column Properties.....	75
Creating Tables.....	76
Computed Columns.....	77
Sparse Columns.....	78
Constraints.....	79
Primary Keys.....	79
Unique Constraints.....	81
Check Constraints.....	82
Default Constraints.....	82
Foreign Keys.....	83
Database Diagrams.....	85
Chapter 5 Quick Reference.....	86
6 Indexes.....	87
Index Structure.....	87
Balanced Trees (B-trees).....	88
Index levels.....	89
Clustered Indexes.....	91
Nonclustered Indexes.....	93
Index Maintenance.....	93
Included Columns.....	95
Covering Indexes.....	95
Filtered Indexes.....	97
Online Index Creation.....	98
Index Management and Maintenance.....	99
Index Fragmentation.....	99
Fill Factor.....	99
Defragmenting an Index.....	100
Disabling an index.....	100
XML Indexes.....	100
Spatial Indexes.....	102
Chapter 6 Quick Reference.....	102
7 Partitioning.....	103
Partition Functions.....	103
Partition Schemes.....	105
Filegroups.....	105

Partitioning Tables and Indexes	106
Partial Backup and Restore	107
Creating a Partitioned Index	107
Managing Partitions	110
SPLIT and MERGE Operators	110
Altering a Partition Scheme	110
Index Alignment	110
SWITCH Operator	111
Chapter 7 Quick Reference	117

Part III Retrieving and Manipulating Data

8 Data Retrieval	121
General SELECT Statement	121
Sorting Results	126
Filtering Data	127
Retrieving from More Than One Table	131
Retrieving Unique Results	134
Chapter 8 Quick Reference	135
9 Advanced Data Retrieval	137
General SELECT Statement	137
Aggregating Data	138
Aggregating Multiple Permutations	140
Filtering Aggregates	143
Running Aggregates	144
Calculating Pivot Tables	145
Ranking Data	146
Aggregating Result Sets	147
Common Table Expressions	149
Querying XML Data	150
Chapter 9 Quick Reference	152
10 Data Manipulation	153
Inserting Data	153
INSERT	153
SELECT INTO	159
Updating Data	159
Deleting Data	161
TRUNCATE	164

MERGE Statement	164
OUTPUT Clause	166
Transaction Handling	167
Tracking Changes	171
Change Tracking	171
Change Data Capture	173
Chapter 10 Quick Reference	176

Part IV Designing Advanced Database Objects

11 Views	179
Creating a View	179
Query Substitution	180
Modifying Data Through a View	181
Creating an Indexed View	182
Query Substitution	184
Chapter 11 Quick Reference	184
12 Stored Procedures	185
Creating Stored Procedures	185
Commenting Code	186
Variables, Parameters, and Return Codes	186
Variables	186
Parameters	188
Return Codes	188
Executing Stored Procedures	189
Control Flow Constructs	189
Error Handling	192
Dynamic Execution	193
Cursors	194
CLR Procedures	196
Building an Administrative Procedure	196
Chapter 12 Quick Reference	198
13 Functions	199
System Functions	199
Creating a Function	200
Retrieving Data from a Function	204
Chapter 13 Quick Reference	212

14	Triggers	213
	DML Triggers	213
	DDL Triggers	215
	Chapter 14 Quick Reference	218
15	Database Snapshots	219
	Creating a Database Snapshot	219
	Copy-On-Write Technology	220
	Reverting Data Using a Database Snapshot	222
	Chapter 15 Quick Reference	223
16	Service Broker	225
	Service Broker Architecture	226
	Messaging Overview	226
	Service Broker Components	227
	Application Interaction	228
	Message Types and Contracts	229
	Message Types	229
	Contracts	232
	Queues and Services	233
	Queues	233
	Services	234
	Conversations	235
	Sending and Receiving Messages	236
	Sending Messages	236
	Receiving Messages	237
	Queue Activation	242
	Prioritization	245
	Chapter 16 Quick Reference	246
17	Full-Text Indexing	247
	Full-Text Catalogs	247
	Full-Text Indexes	249
	Querying Full-Text Data	251
	FREETEXT	252
	CONTAINS	252
	Chapter 17 Quick Reference	255

Part V Database Management

18	Security	259
	Configuring the Attack Surface	259
	Endpoints	260
	Endpoint Types and Payloads	261
	Endpoint Access	261
	TCP Endpoints	262
	Principals, Securables, and Permissions	264
	Principals	264
	Impersonation	273
	Securables	273
	Permissions	274
	Ownership Chains	277
	CLR Security	279
	Data Encryption	280
	Master Keys	281
	Hash Algorithms	282
	Symmetric Keys	284
	Certificates and Asymmetric Keys	285
	Transparent Data Encryption	286
	Encryption Key Management	287
	Chapter 18 Quick Reference	288
19	Policy-Based Management	289
	Overview of Policy-Based Management	289
	Facets	290
	Conditions	290
	Policy Targets	293
	Policies	295
	Policy Categories	297
	Policy Compliance	298
	Chapter 19 Quick Reference	300
20	Data Recovery	301
	Database Backups	301
	Backup Types	302
	Page Corruption	308
	Recovery Models	309
	Database Restores	311

	Restoring a Full Backup.	311
	Restoring a Differential Backup.	314
	Restoring a Transaction Log Backup.	315
	Chapter 20 Quick Reference.	317
21	SQL Server Agent	319
	Creating Jobs.	319
	Jobs Steps	319
	Job Schedules	320
	Operators	321
	Creating Maintenance Plans.	325
	Creating Alerts	329
	Chapter 21 Quick Reference.	332
22	Dynamic Management Views	333
	Overview of DMVs	333
	Retrieving Object Metadata.	334
	Database Diagnostics	336
	Object Size	336
	Indexes.	337
	Query Execution Statistics	339
	Chapter 22 Quick Reference.	341
 Part VI High Availability Overview		
23	High Availability	345
	Failover Clustering	345
	Failover Cluster Instance Components.	345
	Network Configuration.	346
	Disk Configuration.	346
	Security Configuration	346
	Health Checks.	347
	Cluster Failover.	348
	Database Mirroring	348
	Database Mirroring Roles.	349
	Database Mirroring Endpoints.	350
	Operating Modes.	351
	Caching	353
	Transparent Client Redirect	354
	Corrupt Pages.	354

Database Snapshots	354
Initializing Database Mirroring	355
Log Shipping	356
Log Shipping Components	356
Log Shipping Initialization	357
Replication	363
Replication Components	364
Replication Roles	365
Replication Agents	365
Replication Methods	366
Chapter 23 Quick Reference	373

Part VII Business Intelligence

24 SQL Server Integration Services	377
BIDS Overview	378
Tasks	380
Transforms	382
Building a Package	385
Connections	386
Control Flow	390
Data Flow	394
Data Conversion	401
Exception Handling	403
Configuration	406
Deployment	409
Chapter 24 Quick Reference	411
25 SQL Server Reporting Services	413
Configuring Reporting Services	413
Reporting Services Web Site	417
Creating Reports	418
Building a Report	418
Formatting	423
Computations	429
Interactive Elements	431
Parameters	434
Deploy Reports	439
Report Subscriptions	441
Linked Reports	443

Report Caching and Snapshots	444
Chapter 25 Quick Reference	449
26 SQL Server Analysis Services	451
Data Warehousing Overview	451
Online Analytic Processing (OLAP)	452
Dimensional Model	453
Cubes	454
Dimensions, Measures, and Calculations	462
Hierarchies	467
KPIs, Partitions, Perspectives, and Translations	470
Key Performance Indicators (KPIs)	470
Partitions	471
Perspectives	471
Translations	471
Data Mining	472
Algorithms	473
Mining Models and Mining Structures	474
Data Mining Demystified	483
Chapter 26 Quick Reference	485
Index	487

What do you think of this book? We want to hear from you!

Microsoft is interested in hearing your feedback so we can continually improve our books and learning resources for you. To participate in a brief online survey, please visit:

www.microsoft.com/learning/booksurvey/

Acknowledgments

Thank you to all of my readers over the past decade or so; it's hard to believe that this will be the eighth book I've written and it would not be possible without you. I'd like to thank the talented and incredibly patient editorial team at Microsoft Press – Denise Bankaitis and Sally Stickney. I would especially like to thank Ken Jones, who has now gone through four books with me and has an exceptional talent for keeping things moving smoothly through all of the various trials that come up during the authoring and editing processes. My words and thoughts would probably be an unintelligible mess without the help of Randall and Christian, who not only smoothed out the rough edges, but made sense out of many 3 A.M. missives that likely had them questioning my sanity.

Introduction

Microsoft SQL Server has been Microsoft's flagship database for over 15 years. Before the next version of SQL Server launches, SQL Server 2008 will be celebrating SQL Server's 20th birthday. In that time, SQL Server has grown from handling small departmental tasks to serving up the largest databases on the planet. The release of SQL Server 2000 saw a dramatic evolution of Microsoft SQL Server. No longer a simple "database," Microsoft SQL Server is now a complete data architecture solution capable of handling the data storage and manipulation needs of any organization.

SQL Server 2005 extended the data platform with dramatic new capabilities in programming, .NET integration, high availability, management instrumentation, and business intelligence. So dramatic were the feature enhancements, that while retaining the same "look and feel" of Microsoft SQL Server, Microsoft essentially released an entirely new data platform.

SQL Server 2008 picks up the rapid innovation by enhancing hundreds of existing features while at the same time adding hundreds more. SQL Server 2008 also presents a highly scalable, highly adaptable data architecture platform against which you can build any conceivable application. Yet with these new and improved features come challenges for IT professionals. I know from nearly two decades of experience working with and teaching Microsoft SQL Server in organizations of all sizes spanning all industries that if users don't understand how to use the product effectively, they and their organization won't be able to get the full benefits of this powerful product. At the same time, the role and skill set of a DBA is changing rapidly. While it might be acceptable for a consultant to focus on a very narrow area such as performance tuning, the storage engine, or writing reports, most companies are beginning to insist that their DBAs know how to architect, code, and manage solutions that utilize every feature that ships with SQL Server – from building a table through building a data mining model.

I wrote this book because I wanted to provide the first comprehensive tour of the entire feature set available within Microsoft SQL Server, beginning with the relational databases that lie at the core of every organization, through the unique management capabilities, and finishing with a set of extraordinarily powerful analysis platforms which comprise the core of the Microsoft business intelligence platform. Armed with this information, you will be able to:

- Architect, secure, and manage relational databases
- Retrieve and manipulate data
- Expand your application's capabilities with programmable objects
- Secure and recover your business data
- Ensure that your database platform performs well and is resilient
- Transform your business data into actionable business intelligence

Who This Book Is For

The aim of this book is to teach you the fundamentals of the SQL Server 2008 data platform. SQL Server contains features that appeal to anyone involved with the storage or manipulation of data within an organization.

This book is intended for the IT professional who is either new to SQL Server or new to SQL Server 2008. Experienced SQL Server professionals will still find a significant amount of information that is applicable to their jobs.

No book can possibly include all of the intricacies of SQL Server 2008. Instead, the focus is on providing an overview of each feature in sufficient depth to allow you to build SQL Server applications. As you progress through this book, you will learn how to install each of the components, configure and manage instances, and build databases. You will walk through each of the client tools that ship with SQL Server, and I'll explain how each tool enables you to develop and manage your database environment. You will learn how to manipulate data, secure your databases, manage and protect your data, and distribute your data platform to make it more scalable, redundant, and fault tolerant. You will learn how to integrate your SQL Server data with a variety of sources, build reports to serve your lines of business, and finally feed all of this data into powerful analysis and data mining systems to deliver actionable information to your lines of business in near real-time.

How This Book Is Organized

This book is organized in seven parts along with three online articles that allow you to focus on specific portions within the SQL Server platform as well as specific job functions as follows:

- Part One describes the core components available within SQL Server along with instructions on how to install each component. You will also get an overview of each tool that ships with SQL Server that we will be using throughout the book.
- Part Two shows you how to create and configure databases to provide the foundation for the table and index structures that form the backbone of every database application you will create.
- Part Three teaches you how to manipulate and retrieve data.
- Part Four introduces you to the fundamental programming structures, views, stored procedures, functions, and triggers available.
- Part Five explains how to secure, manage, backup, and recover databases.
- Part Six will provide an introduction to the high availability technologies which ship with SQL Server such as clustering, database mirroring, and log shipping.
- Part Seven covers the three business intelligence technologies – Integration Services, Reporting Services, and Analysis Services.

- Bonus material is provided online in three articles, “Performance Analysis and Tuning”, “Performance and Data Capture Tools”, and “Performance Analysis Tools” on the Microsoft Press Online Windows Server and Client Web site at www.microsoft.com/learning/books/online/serverclient

Finding Your Best Starting Point in This Book

This book is designed to help you build skills in a number of essential areas. You can use this book if you are new to SQL Server or if you are switching from another database system. Use the following table to find your best starting point.

If you are a(n)	Follow these steps
Database administrator, database architect, database developer, database engineer, or data analyst	<ol style="list-style-type: none"> 1. Install the practice files as described in the next section, “Installing and Using the Practice Files.” 2. Work through the chapters in Parts One through Four sequentially. 3. Complete Parts Five through Seven along with the three online articles as your level of experience and interest dictates.
Application developer	<ol style="list-style-type: none"> 1. Install the practice files as described in the next section, “Installing and Using the Practice Files.” 2. Skim the chapters in Part One to get an overview of installing SQL Server and the tools available, and then concentrate on the chapters in Parts Two through Four. 3. Complete Parts Five through Seven along with the three online articles as your level of experience and interest dictates.
System administrator, network administrator, or security administrator	<ol style="list-style-type: none"> 1. Install the practice files as described in the next section, “Installing and Using the Practice Files.” 2. Work through the chapters in Part One. 3. Skim the chapters in Parts Two through Four. 4. Work through the chapters in Part Five. 5. Complete Parts Six and Seven along with the three online articles as your level of experience and interest dictates.
Business analyst	<ol style="list-style-type: none"> 1. Install the practice files as described in the next section, “Installing and Using the Practice Files.” 2. Skim through the chapters in Parts One through Six. 3. Work through the chapters in Part Seven as your level of experience and interest dictates. 4. Complete the three online articles as your level of experience and interest dictates.
IT Management	<ol style="list-style-type: none"> 1. Work through Chapter 1. 2. Skim through the rest of the chapters and online articles as your experience and level of interest dictates.
Referencing the book after working through the exercises	<ol style="list-style-type: none"> 1. Use the Index or the Table of Contents to find information about particular subjects. 2. Read the Quick Reference sections at the end of each chapter to find a brief review of the syntax and techniques presented in the chapter.

Conventions and Features in This Book

This book presents information using conventions designed to make the information readable and easy to follow. Before you start, read the following list, which explains conventions you'll see throughout the book and points out helpful features that you might want to use.

Conventions

- Each exercise is a series of tasks. Each task is presented as a series of numbered steps (1, 2, and so on). A round bullet (•) indicates an exercise that has only one step.
- Notes labeled “Tip” provide additional information or alternative methods for completing a step successfully.
- Notes labeled “Important” alert you to information you need to check before continuing.
- Text that you type appears in bold.
- A plus sign (+) between two key names means that you must press those keys at the same time. For example, “Press Alt+Tab” means that you hold down the Alt key while you press the Tab key.

Other Features

- Sidebars throughout the book provide more in-depth information about the exercise. The sidebars might contain background information, design tips, or features related to the information being discussed.
- Each chapter ends with a Quick Reference section. The Quick Reference section contains quick reminders of how to perform the tasks you learned in the chapter.

System Requirements

You'll need the following hardware and software to complete the practice exercises in this book:

- Microsoft Windows Vista Home Basic Edition or higher, Windows Server 2008 Standard edition or higher, Windows Server 2003 SP2 or higher, or Windows XP Professional SP2 or higher.



Note SQL Server 2008 is not supported on Windows Server 2008 Server Core edition.

- Microsoft SQL Server 2008 Evaluation edition, SQL Server 2008 Developer edition, or SQL Server 2008 Enterprise edition.



Note You can use other editions of SQL Server 2008, however, you will be limited by the feature set supported by the SQL Server edition that you have installed.

- 2.0 GHz Pentium III+ processor, or faster
- 1 GB of available, physical RAM
- 2GB of available disk space
- Video (800 × 600 or higher resolution) monitor with at least 256 colors
- CD-ROM or DVD-ROM drive
- Microsoft mouse or compatible pointing device

You will also need to have Administrator access to your computer to configure SQL Server 2008.

Sample Databases

All of the examples within this book utilize the AdventureWorks and AdventureWorksDW sample databases. Sample databases no longer ship with SQL Server and must be downloaded from the CodePlex website at <http://www.codeplex.com/SQLServerSamples>.



Tip In addition to the sample databases, the CodePlex site contains dozens of examples, sample applications, and add-ons that can greatly enhance your SQL Server experience.

Code Samples

The companion CD inside this book contains the code samples that you'll use as you perform the exercises. By using the code samples, you won't waste time creating files that aren't relevant to the exercise. The files and the step-by-step instructions in the lessons also let you learn by doing, which is an easy and effective way to acquire and remember new skills.

Digital Content for Digital Book Readers: If you bought a digital-only edition of this book, you can enjoy select content from the print edition's companion CD. Visit <http://www.microsoftpressstore.com/title/9780735626041> to get your downloadable content. This content is always up-to-date and available to all readers.

Installing the Code Samples

Follow these steps to install the code samples and required software on your computer so that you can use them with the exercises.

1. Remove the companion CD from the package inside this book and insert it into your CD-ROM drive.



Note An end-user license agreement should open automatically. If this agreement does not appear, open My Computer on the desktop or Start menu, double-click the icon for your CD-ROM drive, and then double-click StartCD.exe.

2. Review the end-user license agreement. If you accept the terms, select the accept option and then click Next.

A menu will appear with options related to the book.

3. Click Install Code Samples.
4. Follow the instructions that appear.

The code samples are installed to the following location on your computer:

Documents\Microsoft Press\SQL Server 2008 Step By Step.

Using the Code Samples

Each chapter in this book explains when and how to use any code samples for that chapter. When it's time to use a code sample, the book will list the instructions for how to open the files.

Uninstalling the Code Samples

Follow these steps to remove the code samples from your computer.

1. In Control Panel, open Add Or Remove Programs if running Windows XP or Programs\Uninstall A Program if running Windows Vista.
2. From the list of Currently Installed Programs in Windows XP, or from the list of Uninstall Or Change A Program in Windows Vista, select <Microsoft SQL Server 2008 Step by Step>.
3. Click Remove in Windows XP or click Uninstall/Change in Windows Vista.
4. Follow the instructions that appear to remove the code samples.

Find Additional Content Online

As new or updated material becomes available that complements your book, it will be posted online on the Microsoft Press Online Developer Tools Web site. The type of material you might find includes updates to book content, articles, links to companion content, errata, sample chapters, and more. This Web site is available at www.microsoft.com/learning/books/online/serverclient, and is updated periodically.

Support for This Book

Every effort has been made to ensure the accuracy of this book and the contents of the companion CD. As corrections or changes are collected, they will be added to a Microsoft Knowledge Base article.

Microsoft Press provides support for books and companion CDs at the following Web site:

<http://www.microsoft.com/learning/support/books/default.aspx>.

Questions and Comments

If you have comments, questions, or ideas regarding the book or the companion CD, or questions that are not answered by visiting the sites above, please send them to Microsoft Press via e-mail to:

mspinput@microsoft.com.

Or via postal mail to:

Microsoft Press
Attn: *Programming Microsoft SQL Server 2008 Step by Step* Series Editor
One Microsoft Way
Redmond, WA 98052-6399.

Please note that Microsoft software product support is not offered through the above addresses.

Chapter 3

Using the Tools in SQL Server 2008

After completing this chapter, you will be able to

- Select the appropriate SQL Server 2008 tool for a given task
- Manage SQL Server 2008 services
- Launch, navigate, and utilize SQL Server Management Studio
- Utilize several shortcuts for better productivity
- Configure Database Mail

SQL Server 2008 ships with eight stand-alone tools used to configure, manage, and monitor SQL Server services. Within this group of eight core tools, you can also design SQL Server objects and execute code. The most wide-reaching tool, SQL Server Management Studio (SSMS), contains four additional tools designed for management and monitoring.



Note You can install instances of SQL Server Integration Services, Reporting Services, Analysis Services, and the Database Engine. To simplify the terminology, we will simply refer to a Database Engine instance as a SQL Server instance. All other instances will be referred to as either an SSIS, SSAS, or SSRS instance.

SQL Server Documentation

SQL Server 2008 ships with a very comprehensive Books Online. While many do not consider documentation as a “tool,” the saying “information is power” immediately comes to mind. There is a reason that you will hear someone say: “Read the manual.” Undeservedly, Books Online has received a very bad reputation.

Books Online should be your primary source for information concerning SQL Server 2008, after this book of course. Books Online contains detailed explanations of every feature within SQL Server, syntax on every command, and thousands of code samples that you can apply. Additionally, SQL Server Books Online integrates a vast array of online content into the local documentation in order to provide extensive, constantly updated information that can be applied within your environment. While we will not spend a significant amount of time going through Books Online, we will explain a couple of very useful and often-overlooked features.

Hidden within Books Online, underneath the How Do I Link button, lurks a set of several dozen comprehensive tutorials that walk you through important feature sets such as policy-based management, *hierarchyID* data types, designing OLAP cubes, implementing Data mining models, deploying replication, and building SSIS packages or SSRS reports, as shown in Figure 3-1.

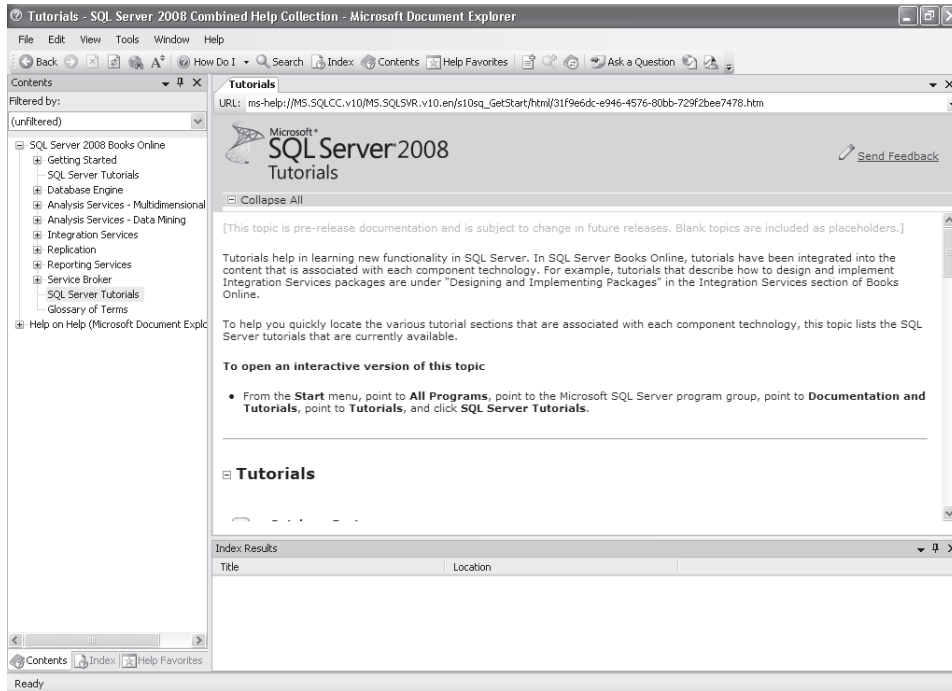


FIGURE 3-1 SQL Server 2008 tutorials

You can also see a set of three links at the far right of the toolbar in Figure 3-1. The first of these three links, Ask A Question, will launch a browser window into the right-hand pane and take you to the Microsoft Developer Network (MSDN) forums. The MSDN forums allow you to ask any question about SQL Server for which you cannot find an answer. Questions are answered by volunteers who include thousands of SQL Server professionals around the world as well as hundreds of members of the SQL Server development team.

The second-to-last link will again take you to the MSDN forums, but will automatically apply a filter to display only those questions you are participating in so that you can easily follow up on the status of your questions.

The last link will launch a browser window in the right-hand pane and take you to Microsoft Connect where you can post feedback on SQL Server or a particular feature

as well as publish bug reports. The bug reports posted are continuously reviewed by the SQL Server development team so that they can proactively supply patches to functionality. The product feedback plays an important role during the planning of the next version of SQL Server.

You can bookmark topics that you reference frequently by using the Help Favorites feature. The Help Favorites feature allows you to also save frequently executed searches for later recall.

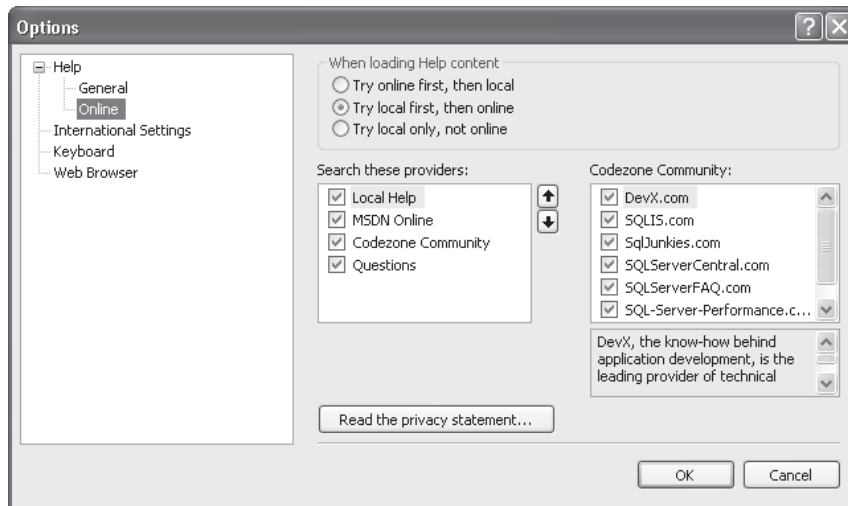


Note Because SQL Server Books Online contains features to integrate the local version of help with various Web sites, you can experience some performance issues. To mitigate any performance issues, you should set Books Online to look at the local help first before trying online resources.

In this procedure, you will configure Books Online for optimal performance.

Configure Books Online

1. Launch Books Online by selecting Start | All Programs | Microsoft SQL Server 2008 | Documentation And Tutorials | SQL Server Books Online.
2. Select Tools | Options.
3. Select Online and configure as shown here.



Management Tools in SQL Server 2008

SQL Server 2008 ships with a set of nine tools for managing SQL Server instances and interacting with data: OSQL, SQLCMD, Tablediff, Bulk Copy Program (BCP), SQLDiag, Resource Governor, SQL Server Configuration Manager, SSMS, and Database Mail.

OSQL is a command line utility which was added to SQL Server 2000 as a replacement for ISQL. OSQL allows you to connect to and execute queries against a SQL Server instance without requiring the overhead of a graphical interface.



Important OSQL has been deprecated as of SQL Server 2005. While OSQL is still available in SQL Server 2008, you should rewrite any OSQL routines to utilize SQLCMD.

SQL Server 2005 introduced SQLCMD as the command line query interface that replaced OSQL. While OSQL allowed you to submit interactive queries from a command line along with very limited automation capabilities, SQLCMD provides a rich automation interface complete with variable substitution and dynamic code creation/execution.



Note A discussion of the specifics of SQLCMD is beyond the scope of this book. For details on SQLCMD, please see the Books Online topic "SQLCMD Utility."

You can use Tablediff.exe to compare the data between two tables. Tablediff can be run to alert you if the data or structure of two tables is different. Additionally, Tablediff can generate a script file containing the statements necessary to bring the destination table into synchronization with the source table. Tablediff is primarily used within replication architectures.

The BCP utility is the oldest utility within the SQL Server product, dating all the way back to the very first version of SQL Server. BCP has been enhanced with each successive version to handle new data types and named instances, but the speed and feature set has not changed. BCP is used to export data from a table to a file as well as import data from a file into a table. If your import and export needs are reasonably simple and straightforward, BCP should be the only utility that you would need. For more advanced import and export capabilities, you should utilize SQL Server Integration Services (SSIS).



Note You will learn about BCP and its import cousin, BULK INSERT, in Chapter 10, "Data Manipulation." You will learn about SSIS in Chapter 24, "Business Intelligence."

SQLDiag is a utility that collects diagnostic information about a SQL Server instance. SQLDiag is designed to capture Windows performance counters, event logs, SQL Server Profiler traces, SQL Server blocking, and SQL Server configuration information. Primarily used as a data

collection engine for Microsoft Customer Service and Support (CSS) to troubleshoot SQL Server issues, the data collected can also be used by a DBA to analyze SQL Server performance and stability issues.



Note The details of SQLDiag are beyond the scope of this book. Please refer to the Books Online article “SQLDiag Utility” for more information.

Resource Governor is a new feature in SQL Server 2005 that is found within SSMS. The purpose of Resource Governor is to allow a DBA to configure rules around resource allocation such as processor or memory that is then applied to specific queries, users, or groups of users. The goal of Resource Governor is to allow high-priority workloads to take priority over other workloads in order to provide the best response based on user expectations.



Note You will learn about Resource Governor in the article, “Performance and Data Capture Tools,” which can be found on the Microsoft Press Online Windows Server and Client Web site at www.microsoft.com/learning/books/online/serverclient.

SQL Server Configuration Manager

Shown in Figure 3-2 on page 36, SQL Server Configuration Manager is responsible for managing SQL Server services and protocols. The primary tasks that you will perform with SQL Server Configuration Manager are:

- Start/Stop/Pause/Restart a service
- Change service accounts and passwords
- Manage the startup mode of a service
- Configure service startup parameters

Once you have completed the initial installation and configuration of your SQL Server services, the primary action that you will perform within SQL Server Configuration Manager is to periodically change service account passwords. When changing service account passwords, you no longer have to restart the SQL Server instance for the new credential settings to take effect.



Important Windows Service Control Applet also has entries for SQL Server services and allows you to change service accounts and passwords. You should never change service accounts or service account passwords using the Windows Service Control Applet. SQL Server Configuration Manager needs to be used, because SQL Server Configuration Manager includes the code to regenerate the service master key that is critical to the operation of SQL Server services.

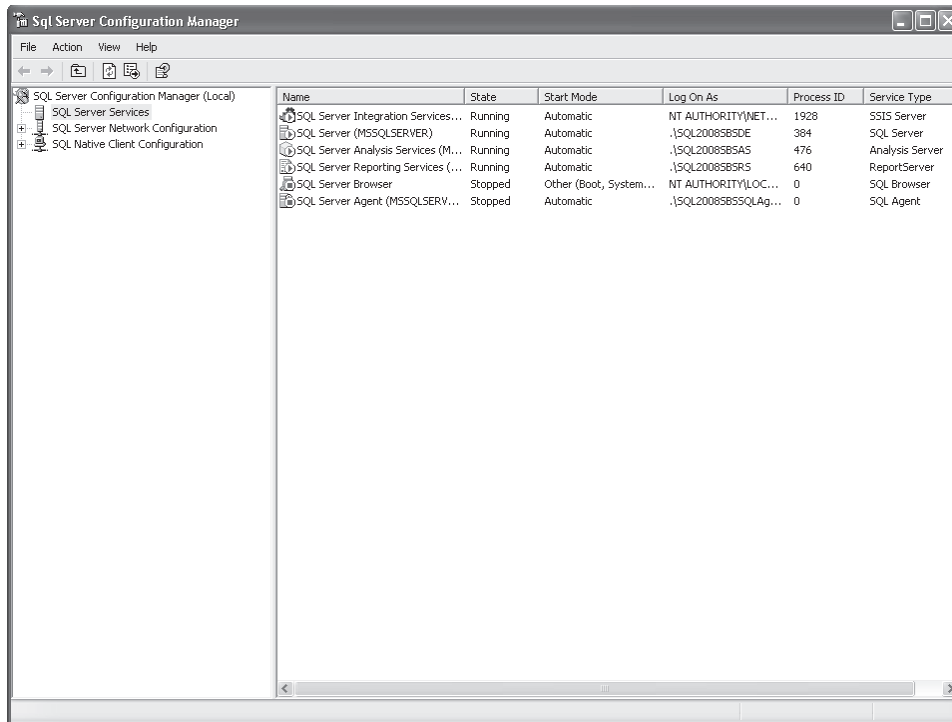


FIGURE 3-2 List of services within SQL Server Configuration Manager

While you can Start, Stop, Pause, and Restart SQL Server services, SQL Server has extensive management features which should ensure that you rarely, if ever, need to shut down or restart a SQL Server service.

In this procedure, you will review the options and settings available for SQL Server services.

Review Service Options

1. Launch SQL Server Configuration Manager by selecting Start | All Programs | Microsoft SQL Server 2008 | Configuration Tools | SQL Server Configuration Manager.
2. In the left-hand pane, highlight SQL Server Services.
3. Double-click the SQL Server service in the right-hand pane to display the Properties dialog box shown here.
4. Review the options on each of the tabs.
5. Click Cancel to close the Properties dialog box without making any changes.



SQL Server Configuration Manager also allows you to configure the communications protocols available to client connections. In addition to configuring protocol-specific arguments, you can also control whether communications are required to be encrypted or whether an instance will respond to an enumeration request, as shown in Figure 3-3.

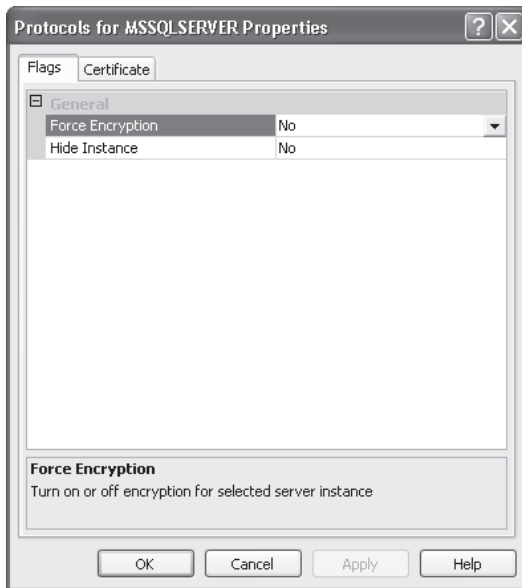


FIGURE 3-3 Protocol properties



Tip Applications can broadcast a special command, called an enumeration request, across a network to locate any SQL Servers that are running on the network. While being able to enumerate SQL Servers is valuable in development and testing environments where instances can appear, disappear, and be rebuilt on a relatively frequent basis, enumeration is not desirable in a production environment. By disabling enumeration responses by setting the Hide Instance to Yes, you prevent someone from using discovery techniques to locate SQL Servers for a possible attack.

SQL Server Management Studio

SQL Server Management Studio is the core tool that you will be spending a large part of your time using. SSMS provides all of the management capabilities for SQL Server services along with the ability to create and execute Transact-SQL (TSQL), Multidimensional Expression (MDX) query language, Data Mining Extensions (DMX), and XML for Analysis (XMLA) code. This section will provide a brief overview of SSMS, as shown in Figure 3-4, to get you started. Each subsequent chapter within this book will extend your knowledge of SSMS capabilities.

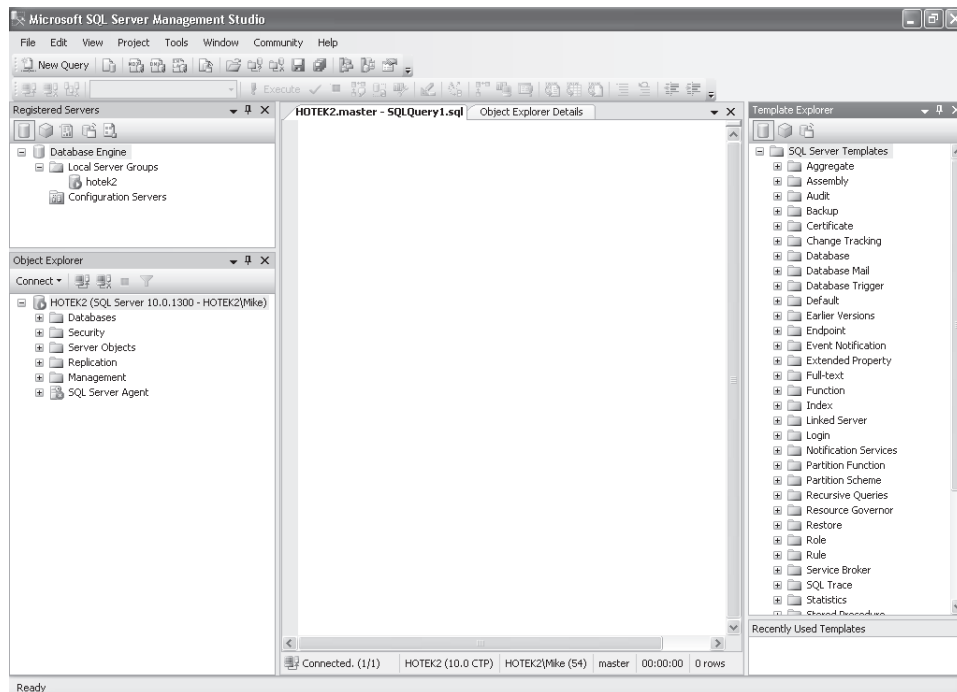


FIGURE 3-4 SQL Server Management Studio

In the following practice, you will launch SSMS and connect to the Database Engine instance that you installed in Chapter 2, “Installing and Configuring SQL Server 2008.”

Launch SSMS and Connecting to an Instance

1. Launch SSMS by selecting Start | All Programs | Microsoft SQL Server 2008 | SQL Server Management Studio.
2. When the Connect To Server dialog box is displayed, accept the default options and click Connect.



Note Because you have only installed a default instance at this point, this dialog should default to Database Engine for the server type, <machinename> for the server name, and Windows Authentication for the authentication option. Now that you have connected to an instance within SSMS, for all remaining exercises in this book, we will assume that you can perform these steps and will not repeat them.

SSMS has a variety of windows that you can open and position within the interface in order to access various feature sets.

The Registered Servers window provides a place to store connection information for all of the SQL Server services within your environment. Once stored, you can right-click any server and launch a connection to the server in either the Object Explorer or a query window.

The Template Explorer, shown in the right-hand pane of Figure 3-4, enables access to hundreds of predefined templates to create, alter, or drop objects as well as query various objects using TSQL, MDX, XMLA, or DMX. You can use the templates that ship with SQL Server, modify the templates to include your organization-specific coding standards, and add additional templates or template groups.

The Community menu on the toolbar allows you to launch a browser window into the center pane to access the MSDN forums and Microsoft Connect in the same way as previously described for Books Online.

The Tools | Options menu on the toolbar will display the Options dialog box, as shown in Figure 3-5 on page 40, so that you can set up the SSMS environment specifically the way you want to work.

Configure the SSMS Environment

1. Select Tools | Options from the toolbar.
2. Expand the Environment tree and select the General node. Use the At Startup drop-down list to configure the startup look and feel of SQL Server Management Studio.
3. Expand Text Editor | All Languages | Tabs.
4. Set the Tab Size to 4.

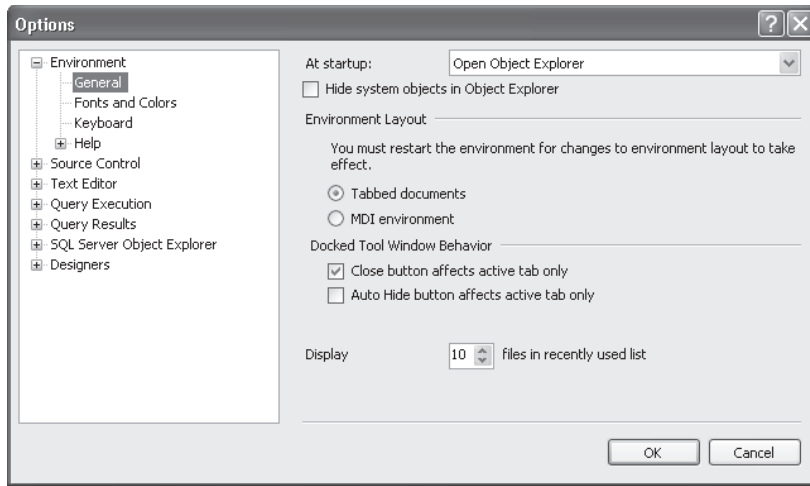


FIGURE 3-5 Options dialog box of the SQL Server Management Studio

5. Set the Indent Size to 4.
6. Select the Insert spaces option.
7. Explore the rest of the options that are available for configuration.
8. Click OK to save your settings.



Tip When you set SSMS to start up with an empty environment, you will not see a Connect To Server dialog and SSMS will immediately start. You will then need to explicitly connect to an instance for the Object Explorer or query window through either the Registered Server pane, File | Connect Object Explorer, or the New Query button. By setting the tab size and insert spaces options, SSMS will automatically replace any tabs with spaces in a query window, allowing you to more easily format and align code even when using a proportional font.

As you can see from Figure 3-6, the Object Explorer provides access to practically any action that you wish to perform against any SQL Server object. You will be using the functionality within the Object Explorer throughout virtually every chapter in this book.

Two additional capabilities of SSMS are object summaries and built-in reporting capabilities. The Object Explorer Details tab will display summary information according to the object that is currently selected within the Object Explorer. SSMS Reports, shown in Figure 3-7, allow you to display either Standard Reports that ship with SQL Server or to access your own custom reports that have been designed using the Reporting Services Report Designer that you will learn about in the article, "Reporting Services," which can be found on the Microsoft Press Online Windows Server and Client Web site at www.microsoft.com/learning/books/online/serverclient.

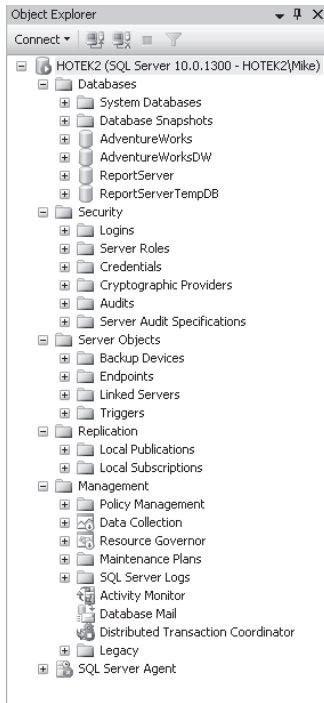


FIGURE 3-6 Object Explorer for a SQL Server instance

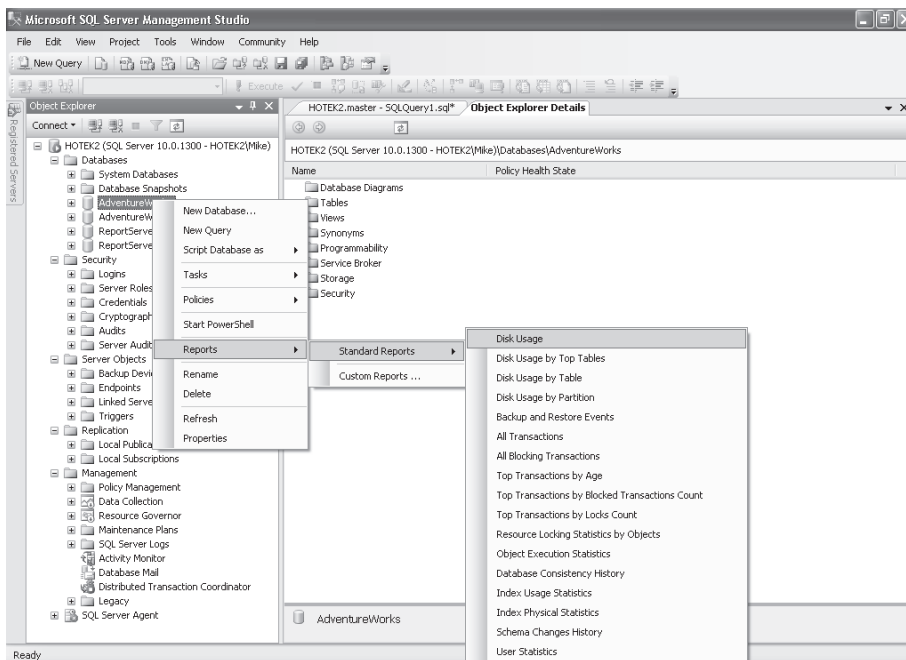


FIGURE 3-7 SSMS Reports

Database Mail

Database Mail enables a SQL Server to send outbound mail messages. While messages can contain the results of queries, Database Mail is primarily used to send alert messages to administrators to notify them of performance conditions or changes that have been made to objects. In the procedure that follows, you will learn how to configure Database Mail.

Configure Database Mail

1. Click the New Query button to open a new query window and execute the following code to enable the Database Mail feature:

```
EXEC sp_configure 'Database Mail XPs',1
GO
RECONFIGURE WITH OVERRIDE
GO
```
2. Within the Object Explorer, open the Management Node, right-click on Database Mail, and select Configure Database Mail.
3. Click Next on the Welcome screen.
4. Select the Set Up Database Mail by Performing the Following Tasks option and click Next.
5. Specify a name for your profile and click Add to specify settings for a mail account.
6. Fill in the account name, e-mail address, display name, reply e-mail, and server name fields on the New Database Mail Account page.
7. Select the appropriate SMTP Authentication mode for your organization and, if using Basic Authentication, specify the user name and password. Your settings should look similar to the following:

Specify name, description, and attributes for your SMTP account.

Account name: mhotek@mssqlserver.com

Description:

Outgoing Mail Server (SMTP)

E-mail address: mhotek@mssqlserver.com

Display name: Michael Hotek

Reply e-mail: mhotek@mssqlserver.com

Server name: Port number: 25

This server requires a secure connection (SSL)

SMTP Authentication

Windows Authentication using Database Engine service credentials

Basic authentication

User name:

Password:

Confirm password:

Anonymous authentication

OK Cancel Help




Note Your screen should look similar to the settings in the graphic. I am using my Internet e-mail account and have purposely left the Server Name, User Name, and Password out of the graphic. You will need to fill in the Server Name field if you are using an internal mail server.

8. Click OK and then click Next.
9. Check the box in the Public column next to the profile you just created and set this profile to Yes in the Default Profile column and click Next.
10. Review the settings on the Configure System Parameters page and click Next.
11. Click OK, then click Next, and then click Finish.
12. The final page should show success for all four configuration steps; click Close.
13. Within Object Explorer, right-click SQL Server Agent item and select Start from the shortcut menu in order to start the SQL Server Agent service, if it is not already running.



Note Database Mail utilizes the services of SQL Server Agent to send messages as a background process. If SQL Server Agent is not running, messages will accumulate in a queue within the msdb database.

14. Right-click Database Mail and select Send Test E-mail from the shortcut menu.
15. Select the Database Mail Profile you just created, enter an e-mail address in the To: line, and click Send Test E-Mail.



16. Go to your e-mail client and verify that you have received the test mail message.

Performance Management Tools

In addition to configuration and management tools discussed previously, SQL Server 2008 ships with three specialized tools for capturing, analyzing, and troubleshooting performance data.

Profiler

SQL Server Profiler is a graphical tool that acts as an interface to the SQL Trace Application Programming Interface (API). Profiler allows you to define SQL Server events, as shown in Figure 3-8, that you want to capture information on. You can also specify filtering options to target your data capture within the events that you have specified. You will learn about Profiler in the article, “Performance and Data Capture Tools,” which can be found on the Microsoft Press Online Windows Server and Client Web site at www.microsoft.com/learning/books/online/serverclient.

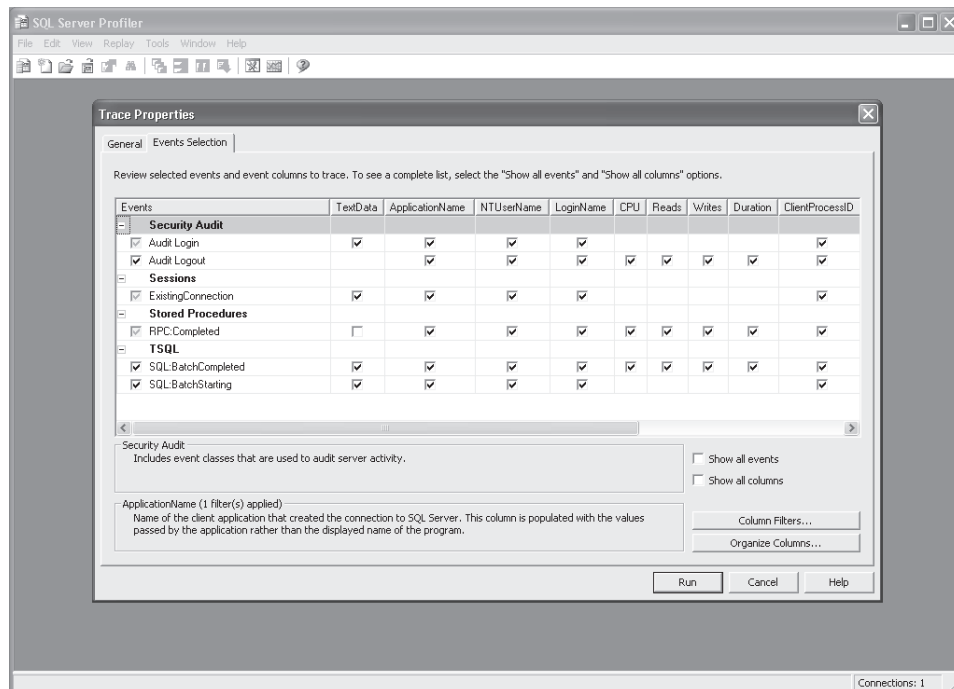


FIGURE 3-8 SQL Server Profiler Trace Properties dialog box Events selection

Database Engine Tuning Advisor

Database Engine Tuning Advisor (DTA) analyzes a query workload and makes recommendations on index and partitioning changes that can improve the performance of your queries (as shown in Figure 3-9). You will learn about indexes in Chapter 6, “Indexes,” partitioning in

Chapter 7, “Partitioning,” capturing a query workload in the article, “Performance and Data Capture Tools,” which can be found on the Microsoft Press Online Windows Server and Client Web site at www.microsoft.com/learning/books/online/serverclient, and how to apply DTA in the article, “Performance Analysis Tools,” which can be found on the Microsoft Press Online Windows Server and Client Web site at www.microsoft.com/learning/books/online/serverclient.

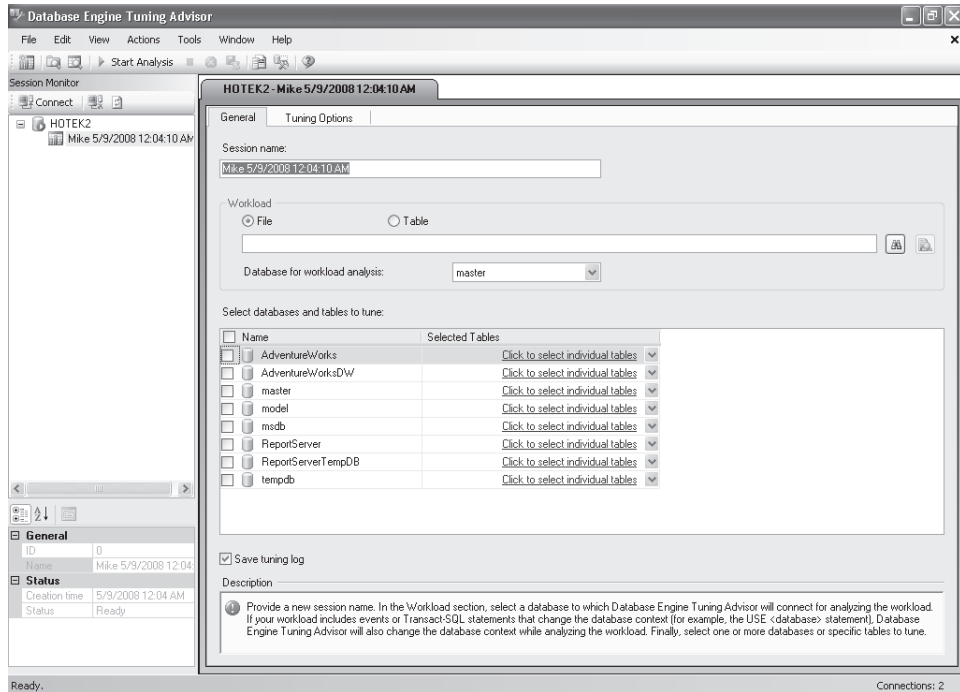


FIGURE 3-9 DTA tuning workload definition

Performance Studio

Performance Studio is the name given to a collection of technologies within SQL Server 2008 that are targeted at the analysis of enterprise-wide performance data. The components of Performance Studio are:

- Performance Data Warehouse
- Data Collectors
- Performance Reports

The Performance Data Warehouse is a database that you create. Data Collectors are SSIS packages, which are executed on a scheduled basis using SQL Server Agent. Performance Reports are a set of Report Designer reports written against the data stored in the Performance Data Warehouse.

You will learn how to configure, manage, and leverage the components of the Performance Studio in the article, "Performance Analysis Tools," which can be found on the Microsoft Press Online Windows Server and Client Web site at www.microsoft.com/learning/books/online/serverclient.

Business Intelligence Tools

Management of SSIS, SSRS, and SSAS occurs within SSMS. However, development of packages, reports, report models, OLAP cubes, and Data mining models occurs within the Business Intelligence Development Studio (BI Dev Studio).

Business Intelligence Development Studio

The BI Dev Studio is the Visual Studio 2008 shell with support for SQL Server 2008 BI projects, as shown in Figure 3-10. Each of the BI projects will be explored in Chapter 24, "Integration Services," Chapter 25, "SQL Server Reporting Services," and Chapter 26, "SQL Server Analysis Services."

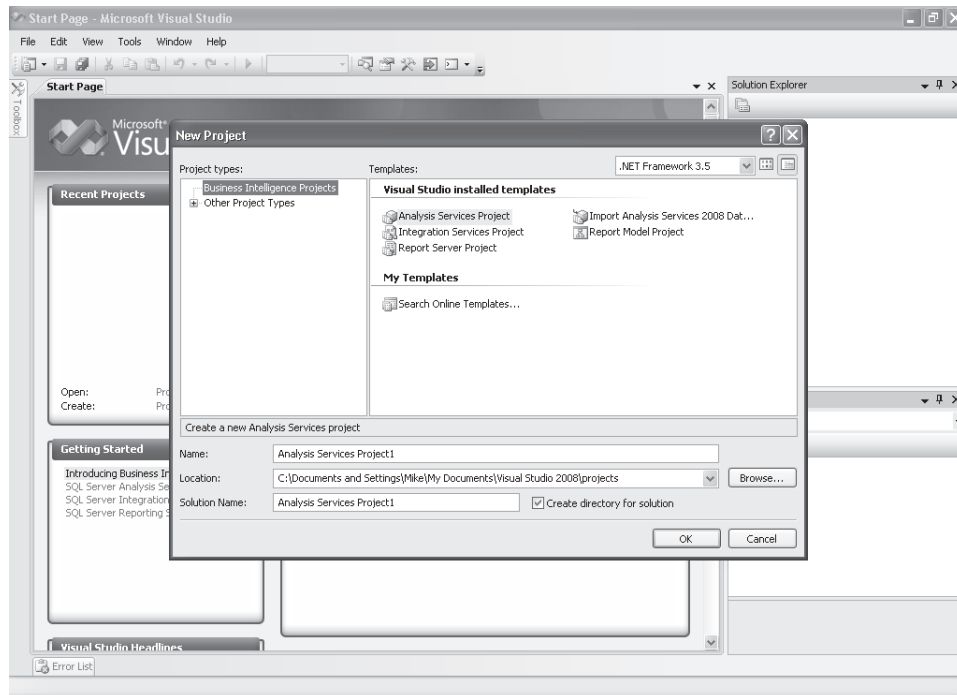


FIGURE 3-10 BI Dev Studio projects

Chapter 3 Quick Reference

To	Do This
Manage a SQL Server instance, OLAP cubes, Data mining models, Integration Services, or Reporting Services	Connect to the appropriate service within the Object Explorer in SQL Server Management Studio
Write and/or execute T-SQL, MDX, or DMX code	Open the appropriate query window (T-SQL, MDX, or DMX) and connect to the instance you want to write or execute code against
Configure and manage Database Mail	Connect to a Database Engine instance
Get help on a topic	Launch SQL Server Books Online. If you are writing a query, you can highlight a term and use SHIFT+F1 to launch Books Online to the highlighted term
Start, Stop, and Pause a service or manage service accounts	Open SQL Server Configuration Manager and either use the Start, Stop, or Pause buttons on the toolbar or double-click the service to access the property sheet
Limit the resources a particular query or user can consume	Configure the Resource Governor within SSMS
Build SSIS, SSRS, or SSAS projects and deploy them to a SQL Server instance	Open BI Dev Studio, create a new project, and design the objects that you wish to deploy
Manage Performance	Use Profiler to capture query workloads, DTA to analyze workloads captured with Profiler, and the Performance Studio to automate the capture of performance metrics

Chapter 14

Triggers

After completing this chapter, you will be able to

- Create DML triggers
- Create DDL triggers

Triggers provide a means to allow you to automatically execute code when an action occurs. Two types of triggers are available in Microsoft SQL Server 2008: DML and DDL. In this lesson, you will learn how to create DML triggers that execute when you add, modify, or remove rows in a table. You will also learn how to create DDL triggers that execute when DDL commands are executed or users log in to an instance.

DML Triggers

Although functions and stored procedures are stand-alone objects, you can't directly execute a trigger. *DML triggers* are created against a table or a view, and are defined for a specific event—*INSERT*, *UPDATE*, or *DELETE*. When you execute the event a trigger is defined for, SQL Server automatically executes the code within the trigger, also known as "firing" the trigger.

The generic syntax for creating a trigger is:

```
CREATE TRIGGER [ schema_name . ]trigger_name
ON { table | view }
[ WITH <dml_trigger_option> [ ,...n ] ]
{ FOR | AFTER | INSTEAD OF }
{ [ INSERT ] [ , ] [ UPDATE ] [ , ] [ DELETE ] }
[ WITH APPEND ]
[ NOT FOR REPLICATION ]
AS { sql_statement [ ; ] [ ,...n ] | EXTERNAL NAME <method specifier [ ; ] > }
```

When a trigger is defined as *AFTER*, the trigger fires after the modification has passed all constraints. If a modification fails a constraint check, such as a check, primary key, or foreign key, the trigger is not executed. *AFTER* triggers are only defined for tables. You can define multiple *AFTER* triggers for the same action.

A trigger defined with the *INSTEAD OF* clause causes the trigger code to be executed as a replacement for *INSERT*, *UPDATE*, or *DELETE*. You can define a single *INSTEAD OF* trigger for a given action. Although *INSTEAD OF* triggers can be created against both tables and views, *INSTEAD OF* triggers are almost always created against views.

Regardless of the number of rows that are affected, a trigger only fires once for an action.

As explained in Chapter 10, “Data Manipulation,” SQL Server makes a pair of tables named inserted and deleted available when changes are executed.

In the following exercise, you will create a DML trigger that populates the FinalShipDate column in the Orders.OrderHeader table when the ShipDate column has been populated for all rows in the Orders.OrderDetail table for an OrderID.

Create a DML Trigger

1. Execute the following code against the SQL2008SBS database (the code is from the Chapter14\code1.sql file in the book’s accompanying samples):

```
CREATE TRIGGER tiud_orderdetail ON Orders.OrderDetail
FOR INSERT, UPDATE, DELETE
AS

UPDATE a
SET a.FinalShipDate = c.FinalShipDate
FROM Orders.OrderHeader a INNER JOIN
    (SELECT od1.OrderID, MAX(od1.ShipDate) FinalShipDate
    FROM Orders.OrderDetail od1 INNER JOIN
        (SELECT od2.OrderID
        FROM Orders.OrderDetail od2 INNER JOIN inserted i ON od2.OrderID = i.OrderID
        WHERE od2.ShipDate IS NOT NULL
        EXCEPT
        SELECT od3.OrderID
        FROM Orders.OrderDetail od3 INNER JOIN inserted i ON od3.OrderID = i.OrderID
        WHERE od3.ShipDate IS NULL) b
    ON od1.OrderID = b.OrderID
    GROUP BY od1.OrderID) c
ON a.OrderID = c.OrderID
GO
```

2. Validate your newly created trigger by setting the ShipDate column for all order detail rows for an order.

In the following exercise, you will create a DML trigger that enforces referential integrity between the SQL2008SBS and SQL2008SBSFS databases.

Create a DML Trigger

1. Execute the following code against the SQL2008SBS database (the code is from the Chapter14\code2.sql file in the book’s accompanying samples):

```
USE SQL2008SBSFS
GO

CREATE TRIGGER tiu_productdocuments ON Products.ProductDocument
FOR INSERT, UPDATE
AS
IF EXISTS (SELECT 1 FROM SQL2008SBS.Products.Product a
    INNER JOIN inserted b ON a.ProductID = b.ProductID)
```

```

BEGIN
    RETURN
END
ELSE
BEGIN
    ROLLBACK TRANSACTION
    RAISERROR('Violation of foreign key',16,1)
END
GO

USE SQL2008SBS
GO

CREATE TRIGGER td_product ON Products.Product
FOR DELETE
AS
IF EXISTS (SELECT 1 FROM SQL2008SBSFS.Products.ProductDocument a
           INNER JOIN deleted b ON a.ProductID = b.ProductID)
BEGIN
    ROLLBACK TRANSACTION
    RAISERROR('You must first delete all documents for this product',16,1)
END
ELSE
BEGIN
    RETURN
END
GO

```

2. Validate your newly created trigger by attempting to insert a document with a ProductID that does not exist.

DDL Triggers

DDL triggers execute under the following circumstances:

- DDL is executed.
- A user logs into an instance.

The general syntax for creating a DDL trigger is as follows:

```

CREATE TRIGGER trigger_name
ON { ALL SERVER | DATABASE }
[ WITH <ddl_trigger_option> [ ,...n ] ]
{ FOR | AFTER } { event_type | event_group } [ ,...n ]
AS { sql_statement [ ; ] [ ,...n ] | EXTERNAL NAME <method_specifier> [ ; ] }

<ddl_trigger_option> ::=
    [ ENCRYPTION ] [ EXECUTE AS Clause ]

<method_specifier> ::=
    assembly_name.class_name.method_name

```

DDL triggers can be scoped at either the database or instance level. To scope a DDL trigger at the instance level, you utilize the ON ALL SERVER option. To scope a DDL trigger at the database level, you utilize the ON DATABASE option.

The following is an example of a DDL trigger:

```
CREATE TRIGGER tddl_tabledropalterprevent
ON DATABASE
FOR DROP_TABLE, ALTER_TABLE
AS
    PRINT 'You are attempting to drop or alter tables in production!'
    ROLLBACK;
```



Note Almost all DDL commands run within the context of a transaction. Since a DDL trigger also runs within the same transaction context, any DDL statement running in the context of a transaction can be rolled back. *ALTER DATABASE* is one of the commands which does not execute in the context of a transaction, because the command affects objects outside of SQL Server that do not obey transactional semantics. Therefore an *ALTER DATABASE* command cannot be rolled back.

The value for the event type is derived from the DDL statement being executed, as listed in Table 14-1.

TABLE 14-1 DDL Trigger Event Types

DDL Command	Event Type
<i>CREATE DATABASE</i>	CREATE_DATABASE
<i>DROP TRIGGER</i>	DROP_TRIGGER
<i>ALTER TABLE</i>	ALTER_TABLE

Event types roll up within a command hierarchy called *event groups*. For example, the CREATE_TABLE, ALTER_TABLE, and DROP_TABLE event types are contained within the DDL_TABLE_EVENTS event group. Event types and event groups allow you to create flexible and compact DDL triggers.



More Info The events and associated event groups that are valid for a DDL triggers can be found in the Books Online article, “Event Groups for Use with DDL Triggers.”

Although DML triggers have access to the inserted and deleted tables, DDL triggers have access to the *EVENTDATA()* function which returns the following XML document that can be queried by using the *value()* method available through XQUERY:

```
<EVENT_INSTANCE>
  <EventType>type</EventType>
  <PostTime>date-time</PostTime>
```



```

<SPID>spid</SPID>
<ServerName>name</ServerName>
<LoginName>name</LoginName>
<UserName>name</UserName>
<DatabaseName>name</DatabaseName>
<SchemaName>name</SchemaName>
<ObjectName>name</ObjectName>
<ObjectType>type</ObjectType>
<TSQLCommand>command</TSQLCommand>
</EVENT_INSTANCE>

```

You can retrieve the database, schema, object, and command that you executed, through the following query:

```

SELECT EVENTDATA().value
    ('(/EVENT_INSTANCE/DatabaseName)[1]', 'nvarchar(max)'),
EVENTDATA().value
    ('(/EVENT_INSTANCE/SchemaName)[1]', 'nvarchar(max)'),
EVENTDATA().value
    ('(/EVENT_INSTANCE/ObjectName)[1]', 'nvarchar(max)'),
EVENTDATA().value
    ('(/EVENT_INSTANCE/TSQLCommand)[1]', 'nvarchar(max)')

```

In the following exercise, you create a DDL trigger to prevent accidentally dropping tables in a production environment.

Create a Database Level DDL Trigger

1. Execute the following code against the SQL2008SBS database (the code is from the Chapter14\code3.sql file in the book's accompanying samples):

```

CREATE TRIGGER tddl_preventdrop
ON DATABASE
FOR DROP_TABLE
AS
    PRINT 'Please disable DDL trigger before dropping tables'
    ROLLBACK TRANSACTION
GO

```

2. Validate your trigger by attempting to drop a table in the SQL2008SBS database.

In the following exercise, you create a logon trigger to limit the number of concurrent connections to a user.

Create an Instance Level DDL Trigger

1. Execute the following code (the code is from the Chapter14\code4.sql file in the book's accompanying samples):

```

CREATE TRIGGER tddl_limitconnections
ON ALL SERVER
FOR LOGON

```

```

AS
BEGIN
IF (SELECT COUNT(*) FROM sys.dm_exec_sessions
    WHERE is_user_process = 1 AND
        login_name = suser_sname()) > 5

    PRINT 'You are only allowed a maximum of 5 concurrent connections'
    ROLLBACK
END
GO

```

2. Validate your trigger by attempting to create more than five concurrent connections.



Note You have to be careful with a logon trigger, especially one that prevents logging on to the instance. In the exercise above, you had the trigger apply to **all** logins. You should always exclude logins that are members of the sysadmin role, because you do not want to cause a sysadmin to not be able to log in to an instance.

Chapter 14 Quick Reference

To	Do This
Execute code when a DML command is executed	Create a DML trigger
Execute code when a DDL command is executed	Create a DDL trigger

Index

Symbols and Numbers

.ldf, 53
.mdf, 53, 57
.ndf, 53
.NET Framework 2.0, 17
.NET Framework 3.5, 22

A

accent sensitivity, 19
accounts, service, 17–19, 35–38
ACTIVATION, 242
ActiveX scripts, 319–20, 381
Ad Hoc Distributed Queries, 260
administrative procedures, 196
administrative users, 269–73
AdventureWorks2008, 28
AdventureWorksDW2008, 28
AFTER, 213
after image, 167
aggregate functions, 199, 279–80
Aggregate transform, 383
aggregating data, 123, 138–39
 derived tables, 139–40
 filtering aggregates, 143–44
 multiple permutations, 140–42
 OLAP (Online Analytic Processing), 453–54
 pivot tables, calculating, 145
 ranking data, 146–47
 result sets, 147–49
 running aggregates, 144–45
alerts, 6. *See also* SQL Server Agent
 creating, 329–32
 Database Mail, 42–43
 jobs, 321–25
 monitor server, 357
algorithms
 data mining, 472
 categories, 473–74
 hash algorithms, 282–84
aliases, 123
ALL, 105, 263
ALTER, 76–77, 208, 275
ALTER DATABASE, 59, 216
ALTER FUNCTION, 208
ALTER INDEX, 99–100
ALTER INDEX...REBUILD, 309–10
ALTER PARTITION SCHEME, 105
ALTER PROCEDURE, 208
ALTER TABLE, 216
ALTER TRIGGER, 208
ALTER VIEW, 208
Analysis Services. *See* SQL Server Analysis Server (SSAS)
Analysis Services Execute DDL Task, 380
ANSI character set, 70
ANSI_Nulls, 183
API (SQL Trace Application Programming Interface), 44
 applications
 client redirect, 354
 line of business, 9–10
architecture
 data storage, 56
 processor, 69
 Service Broker, 226–28
archive tables, 108
arrays, 129
articles, 364
AS clause, 104
AS PARTITION clause, 105
AS SNAPSHOT OF, 219
ascending sort, 126
assemblies, 279–80
association algorithm, 473
Association Rules, 474
Asymmetric Key, login, 264–68
asymmetric keys, 265, 281–82, 285–86
asynchronous message queuing
 architecture, 226–28
 contracts, 232–33
 conversations, 235–36
 message types, 229–32
 overview, 225–26
 prioritization, 245
 queues, 233–34, 242–44
 sending/receiving messages, 236–42
 services, 234–35
attach, database moving, 59–60
attribute-centric XML format, 151
attributes, OLAP, 453–54
attributes, regression algorithms, 473–74
auditing
 data capture, 173–75
 sys.database_audit, 334
 tracking changes, 171–73
authentication modes, 19, 259
AUTO, 151
availability
 database mirroring, 348
 caching, 353–54
 client redirect, 354
 Database Snapshots, 349, 354
 endpoints, 350–51
 initializing, 355–56

- operating modes, 351–53
 - page corruption, 354
 - roles, 349–50
 - failover clustering, 345–48
 - log shipping, 356
 - components, 356–57
 - initialization, 357–63
 - overview, 7
 - replication, 363
 - agents, 365–66
 - components, 364
 - configuring, 367–73
 - methods, 366–67
 - roles, 365
- AVG, 183, 383

B

- background processing, reports, 441
- backup job, 358
- BACKUP MASTER KEY, 282
- backups, 301
 - database mirroring, 355–56
 - database restores, 311–16
 - Database Snapshot, 220
 - db_backupoperator, 269
 - differential backups, 307–8, 314–15
 - filegroups, 308
 - full backups, 302–5, 311–12
 - log shipping, 358–59
 - maintenance plans, 325–28
 - page corruption, 308–9
 - partial, 107
 - recovery models, 309–11
 - transaction log backups, 305–7, 315–16
- balanced trees (B-trees), 88–89
- BCP (Bulk Copy Program), 34, 309–10, 366
- before image, 167
- BEGIN statements, 190–91
- BEGIN TRAN, 167–68
- BEGIN...END statements, 190–91
- BI (business intelligence), 9–14
- BIDS (Business Intelligence Development Studio), 46
 - overview, 378–80
 - tasks, 380–82
 - transforms, 382–85
- bigint, 68
- binary data, 72
- bit, 72
- BLOBs (Binary Large Objects), 74
- blocking, tables, 112
- bookmarks, reports, 431
- Books Online, 31–33
- boundary points, 103–5, 128
- BREAK, 191
- B-trees (balanced trees), 88–89
- Bulk Copy Program (BCP), 34, 309–10, 366
- BULK INSERT, 309–10, 382
- Bulk Insert task, 382
- Bulk Logged recovery model, 309–11
- bulkadmin, 266
- Business Intelligence (BI), 9–14
- Business Intelligence Development Studio (BIDS), 46
 - overview, 378–80
 - tasks, 380–82
 - transforms, 382–85

C

- CA (certificate authority), 285–86
- Cache transform, 383
- caches, 353–54, 444–49
- calculations
 - data retrieval, 122
 - functions, 199–205
 - key performance indicators (KPIs), 470
 - measures, OLAP, 454
 - OLAP, 462–66
 - pivot tables, 145
 - processors and, 69
 - variables and, 187
- CALLED ON NULL INPUT, 201–2
- CALLER, 202
- CASE function, 125–26
- case sensitivity, 19
- CAST, 123–26
- catalogs, full-text, 223
- CATCH, 192–93
- CDOC (Cross Database Ownership Chaining), 260
- certificate authority (CA), 285–86
- Certificate, login, 264–68
- certificates, 265, 281–82, 285–86
- change tables, 171–73
- change tracking, 171–75, 333
- CHANGE_TRACKING, 250
- char, 70
- char(n), 70
- CHAR/VARCHAR, 249
- CHAR/VARCHAR/XML, 249
- character data, 19, 69–71, 129, 384
- Character Map, 384
- charts, formatting, 423–28
- CHECK, 204–11
- check constraints, 82
- CHECK_EXPIRATION, 265
- CHECK_POLICY, 265
- CHECKSUM, 304, 309
- classification algorithms, 473
- client connections, 37
- client redirection, 354
- CLOSE, 194–96
- CLOSE MASTER KEY, 282
- CLR (Common Language Runtime), 196, 260, 279–80, 333
- CLR Enabled, 260
- clustered indexes, 91–93

- clustering key, 91–93
- clustering, failover, 7, 345–48
- clustering, SSRS, 416
- COALESCE, 124
- Code Access Security, 279–80
- CodePlex web site, 28
- COLLATE, 75
- collation, 75
- collation sequences, 19
- column filter, 364
- columns, table
 - clustered indexes, 91–93
 - computed, 77–78
 - data retrieval, 122–23, 126
 - included columns, indexes, 95–96
 - properties, 75
 - sparse, 78–79
 - views, 181–82
- command line utilities
 - OSQL, 34
 - SQLCMD, 34
- commenting code, 186
- COMMIT, 167
- COMMIT TRAN, 167–68
- Common Language Runtime (CLR), 196, 260, 279–80, 333
- common table expressions (CTEs), 149–50
- communications protocols
 - endpoints, 260–63
 - Service Broker
 - architecture, 226–28
 - contracts, 232–33
 - conversations, 235–36
 - message types, 229–32
 - overview, 225–26
 - prioritization, 245
 - queues, 233–34, 242–44
 - sending/receiving messages, 236–42
 - services, 234–35
 - SQL Server Configuration Manager, 37
- compound criteria, filtering data, 127–30
- compression, backups, 304
- computations
 - Aggregate transforms, 383
 - data retrieval, 125
 - views, 181
- COMPUTE, 179, 183
- COMPUTE BY, 179, 183
- concatenated data, 124
- conditional expressions, 125
- Conditional Split, 384
- conditions, 290–93
- configuration file, SSIS packages, 406–9
- configuration functions, 199
- Configuration Manager, 35–38
- configuring. *See also* Configuration Manager
 - Database Mail, 42–43
 - SQL Server Management Studio (SSMS), 39–40
- CONNECT, 261–62
- connection strings, 386–90
- connections, server
 - endpoints, 260–63
 - query execution stats, 339–40
 - SSIS, 386–90
- consistency checks, 325–28
- constants, 122
- constraints
 - check, 82
 - clustered indexes, 92
 - default, 82–83
 - DML triggers, 213
 - foreign keys, 83–84
 - primary keys, 79–81
 - unique, 81
 - XML indexes, 101
- CONTAINS, 251–55
- CONTAINSTABLE, 251–55
- CONTINUE, 191
- CONTINUE_PAST_ERROR, 304
- CONTROL, 275
- control flow, 390–94
- control flow constructs, 189–92
- control flow options, 320
- Control Flow, SSIS tab, 380
- CONVERT, 123–26
- COOKIE INTO, 202
- cookies, 202
- copy job, 358
- Copy Objects Wizard, 22
- copying
 - BCP (Bulk Copy Program), 34, 309–10, 366
 - Object Explorer, 95
- Copy-On-Write, 220–21
- correlated sub-queries, 133
- COUNT, 125
- covering indexes, 95–96
- CREATE, 76–77
- CREATE BROKER PRIORITY, 245
- CREATE DATABASE, 52, 57–58, 216, 219–20
- CREATE ENDPOINT, 262
- CREATE INDEX, 309–10
- CREATE PARTITION SCHEME, 105
- CREATE TABLE, 159
- CREATE USER, 268–69
- creating
 - alerts, 329–32
 - clustered indexes, 92–93
 - cubes, 454–62
 - database diagrams, 85–86
 - databases, 51–58
 - functions, 200–3
 - indexes, 93–94, 96, 98–99, 101, 107–8
 - jobs, 319–25
 - maintenance plans, 325–28
 - reports, SSRS, 418–38
 - schemas, 64–66, 72–74

- service accounts, 17–19
- stored procedures, 185
- tables, 76–77, 80–82
- XML indexes, 101
- Cross Database Ownership Chaining (CDOC), 260
- CROSS JOIN, 131–33, 181
- CROSS/OUTER APPLY, 183, 204–11
- cryptographic functions, 199
- CSS (Microsoft Customer Service and Support), 34–35
- CTEs (common table expressions), 149–50
- CUBE operator, 141–42
- cubes, OLAP (Online Analytic Processing), 13–14, 454–62
 - dimensions, measures, calculations, 462–66
 - hierarchies, 467–70
 - partitions, 471
 - perspectives, 471
 - translations, 471–72
- cursors, 194–96, 199
- cut, Object Explorer, 95

D

- DAC (Dedicated Admin Connection), 260
- data
 - availability, 7
 - BLOBs (Binary Large Objects), 74
 - Bulk Copy Program (BCP), 34
 - character, 69–71
 - data flow, SSIS, 394–403
 - date and time, 124
 - loss exposure, 358
 - numeric, 67–69
 - SPATIAL, 74–75
 - storage, 56, 72–75
 - Tablediff, 34
 - types of, 67
- data cache, 353
- Data Collectors, Performance Studio, 45–46
- Data Conversion transform, 384
- Data Definition Language (DDL), 76–77, 269, 380
- data encryption. *See* encryption
- data filegroups, 56
- data files, 53–57
- Data Flow, SSIS, 380, 382–85
- data manipulation
 - deleting data, 161–64
 - functions, 123
 - inserting data, 153–59
 - MERGE, 164–66
 - OUTPUT, 166–67
 - tracking changes, 171–75
 - transaction handling, 167–71
 - updating data, 159–61
- Data Manipulation Language (DML), 76–77, 236–37, 290–93
- data mining, 380
 - algorithms, 473–74
 - demystified, 483–84
 - models and structures, 474–83
 - overview, 472
- Data Mining Engine, 14
- Data Mining Extensions (DMX), 38
- Data Mining Query, 380
- data recovery
 - backups, 301
 - database restores, 311–16
 - differential backups, 307–8, 314–15
 - filegroup backups, 308
 - full backups, 302–5, 311–12
 - page corruption, 308–9
 - recovery models, 309–11
 - transaction log backups, 305–7, 315–16
- data retrieval
 - aggregating data, 138–39
 - derived tables, 139–40
 - filtering aggregates, 143–44
 - multiple permutations, 140–42
 - pivot tables, calculating, 145
 - ranking data, 146–47
 - result sets, 147–49
 - running aggregates, 144–45
 - common table expressions (CTEs), 149–50
 - filtering data, 127–30
 - functions, 204–11
 - multiple tables, 131–33
 - overview, 121
 - SELECT statement, 126, 137–38
 - sorting results, 126
 - unique results, 134–35
 - views
 - creating, 179–80
 - indexed views, 182–84
 - modifying data, 181–82
 - XML data querying, 150–51
- Data Transformation Services (DTS), 10
- data types
 - changing, 123–26
 - security, 279–80
 - Visual Studio to SQL Server mapping, 384
- data warehousing, 451–52
- data, unstructured. *See* full-text indexing
- DATA_BASE_MIRRORING, 350–51
- Database Diagrams, 85–86
- Database Engine
 - availability, 7
 - creating account, 17–19
 - installing, 22–29
 - overview, 3
 - programming interface, 5–6
 - replication, 6–7
 - security subsystem, 4–5
 - Service Broker, 6
 - SQL Server Agent, 6
 - Storage Engine, 4
- Database Engine Tuning Advisor (DTA), 44–45
- database level principals, 268–69
- Database Mail, 42–43, 260

- database master key, 281–82
- database mirroring, 263, 348
 - caching, 353–54
 - client redirect, 354
 - Database Snapshots, 354
 - endpoints, 350–51
 - initializing, 355–56
 - operating modes, 351–53
 - overview, 7
 - page corruption, 354
 - roles, 349–50
 - TCP endpoints, 262
- Database Snapshots, 221–22
 - Copy-On-Write, 220–21
 - database mirroring, 349, 354
 - overview, 219–20
 - reverting data, 222–23
- DATABASE_MIRRORING, 261
- databases
 - creating, 57–58
 - designing, 64–67
 - diagrams, 85–86
 - file structure, 53–57
 - maintenance plans, 325–28
 - moving, 59–60
 - attach, 60
 - overview, 51–53
 - sample, installing, 28–29
 - vs. spreadsheets, 67
- data-driven subscriptions, 441
- DATEADD, 208
- dates, 71–72, 124, 199, 424
- datetime, 71–72
- datetime2, 71–72
- datetimeoffset, 71–72
- db_accessadmin, 269
- db_backupoperator, 269
- db_datareader, 269
- db_datawriter, 269
- db_ddladmin, 269
- db_denydatareader, 269
- db_denydatawriter, 269
- db_owner, 269–73
- db_securityadmin, 269
- dbcreator, 266
- dbo.syspolicy_management_facets table, 290
- DCM (Differential Change Map), 307–8
- DDL (Data Definition Language), 76–77, 269, 380
- DDL triggers, 216, 218, 295
- deadlocking, tables, 112
- DEALLOCATE, 194–96
- debugging
 - IF statements, 191
 - WHILE statements, 191
- decimals, 67–69
- DECLARE, 186, 194–96
- decryption, 284. *See also* encryption
- Dedicated Admin Connection (DAC), 260
- DEFAULT, 154, 204–11
- default constraints, 82–83
- DEFAULT filegroup, 56–58
- defense in depth, 259
- defragmenting, indexes, 99–100
- DELETE
 - db_datawriter, 269
 - db_denydatawriter, 269
 - deleting data, 161–63
 - DML triggers, 213
 - locks, 168
 - MERGE, 164–66
 - OUTPUT, 167
 - performance degradation, 93–94
 - permissions, 275
 - transaction handling, 167–68
- deleted tables, 166–67
- deleting data, 161–64
- DENSE_RANK, 146
- dependent objects, 202
- deploying packages, SSIS, 409–10
- deploying reports, SSRS, 439–40
- Derived Column transform, 384
- derived tables, 139–40
- descending sort, 126
- designing databases, 85–86
- designing tables, 63–64
 - binary data, 72
 - columns, 75, 77–79
 - constraints, 79–84
 - creating tables, 76–77
 - data types, 67–71
 - database design, 64–67
 - database diagrams, 85–86
 - date and time data, 71–72
 - FILESTREAM, 74
 - hierarchyID data, 75
 - naming objects, 64
 - schemas, 64–66
 - SPATIAL data, 74–75
 - XML data, 72–74
- detach, database moving, 59–60
- deviation analysis, 474
- diagnostics
 - object size, 336
 - SQLDiag, 34–35
 - sys.dm_os, 334
- DIALOG, 236
- dialog handles, 236
- differential backups, 307–8, 314–15
- Differential Change Map (DCM), 307–8
- Digital Rights Management (DRM), 201
- digital signatures, 285–86
- dimension tables, 454
- dimensions, OLAP, 453–54, 462–66
- directories, permissions, 347
- DISABLED, 261–62
- disabling, index, 100

disaster recovery. *See also* database mirroring

- backups, 301
- database restores, 311–16
- differential backups, 307–8, 314–15
- filegroup backups, 308
- full backups, 302–5, 311–12
- page corruption, 308–9
- recovery models, 309–11
- transaction log backups, 305–7, 315–16

discretized attributes, 474

disk configuration, 346

disk drives, failover clustering, 345–48

disk space, installation requirements, 17

diskadmin, 266

DISTINCT, 134–35, 138, 181, 183

Distrib.exe, 366

Distribution Agent, 366–67

distribution statistics, 97

distributor, 365

DML (Data Manipulation Language), 76–77, 236–37, 290–93

DMVs (dynamic management views)

- indexes, 337–39
- object size, 336
- overview, 333–34
- query execution stats, 339–40
- retrieving object metadata, 334–35

DMX (Data Mining Extensions), 38

documentation, 31–33

documents, XML, 72–73

drag and drop, Object Explorer, 95

DRM (digital rights management), 201

DROP, 76–77

DROP TRIGGER, 216

DROP/CREATE, 208

DTA (Database Engine Tuning Advisor), 44–45

DTS (Data Transformation Services), 10

dynamic execution, 193–94

dynamic management views (DMVs)

- indexes, 337–39
- object size, 336
- overview, 333–34
- query execution stats, 339–40
- retrieving object metadata, 334–35

dynamic row filter, 364

dynamic scalability, 242

E

editions, SQL Server, 15–16

EKM (Extensible Key Management), 287

EKM (External Key Management), 260

element-centric XML data, 151

ELEMENTS, 151

ELSE statements, 190–91

e-mail

- architecture, 226–28
- contracts, 232–33

- conversations, 235–36
- message types, 229–32
- overview, 225–26
- prioritization, 245
- queues, 233–34, 242–44
- sending/receiving messages, 236–42
- services, 234–35

EMPTY, validation option, 229

encryption

- asymmetric keys, 285–86
- certificates, 285–86
- Extensible Key Management (EKM), 287
- hash algorithms, 282–84
- key management, 260
- master keys, 281–82
- overview, 280–81
- SQL Server Configuration Manager, 37
- SSRS, 415
- symmetric keys, 284
- TCP endpoints, 263
- Transparent Data Encryption (TDE), 286–87

ENCRYPTION, 201, 236

Encryption Keys, 415

END statements, 190–91

endpoints, 260–63, 350–51

equality, 127–28

equality joins, 133

errors, 192–93, 329–32

evaluation order, 125

event groups, 216

Event Handlers, SSIS, 380

event logs, SQLDiag, 34–35

EVENTDATA(), 216

events

- alerts, creating, 329–32
- configuration, 334
- notification policies, 295
- SQL Server Profiler, 44
- WMI, BIDS tasks, 382

Excel documents, 249

EXCEPT, 148–49, 181, 183

exception handling, SSIS, 403–6

exclusive locks, 168

EXEC, 189, 193–94

EXECUTE, 202, 275

EXECUTE AS, 201–2, 273

Execute tasks, 381

execution, dynamic, 193–94

EXISTS, 133

EXPLICIT, 151

Export Column, 384

exporting data, 34

Extensible Key Management, 287

Extensible Markup Language (XML).

- See* XML (Extensible Markup Language)

extents, 55

External Key Management (EKM), 260

EXTERNAL_ACCESS, 279–80

F

facets, 290
 fact tables, 454
 failover clustering, 7, 345–48
 fatal error codes, 193
 FETCH, 194–96
 File System task, 382
 FILEGROUP, 248
 filegroups, 56–57
 backups, 308
 database mirroring, 348
 message queues, 233
 partitions, 105, 110
 FILEGROWTH, 54–55
 files. *See* FILESTREAM
 files, structure of, 53–57
 FILESTREAM, 219. *See also* full-text indexing
 creating, 57–58
 data storage, 74
 database mirroring, 348
 Database Snapshot, 223
 filegroups, 56–57
 index files, 55
 security, 260
 Filestream Access Level, 260
 fill factor, indexes, 99–100
 filtered indexes, 97–98
 filtering
 aggregate filtering, 143–44
 data retrieval, 127–30
 EXISTS, 133
 full text indexes, 249
 replication, 364
 table joins, 131–33
 firewalls, endpoints, 261
 float(n), 68
 footers, formatting, 423–28
 For Loop, 381
 FOR PATH, 101
 FOR PROPERTY, 101
 FOR VALUE, 101, 104
 FOR XML, 151
 Foreach Loop, 381
 forecasting algorithm, 474
 foreign keys, 83–84
 formatting
 date and time data, 124
 reports, SSRS, 423–28
 XML data queries, 151
 FORMSOF, 254
 forwarding pointers, 91
 fragmentation, 99–100, 337
 FREETEXT, 251–55
 FREETEXTTABLE, 251–55
 FROM, 138–39
 data retrieval, 122
 deleting data, 161–261

 functions, 204–11
 pivot tables, 145
 updating data, 159–61
 FROM clause, 131–33
 FTP tasks, 382
 FULL OUTER JOIN, 131–33
 Full recovery model, 309–11
 full text catalogs, 57, 223, 247–48
 full text filegroups, 56–57
 full text indexing, 249–51
 Database Snapshot, 220
 Dynamic Management Views (DMVs), 334
 file structure, 55
 filegroups, 57
 partitions, 112
 querying, 251–55
 Full Text Search, 17–19. *See also* full-text indexing
 functions, 199
 CLR security, 279–80
 creating, 200–3
 nesting in queries, 125
 retrieving data from, 204–11
 system functions, 199
 Fuzzy Grouping, 383
 Fuzzy Lookup, 383

G

geography data, 74, 102
 geometry data, 74, 102
 global variables (@@), 186–88
 GOTO, 192
 GROUP BY, 138–39, 142, 181
 grouping data, reports, 423–28
 GROUPING SETS, 142

H

hash algorithms, 282–84
 HAVING, 144–45, 181
 headers, formatting, 423–28
 health checks, 347
 heap, 91
 hidden tables, 228, 233–34
 hierarchies
 encryption, 281–82
 OLAP, 453–54, 467–70
 permissions, 274–77
 ROLLUP operator, 142
 hierarchyID, 75
 high availability operating mode, 351–52
 high performance operating mode,
 352–53
 high safety operating mode, 353
 histogram, 97
 HTML, full text indexes, 249
 HTTP, endpoints, 261

I

identifiers, object, 64
 IDENTITY, 75
 IF statements, 190–91
 image, 72, 104, 423–28
 IMAGE data type, 126, 249
 image functions, 200
 IMPERSONATE, 273
 Import Column, 384
 importing data, 34
 IN, 129
 IN PATH, 248
 included columns, indexes, 95–96
 INCREMENT, 75
 indexed views, 103–8, 182–84
 indexes

- administrative procedures, 196
- clustered indexes, 91–93
- covering indexes, 95–96
- Database Engine Tuning Advisor (DTA), 44–45
- dynamic management views (DMVs), 337–39
- filtered indexes, 97–98
- full text, 55, 57, 220, 249–51, 334
- full text catalogs, 247–48
- included columns, 95–96
- index key, 129
- maintenance plans, 325–28
- management and maintenance, 93–95, 99–100
- nonclustered, 93–95
- online index creation, 98–99
- partitions, 103–8, 110–16
- querying full text data, 251–55
- searching, 129
- spatial, 102
- structure, 87–91
- XML, 100–1

inequality, 127–28
 INFLECTIONAL, 254
 infrastructure requirements, 17
 INIT/NOINIT, 304
 initiator, communication, 227
 injection attack, SQL, 194
 inline table valued functions, 201
 INNER JOIN, 131–33, 149
 in-place upgrades, 20–21
 input parameters, 189
 INSERT, 153–59

- Database Snapshots, 222
- db_datawriter, 269
- db_denydatawriter, 269
- DML triggers, 213
- locks, 168
- MERGE, 164–66
- OUTPUT, 167
- performance degradation, 93–94
- permissions, 275
- SELECT INTO, 159
- transaction handling, 167–68

inserted tables, 166–67
 installing SQL Server 2008, 22–29

- authentication modes, 19
- collation sequences, 19
- edition descriptions, 15–16
- infrastructure requirements, 17
- instances, 20
- sample databases, 28–29
- service accounts, 17–19
- upgrading, 20–22

instance level principals, 264–68
 instances, 20, 216, 345–48. *See also* policy-based management
 instance-to-disk ratio, 346
 INSTEAD OF, 181, 213
 int, 68
 integers, 67–69
 Integration Services. *See* SQL Server Integration Server (SSIS)
 interactive elements, reports, 431–34
 INTERSECT, 148–49, 181, 183
 INITIATOR, 232
 INTO, 159, 165, 179
 IP addresses, 262–63, 345–48
 IsAlive test, 347
 ISNULL, 124, 130
 isolation levels, 170–71
 ISQL, 34

J

jobs, 319–25, 358
 join filter, 364
 join operators, 131–33, 144–45, 149, 204–5

K

kana sensitivity, 19
 KEY INDEX, 250
 key performance indicators (KPIs), 470
 keys

- Asymmetric Key login, 264–68
- asymmetric keys, 265, 281–82, 285–86
- clustering key, 91–93
- Encryption Keys, 415
- foreign keys, 83–84
- index key, 129
- KEY INDEX, 250
- management of, 260, 287
- master keys, 281–82
- partitioning key, 107, 113
- primary keys, 79–81, 92, 101, 113, 171
- private keys, 285
- public keys, 285–86
- symmetric keys, 284

KPIs (key performance indicators), 470

L

language files, 249, 471–72
 leaf page, indexes, 88–91, 93–95, 99–100
 LEFT OUTER JOIN, 131–33
 LIFETIME, 236
 LIKE operator, 129
 lines, formatting, 423–28
 linked reports, 443–44
 LISTENER_IP, 262–63
 LISTENER_PORT, 262–63
 local SID, 264–68
 local variables (@), 186–88
 lock escalation, 169–70
 locks
 partitions and, 112
 shared, 168
 transaction handling, 168–70
 log files
 performance, 309–10
 transaction log, 53
 Log Reader Agent, 365–66
 Log Sequence Number (LSN), 305–7, 355
 log shipping, 356
 components, 356–57
 database mirroring, 353–54
 initialization, 357–63
 overview, 7
 LOGIN, 202
 loginless users, 271
 logins
 administrative users, 269–73
 instance level principals, 264–68
 security of, 260
 logon triggers, 218
 Logread.exe, 365–66
 LooksAlive test, 347
 lookup tables, 84
 Lookup transform, 383
 LSN (Log Sequence Number), 173, 305–7

M

mail. *See* e-mail
 maintenance
 indexes, 99–100
 plans, creating, 325–28
 Manage Schedules, 320
 management
 disk storage, 54–55
 dynamic management views (DMVs)
 indexes, 337–39
 object size, 336
 overview, 333–34
 query execution stats, 339–40
 retrieving object metadata, 334–35
 encryption keys, 415
 indexes, 93–95, 99–100
 locks, 169–70
 membership roles and permissions, 269
 messaging process, 235–36
 partitions, 110–16
 performance, 44–46
 policy-based management
 compliance, 298
 conditions, 290–93
 facets, 290
 overview, 289–90
 policies, 295–98
 targets, 293–94
 transaction log data, 56
 Management Studio, SQL Server (SSMS),
 38–40, 290–93
 management system functions, 199
 management tools
 Bulk Copy Program (BCP), 34, 309–10, 366
 Database Mail, 42–43
 OSQL, 34
 Resource Governor, 35, 334
 SQL Server Configuration Manager, 35–38
 SQL Server Management Studio (SSMS), 38–40,
 290–93
 SQLCMD, 34
 SQLDiag, 34–35
 Tablediff, 34
 manipulating data. *See* data manipulation
 mantissa, 68
 master database, 52
 master keys, 281–82. *See also* keys
 materialized views (indexed views), 103–8, 182–84
 mathematical functions, 200
 matrices, formatting, 423–28
 MAX, 183, 383
 MAX_QUEUE_READERS, 242
 MAXSIZE, 55
 MDAC (Microsoft Data Access Components), 354
 MDX (Multidimensional Expression), 38
 measures, OLAP, 453–54
 memory, 35, 67
 memory buffers, 351
 MERGE, 110, 164–66
 Database Snapshots, 222
 db_datawriter, 269
 db_denydatawriter, 269
 inserted and deleted tables, 167
 performance degradation, 93–94
 Merge Agent, 366
 Merge Join, 384
 merge replication, 364, 367
 Merge transform, 384
 Message Queue, 382
 MESSAGE TYPE, 237
 MESSAGE_FORWARDING, 263
 MESSAGE_FORWARDING_SIZE, 263
 messages, 6, 42–43, 229–32
 architecture, 226–28
 contracts, 232–33
 conversations, 235–36

- message types, 229–32
- overview, 225–26
- prioritization, 245
- queues, 233–34, 242–44
- sending/receiving messages, 236–42
- services, 234–35
- metadata
 - database mirroring, 353–54
 - functions, 200
 - operation, 111
 - retrieving object metadata, 334–35
 - security, 277–78
 - SSRS, 415
- Microsoft Clustering, 474
- Microsoft Connect, 32–33
- Microsoft Customer Service and Support (CSS), 34–35
- Microsoft Data Access Components (MDAC), 354
- Microsoft Decision Trees, 474
- Microsoft Developer Network (MSDN), 32
- Microsoft Linear Regression, 474
- Microsoft Regression Trees, 474
- MIN, 183, 383
- mirror databases, 220
- mirror role, database, 349
- MIRROR TO, 304
- mirrored backup, 304
- model database, 52
- monetary data, 67–69
- monitor server, 357
- MOVE, 312
- moving, databases, 59–60
- msdb database, 52
- MSDN (Microsoft Developer Network), 32
- Multicast, 385
- Multidimensional Expression (MDX), 38, 470
- multi-platform processing, 225–26
- multiple-instance cluster, 345

N

- naming objects, 64
- Nave Bayes, 474, 483–84
- nchar(n), 70
- NEAR, 254
- network configuration, failover clustering, 346
- network names, failover clustering, 345–48
- Neural Networks, 474
- NEXT USED flag, 110
- NO REVERT, 202
- NODE, 101
- noise words, 249
- nonclustered indexes, 93–95
- NONE, validation option, 229
- NORECOVERY, 312
- NOT NULL, 75
- NOT, search queries, 130
- notifications, jobs, 321–25
- NTEXT, 70, 104, 126
- NTILE, 147
- NULL, 75
 - CUBE operator, 141
 - data retrieval, 124
 - functions, 201–2
 - search queries, 130
 - transforms, 384
 - XML data queries, 151
- nullability, 75
- numeric data, 67–69
- nvarchar(max), 70–71
- nvarchar(n), 70

O

- Object Explorer, 40, 95
- objects
 - dependent, 202
 - dynamic management views (DMVs), 336
 - metadata, retrieving, 334–35
 - naming, 64
 - permissions, 278–79
 - policy-based management
 - conditions, 290–93
 - facets, 290
 - overview, 289–90
 - policies, 295–97
 - policy categories, 297–98
 - policy compliance, 298
 - policy targets, 293–94
 - schemas, 273–74
- OFFLINE, indexes, 98–99
- OLAP engine (Online Analytic Processing). *See also* Database Engine
 - cubes, 454–62
 - dimensional model, 453–54
 - dimensions, measures, calculations, 462–66
 - hierarchies, 453–54, 467–70
 - key performance indicators (KPIs), 470
 - overview, 13–14, 452–53
 - partitions, 471
 - perspectives, 471
 - translations, 471–72
- OLE Automation Procedures, 260
- OLE DB, 385
- OLTP (online transaction processing), 13–14. *See also* Database Engine
- ON, 92
- ON ALL SERVER, 216
- ON clause, 106–7, 131–33
- ON DATABASE, 216
- Online Analytic Processing (OLAP)
 - cubes, 454–62
 - dimensional model, 453–54
 - dimensions, measures, calculations, 462–66
 - hierarchies, 453–54, 467–70
 - key performance indicators (KPIs), 470

- overview, 13–14, 452–53
- partitions, 471
- perspectives, 471
- translations, 471–72
- online indexes, 98–99
- online transaction processing (ONTP), 93–94.
 - See also* Database Engine
- ONLINE, indexes, 98–99
- OPEN, 194–96
- OPEN MASTER KEY, 282
- OPENDATASOURCE, 260
- OPENROWSET, 260
- operating system, 56–57, 260, 319–20
- operations, order of, 139
- operators, jobs, 321–25
- Optimizer, 97, 169–70, 201, 338
- OPTION, 179
- ORDER BY, 126, 130, 138, 146, 179, 183
- order of operations, 139
- OSQL, command line utility, 34
- OUT, 189
- outbound mail messages, 42–43. *See also* e-mail
- OUTPUT, 166–67, 188–89

P

- Package Explorer, SSIS, 380
- packages
 - configuration file, 406–9
 - connections, 386–90
 - control flow, 390–94
 - data flow, 394–403
 - deploying, 409–10
 - exception handling, 403–6
- page splitting, 90
- PageNumber, 423
- pages, 55
 - corruption of, 354
 - Database Snapshot, 220–21
 - fill factor, index, 99–100
 - full backups, 302–5
 - indexes, 88–91, 93–95
 - nonclustered indexes, 93–95
 - object size, 336
 - page corruption, 308–9
 - page-level locks, 169
 - partitioning, pointers, 111–13
- parameters, reports, 445
- parameters, stored procedures, 188–89
- parent-child relationship, 84
- parsing data, 123
- PARTITION BY, 146
- partitioning
 - Database Engine Tuning Advisor (DTA), 44–45
 - managing partitions, 110–16
 - overview, 103
 - partition functions, 103–5
 - partition schemes, 105
 - SSAS (SQL Server Analysis Services), 471
 - sys.partition, 334, 336
 - tables and indexes, 106–8
- partitioning key, 107, 113
- PARTNER, 263
- passwords
 - instance level principals, 264–68
 - security of, 260
 - SQL Server Configuration Manager, 35–38
- paste, Object Explorer, 95
- patches, 20
- PATH, 101, 151
- PATH XML, 101
- payloads, endpoints, 261
- percent (%), wildcard searches, 129
- Percentage transform, 385
- performance
 - alerts, creating, 329–32
 - backups, 304
 - cubes, processing, 471
 - database mirroring, 351–54
 - indexed views, 183
 - key performance indicators (KPIs), 470
 - locks, 169–70
 - logging, 309–10
 - management tools, 44–46
 - Performance Studio, 45–46
 - queries, 93–94, 133, 184, 204–5, 252
 - report caching, 444–49
 - report parameters, 434–38
 - Resource Governor, 35
 - searching, 129
 - SQL Server Profiler, 44
 - stored procedures, 340
- performance counters, 34–35
- Performance Data Warehouse, 45–46
- Performance Reports, 45–46
- Performance Studio, 45–46
- permissions, 274–77
 - access permissions, 279–80
 - directories, 347
 - objects, 278–79
 - service accounts, 347
- perspectives, 471
- phantom read, 170
- ping, 347
- PIVOT, 145, 183
- pivot tables, 145
- Pivot transform, 385
- pointers, 91, 111–13
- policies, 265, 295–97. *See also* policy-based management
- policy-based management
 - conditions, 290–93
 - facets, 290
 - overview, 289–90
 - policies, 295–97

port 1433

- policy categories, 297–98
- policy compliance, 298
- policy targets, 293–94
- port 1433, 262–63
- ports, 262–63
- PowerPoint documents, 249
- primary database, 357
- PRIMARY filegroup, 56–57
- primary keys, 79–81
 - change tracking, 171
 - clustered indexes, 92
 - partitions, 113
 - XML indexes, 101
- PRIMARY, XML indexes, 100–1
- principal role, database, 349
- principals
 - administrative users, 269–73
 - database level, 268–69
 - impersonation, 273
 - instance level, 264–68
- private keys, 285
- procedures, stored, 260, 277, 279–80
- processadmin, 266
- processing, multi-platform, 225–26
- processors, 35, 69
- Profiler, 44
- profiler traces, 34–35
- programming interface, 5–6
- PROPERTY, 101
- proximity searching, 254
- public key certificates, 285–86
- public keys, 285
- public, fixed database role, 269
- publications, 364
- publisher, 365

Q

- Qrdrsvc.exe, 366
- quality assurance, testing, 20
- queries. *See also* indexes
 - Data Mining Query, 380
 - Database Engine Tuning Advisor (DTA), 44–45
 - execution stats, 339–40
 - full text indexing, 251–55
 - functions, performance of, 204–5
 - indexed views, 184
 - multiple tables, 131–33
 - NULL values, 124, 130
 - Optimizer, 97
 - OSQL, command line utility, 34
 - performance of, 184
 - SQLCMD, 34
 - sub-queries, 133
 - views, substitution, 180
 - WHERE clause, 138
 - XML data, 150–51
- query cache, 340, 353

- query() method, 150
- Queue Reader Agent, 366
- queuing, messages, 6
 - architecture, 226–28
 - contracts, 232–33
 - conversations, 235–36
 - message types, 229–32
 - overview, 225–26
 - prioritization, 245
 - queues, 233–34, 242–44
 - sending/receiving messages, 236–42
 - services, 234–35

R

- RANGE LEFT, 104
- RANGE RIGHT, 104
- RANK, 146, 254
- ranking data, 146–47
- ranking functions, 200
- RAW, 151
- Read Committed isolation, 170
- Read Uncommitted isolation, 170
- read/write data files, 55
- real, 68
- RECEIVE, 237
- RECOVERY, 312
- recovery, data. *See also* database mirroring
 - backups, 301
 - database mirroring model, 355–56
 - database restores, 311–16
 - differential backups, 307–8, 314–15
 - filegroup backups, 308
 - full backups, 302–5, 311–12
 - page corruption, 308–9
 - recovery models, 309–11
 - transaction log backups, 305–7, 315–16
- redirection, clients, 354
- redundancy. *See* database mirroring; replication
- REFERENCES, 275
- Registered Servers, 39
- registry keys, 347
- regression algorithm, 473
- RELATED_CONVERSATION, 236
- RELATED_CONVERSATION_GROUP, 236
- Relational Engine, 8
- relational indexes, 91–95
- Remote Admin Connections, 260
- remote connections, 259–60
- Repeatable read isolation, 170
- replication, 363. *See also* database mirroring
 - agents, 365–66
 - caching, 353–54
 - components, 364
 - configuring, 367–73
 - creating jobs, 319–20
 - Database Engine, 6–7
 - methods, 366–67
 - roles, 365

ReplMerge.exe, 366
 Report Designer, 12
 Report Manager URL, 413
 reporting

- building reports, 418–23
- computations, 429–31
- configuring, 413–16
- deploying, 439–40
- formatting, 423–28
- interactive report elements, 431–34
- linked reports, 443–44
- parameters, 434–38
- Performance Studio, 45–46
- report caching and snapshots, 444–49
- SQL Server Management Studio (SSMS), 40
- subscriptions, 441–43
- Web site, 417

 Reporting Services. *See also* SQL Server Reporting Services

- ReportServer, 415
- ReportServerTemp, 415
- ReportServerTempDB, 444–49

 requirements, SQL Server 2008 installation, 17
 Resource Governor, 35, 334
 resources, 169–70, 242, 245
 restore

- database mirroring, 355–56
- job, 358
- log shipping, 358–59
- partial, 107

 RESTORE, 312
 RESTORE MASTER KEY, 282
 restore paths, 312–14
 RETENTION, 233, 237
 retrieving data

- aggregating data, 138–39
 - derived tables, 139–40
 - filtering aggregates, 143–44
 - multiple permutations, 140–42
 - pivot tables, calculating, 145
 - ranking data, 146–47
 - result sets, 147–49
 - running aggregates, 144–45
- common table expressions (CTEs), 149–50
- filtering data, 127–30
- functions, 204–11
- multiple tables, 131–33
- overview, 121
- SELECT statement, 126, 137–38
- sorting results, 126
- unique results, 134–35
- views, 179–84
- XML data querying, 150–51

 RETURN, 190
 return codes, 188
 RETURNS NULL ON NULL INPUT, 201–2

reverting data, Database Snapshots, 222–23
 RIGHT OUTER JOIN, 131–33
 ROLLBACK, 167
 ROLLBACK TRAN, 167–68
 ROLLUP operator, 142
 root pages, indexes, 88–91, 93–95
 Row Count transform, 385
 row filter, 364
 Row Sampling, 385
 ROW_NUMBER(), 146
 rows, table

- cursors, 194–96
- row-level locks, 169
- TRUNCATE, 164

 rowset functions, 200
 run-time activity, 337

S

SAFE, 279–80
 salting a hash, 284
 sample databases, 28–29
 SARG (search argument), 127–31
 scalability, dynamic, 242
 scalar functions, 200, 203
 scheduling, 6, 320. *See also* SQL Server Agent
 SCHEMABINDING, 183, 201–2
 schemas, 64–66, 72–74, 273–77, 366
 SCOPE_IDENTITY(), 158
 Script, 381
 script generators, 335
 Script transform, 385
 scripting tasks, 381
 SEARCH, 134–35
 search argument (SARG)

- filtering data, 127–30
- join clauses, 131
- WHERE clause, 138

 searching. *See also* full-text indexing

- character data, fragments, 129
- full text queries, 251–55
- Full Text Search, 17–19
- index pages, 88–91
- weighted proximity searches, 254
- wildcards, 253

 secondary database, 357
 SECONDARY, XML indexes, 100–1
 securables, 273–77
 security

- attack surface, configuring, 259–60
- authentication modes, 19
- backups, 301
- certificates, 285–86
- CLR security, 279–80
- cubes, OLAP, 454–62
- Database Engine, 4–5
- database level principals, 268–69

- database mirroring, 351–53
- encryption
 - asymmetric keys, 285–86
 - Extensible Key Management (EKM), 287
 - hash algorithms, 282–84
 - master keys, 281–82
 - overview, 280–81
 - symmetric keys, 284
 - TCP endpoints, 263
 - Transparent Data Encryption (TDE), 286–87
- endpoints, 260–63
- EXECUTE AS, 202
- failover clustering, 346–47
- impersonation, 273
- job steps, 320
- metadata, 277–78
- ownership chains, 277–79
- permissions, 274–77
- principals, 264–73
- reports, linked, 443–44
- schemas, 273–74
- securables, 273–74
- SQL Server Configuration Manager, 35–38
- security functions, 200
- Security Identifier (SID), 346
- securityadmin, 266
- SEED, 75
- SELECT
 - Database Snapshots, 222
 - db_datareader, 269
 - db_denydatareader, 269
 - DMVs (dynamic management views), 334
 - functions, retrieving data, 204–11
 - locks, 168
 - permissions, 275
 - recovering data, 312
 - script generators, 335
 - SELECT INTO, 159
 - variables, 187
- SELECT INTO, 139, 159, 309–10
- SELECT statements, 137–38
 - embedded (sub-queries), 133
 - filtering data, 127–30
 - general syntax, 121–26
 - multiple tables, data from, 131–33
 - sorting results, 126
 - views
 - creating views, 179–84
- self-signed certificates, 285
- SEND, 236–37
- Send Mail, 382
- Sequence, 381
- sequence analysis, 474
- Sequence Clustering, 474
- Serializable isolation level, 170
- server blocking, 34–35
- server configuration, 34–35
- server connections, 260–63
- Server Management Objects (SMO), 290
- server roles, 264–68
- serveradmin, 266
- service accounts
 - creating, 17–19
 - failover clustering, 346–48
 - SQL Server Configuration Manager, 35–38
- Service Broker
 - architecture, 226–28
 - contracts, 232–33
 - conversations, 235–36
 - endpoints, 263
 - message types, 229–32
 - overview, 6, 225–26
 - prioritization, 245
 - queues, 233–34, 242–44
 - sending/receiving messages, 236–42
 - services, 234–35
 - TCP endpoints, 262
- service master key, 281–82
- service packs, 20
- SERVICE_BROKER, 261
- SET, 159–61, 187, 192–93
- setupadmin, 266
- shared lock, 168
- SID, local, 264–68
- side-by-side upgrades, 21–22
- signatures, digital, 285–86
- Simple recovery model, 309–11
- single-instance cluster, 345
- Slowly Changing Dimensions, 385
- smalldatetime, 71–72
- smallint, 68
- smallmoney, 68
- SMO (Server Management Objects), 290
- Snapshot Agent, 365–67
- Snapshot isolation level, 170–71
- Snapshot.exe, 365
- snapshots, 221–22
 - Copy-On-Write, 220–21
 - Database Engine, 6–7
 - database mirroring, 349, 354
 - overview, 219–20
 - replication, 366–67
 - reports, 444–49
 - reverting data, 222–23
- SOAP, 261
- sorting, 126
- sp_configure, 260
- sp_executesql, 193–94
- space utilization, 336
- SPARSE, 78–79
- sparse files, 220
- SPATIAL data, 74–75
- spatial indexes, 102
- SPLIT operator, 110
- spreadsheets, 67
- SQL injection attack, 194

- SQL Mail XPs, 260
- SQL Server Agent
 - account, creating, 17–19
 - alerts, creating, 329–32
 - change tracking, 173
 - jobs, creating, 319–25
 - maintenance plans, creating, 325–28
 - msdb database, 52
 - overview, 6
 - schedule policies, 295
- SQL Server Analysis Services (SSAS), 13–14
 - Analysis Services Execute DDL Task, 380
 - creating account, 17–19
 - data mining, 472
 - algorithms, 473–74
 - demystified, 483–84
 - models and structures, 474–83
 - data warehousing, 451–52
 - installing, 22–29
- OLAP (Online Analytic Processing)
 - cubes, 454–62
 - dimensional model, 453–54
 - dimensions, measures, calculations, 462–66
 - hierarchies, 467–70
 - key performance indicators (KPIs), 470
 - overview, 452–53
 - partitions, 471
 - perspectives, 471
 - translations, 471–72
- SQL Server Compact Edition, 16
- SQL Server Configuration Manager, 35–38
- SQL Server Developer Edition, 16–17
- SQL Server Enterprise Edition, 16–17
- SQL Server Evaluation Edition, 16
- SQL Server Express Edition, 16
- SQL Server Integration Services (SSIS), 10–11
 - account, creating, 17–19
 - BIDS, 378–80
 - tasks, 380–82
 - transforms, 382–85
 - jobs, creating, 319–20
 - msdb database, 52
 - overview, 377
 - package building
 - configuration file, 406–9
 - connections, 386–90
 - control flow, 390–94
 - data flow, 394–403
 - deploying, 409–10
 - exception handling, 403–6
- SQL Server Management Studio (SSMS), 38–40, 290–93
- SQL Server Profiler, 44
- SQL Server Reporting Services (SSRS), 12–13
 - building reports, 418–23
 - computations, 429–31
 - configuring, 413–16
 - creating account, 17–19
 - deploying reports, 439–40
 - formatting, 423–28
 - installing, 22–29
 - interactive report elements, 431–34
 - linked reports, 443–44
 - parameters, 434–38
 - report caching and snapshots, 444–49
 - subscriptions, 441–43
 - Web site, 417
- SQL Server Standard Edition, 16
- SQL Server Workgroup Edition, 16
- SQL Trace Application Programming Interface (API), 44
- SQLCMD, 34
- SQLDiag, 34–35
- SSAS (SQL Server Analysis Services), 13–14
 - Analysis Services Execute DDL Task, 380–85
 - BIDS, 380
 - creating account, 17–19
 - data mining, 472
 - algorithms, 473–74
 - demystified, 483–84
 - models and structures, 474–83
 - data warehousing, 451–52
 - installing, 22–29
- OLAP (Online Analytic Processing)
 - cubes, 454–62
 - dimensional model, 453–54
 - dimensions, measures, calculations, 462–66
 - hierarchies, 467–70
 - key performance indicators (KPIs), 470
 - overview, 452–53
 - partitions, 471
 - perspectives, 471
 - translations, 471–72
- SSIS (SQL Server Integration Services)
 - BIDS, 378–80
 - tasks, 380–82
 - transforms, 382–85
 - creating account, 17–19
 - creating jobs, 319–20
 - installing, 22–29
 - msdb database, 52
 - overview, 10–11, 377
 - package, building
 - configuration file, 406–9
 - connection managers, 386–90
 - control flow, 390–94
 - data flow, 394–403
 - deploying, 409–10
 - exception handling, 403–6
- SSMS (SQL Server Management Studio), 38–40, 290–93
- SSMS Reports, 40
- SSRS (SQL Server Reporting Services), 12–13
 - BIDS, 378
 - building reports, 418–23
 - computations, 429–31
 - configuring, 413–16

- creating account, 17–19
- deploying reports, 439–40
- formatting, 423–28
- installing, 22–29
- interactive report elements, 431–34
- linked reports, 443–44
- parameters, 434–38
- report caching and snapshots, 444–49
- subscriptions, 441–43
- Web site, 417
- SSRS Configuration Manager, 413
- Standard SQL Server login, 264–68
- STANDBY, 312
- STARTED, 261–62
- startup mode, service, 35–38
- static computations, 205
- static row filter, 364
- statistics, 325–28, 337, 339–40. *See also* algorithms
- STATUS, 233
- STDEV, 183
- stemmers, 249
- stop words, 249
- STOP_ON_ERROR, 304
- STOPPAT, 315
- STOPPATMARK, 315
- STOPBEFOREMARK, 315
- STOPPED, endpoints, 261–62
- storage. *See also* message queuing
 - binary data, 72
 - character data, 69–71
 - data, architecture, 56
 - date and time data, 71–72
 - FILESTREAM data, 74
 - hierarchyID, 75
 - histograms, 97
 - index pages, 88–91, 93–95
 - management of, 54–55, 110–16
 - memory and, 67
 - numeric data, 67–69
 - object size, 336
 - partitions, 103–8, 110–16
 - sparse columns, 78–79
 - SPATIAL data, 74–75
 - temporary, 52–53
 - XML data, 72–74
- Storage Engine, 4
- stored procedures, 185
 - administrative procedures, 196
 - CLR procedures, 196
 - CLR security, 279–80
 - commenting code, 186
 - control flow constructs, 189–92
 - creating, 185
 - cursors, 194–96
 - dynamic execution, 193–94
 - error handling, 192–93
 - executing, 189
 - message queue, 242
 - ownership chains, 277
 - parameters, 188–89
 - performance, 340
 - return codes, 188
 - variables, 186–88
- string functions, 200
- striped backup, 303
- sub-queries, 133
- sub-reports, formatting, 423–28
- subscriber, 365
- subscriptions, reports, 441–43
- SUM, 383
- SWITCH operator, 111–13
- symmetric keys, 281–82, 284
- synchronizing files, 164–66. *See also* FILESTREAM
- sys.assembly, 333
- sys.change_tracking, 333
- sys.conversation, 333
- sys.database_files, 336
- sys.database_mirroring, 334
- sys.database_mirroring_witnesses, 349
- sys.dm.db, 334
- sys.dm_audit, 334
- sys.dm_broker, 334
- sys.dm_clr, 333
- sys.dm_db_index_operational_stats, 337
- sys.dm_db_index_physical_stats, 337
- sys.dm_db_index_usage_stats, 337
- sys.dm_db_missing_index, 338
- sys.dm_db_missing_index_group_stats, 338
- sys.dm_exec, 334
- sys.dm_exec_cached_plans, 340
- sys.dm_exec_connections, 340
- sys.dm_exec_procedure_stats, 340
- sys.dm_exec_query_plans, 340
- sys.dm_exec_query_stats, 340
- sys.dm_exec_requests, 340
- sys.dm_exec_sessions, 340
- sys.dm_exec_sql_text, 340
- sys.dm_fts, 334
- sys.dm_os, 334
- sys.dm_resource_governor, 334
- sys.dm_xe, 334
- sys.fulltext, 334
- sys.partition, 334, 336
- sys.resource_governor, 334
- sys.server_events, 334
- sys.service, 334
- sys.sp_cdc_enable_db, 173
- sys.trace, 334
- sys.xml, 334
- sysadmin, 266, 269–73
- system functions, 199–200
- system statistics, 200
- system stored procedures, 260

T

- Tablediff, 34
- tables. *See also* indexes; also Tablediff
 - change, 171–73
 - columns, 91–93, 95–96
 - creating, 76–77
 - data retrieval
 - aliases, 123
 - derived tables, 139–40
 - multiple tables, 131–33
 - SELECT, general syntax, 121–26
 - table joins, 131–33
 - deleted, 166–67
 - designing, 63–64
 - binary data, 72
 - character data, 69–71
 - columns, 75, 77–79
 - constraints, 79–84
 - data types, 67
 - database design, 64–67
 - database diagrams, 85–86
 - date and time data, 71–72
 - FILESTREAM data, 74
 - hierarchyID data, 75
 - naming objects, 64
 - numeric data, 67–69
 - schemas, 64–66
 - SPATIAL data, 74–75
 - XML data, 72–74
 - formatting reports, SSRS, 423–28
 - inline table valued functions, 201
 - inserted, 166–67
 - lookup, 84
 - partitions, 103–8, 110–16
 - pivot, calculating, 145
 - table-level locks, 169
 - temporary tables, 139–40
 - triggers
 - DDL triggers, 216, 218, 295
 - DML triggers, 213–15
 - views
 - creating views, 179–80
 - indexed views, 182–84
 - modifying data, 181–82
- table-valued functions, 203
- TAKE OWNERSHIP, 275
- TARGET, 232
- target, communication, 227
- TCP endpoints, 261–63
- TDE (Transparent Data Encryption), 286–87
- Template Explorer, 39
- templates, model database(s), 52
- temporary storage, 52–53
- temporary tables, 139–40
- Term Extraction, 385
- Term Lookup, 385
- terminating executions, 190
- testing, quality assurance, 20
- text, 70
 - full, 56–57
 - partitions, 104
 - report formatting, 423–28
- TEXT, 126, 249
- text functions, 200
- text index files, 55, 57, 112
- text mining, 385
- THESAURUS, 254
- thesaurus files, 251, 254
- time, 71–72, 124, 424
- time functions, 199
- Time Series, 474
- timestamp, 104
- tinyint, 68
- tmpdb database, 52–53
- TO clause, 105
- Toolbox, SSIS, 380
- tools
 - BI Dev Studio, 46
 - Bulk Copy Program (BCP), 34
 - Database Engine Tuning Advisor (DTA), 44–45
 - Database Mail, 42–43
 - documentation, 31–33
 - OSQL, 34
 - Performance Studio, 45–46
 - Resource Governor, 35
 - SQL Server Configuration Manager, 35–38
 - SQL Server Management Studio (SSMS), 38–40
 - SQL Server Profiler, 44
 - SQLCMD, 34
 - SQLDiag, 34–35
 - Tablediff, 34
- TOP, 183
- TOP (n), 237
- TOP operator, 130, 138, 179, 181
- TotalPages, 423
- Trace Application Programming Interface (API), 44
- traces, 334
- tracking changes, 171–75
- transaction handling, 167–71, 216
- transaction log, 351
 - backups, 305–7, 314, 325–28
 - database mirroring, 349
 - recovery models, 309–11
 - restoring backup, 315–16
- transaction log file, 53, 56, 223
- transactional replication, 367
- Transact-SQL (T-SQL), 5, 38, 319–20
- transforms, SSIS, 382–85
- translations, 471–72
- Transparent Data Encryption (TDE), 286–87
- transport, endpoints, 261
- triggers
 - CLR security, 279–80
 - DDL triggers, 216, 218, 295
 - DML triggers, 213–15

Triple DES

Triple DES, 281–82
 troubleshooting
 IF statements, 191
 SQLDiag, 34–35
 WHILE statements, 191
 TRUNCATE, 164
 TRUNCATE TABLE, 164
 TRY, 192–93
 T-SQL (Transact-SQL), 5, 262–63
 TYPE, 151
 TYPE COLUMN, 250

U

underscore (_), wildcard searches, 129
 Unicode data, 70, 252
 UNION, 147–49, 179, 181, 183
 UNION ALL, 147–49, 179, 181
 unique constraints, 81, 92
 unique indexes, 92
 uniqueidentifier, 236
 UNPIVOT, 183
 Unpivot transform, 385
 UNSAFE, 279–80
 unstructured data, 74
 UPDATE
 Database Snapshots, 222
 db_datawriter, 269
 db_denydatawriter, 269
 DML triggers, 213
 locks, 168
 MERGE, 164–66
 OUTPUT, 167
 performance degradation,
 93–94
 permissions, 275
 transaction handling,
 167–68
 update locks, 168
 updating data, 159–61
 upgrading, 20–22
 USER, 202
 USING...ON, 165
 utilities, 34

V

VALID_XML WITH SCHEMA COLLECTION, 229
 VALIDATION, 229
 VALUE, 101
 value() method, 216
 VALUES, 154
 VAR, 183
 VARBINARY, 72, 249
 VARBINARY(MAX), 104, 249–50
 varchar(max), 70–71, 104
 varchar(n), 70
 variables, stored procedures, 186–88

VARP, 183
 verb tenses, 249
 verification, 309
 VIEW DEFINITION, 275
 views
 dynamic management views (DMV)
 indexes, 337–39
 object size, 336
 overview, 333–34
 query execution stats, 339–40
 retrieving object metadata, 334–35
 indexed, 182–84
 modifying data, 181–82
 triggers
 DDL triggers, 216, 218
 DML triggers, 213–15
 virtual tables, 139–40
 Visual Studio, 378, 384

W

WAITFOR, 191
 WAITFOR DELAY, 191
 WAITFOR RECEIVE, 191
 WAITFOR TIME, 191
 Web Service URL, 413
 Web Services, 382
 web sites
 CodePlex, 28
 SSRS Web site, 417
 weighted proximity searches, 254
 WELL_FORMED_XML, 229
 WHEN, 165
 WHERE clause, 181
 conditions, 290–93
 data retrieval, 127–30
 filtering data, 133,
 143–44
 functions, retrieving data, 204–11
 join clauses, 131
 queries, 138
 updating data, 159–61
 WHILE, 191, 195
 wildcard characters, searching, 129, 253
 Windows accounts, 347
 Windows group, login, 264–68
 Windows login, 264–68
 Windows Management Instrumentation (WMI),
 329–32, 382
 Windows password policy, 265
 Windows performance counters,
 34–35
 Windows security API, 265
 Windows Server 2003, 17, 265
 Windows Server 2008, 17, 265
 Windows Service Control Applet, 35
 Windows versions, 19
 Windows Vista, 17

- Windows XP Professional, 17, 19
- WITH CHECK OPTION, 181
- WITH FORMAT, backup, 304
- WITNESS, 263
- witness server, 349, 352–53
- WMI (Windows Management Instrumentation), 329–32, 382
- WMI Data Reader task, 382
- WMI Event Watcher, 382
- word breakers, 249
- Word documents, 249
- words, full text searches, 251–55
- workload, 35
- write/read data files, 55

X

- X.509 standard, 285
- XACT_ABORT, 192–93
- XML (Extensible Markup Language), 5–6.
 - See* full-text indexing
 - BIDS tasks, 382
 - data querying, 150–51
 - data storage, 72–74
 - documents,
 - 72–73, 216
 - full text indexes, 249
 - indexes, 100–1
 - message types, 229
 - schemas, 72–74, 151
 - sorting results, 126
 - sys.xml, 334
 - thesaurus files, 254
- XML for Analysis (XMLA), 38
- XMLA (XML for Analysis), 38
- Xp_cmdshell, 260
- XQUERY, 216
- XSNIL directive, 151