

Microsoft®

MCTS EXAM

70-432

# Microsoft® SQL Server® 2008— Implementation and Maintenance



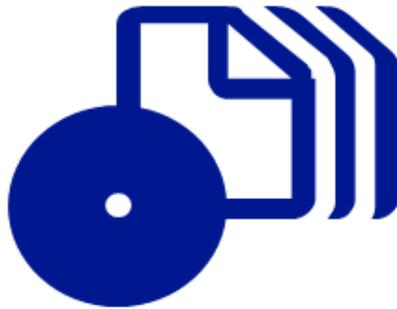
SELF-PACED

Mike Hotek

# Training Kit



# 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/626058/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](mailto:mspinput@microsoft.com)

Microsoft Press



# Exam 70-432: Microsoft SQL Server 2008— Implementation and Maintenance

OBJECTIVE	LOCATION IN BOOK
<b>INSTALLING AND CONFIGURING SQL SERVER 2008</b>	
Configure additional SQL Server components.	Chapter 1, Lessons 3 and 4 Chapter 5, Lessons 1, 2 and 3
Configure SQL Server instances.	Chapter 1, Lesson 3
Configure SQL Server services.	Chapter 1, Lesson 3
Install SQL Server 2008 and related services.	Chapter 1, Lesson 3 Chapter 5, Lessons 1, 2 and 3
Implement database mail.	Chapter 1, Lesson 4
Configure full-text indexing.	Chapter 5, Lessons 1, 2 and 3
<b>MAINTAINING SQL SERVER INSTANCES</b>	
Manage SQL Server Agent jobs.	Chapter 10, Lesson 2
Manage SQL Server Agent alerts.	Chapter 10, Lesson 4
Manage SQL Server Agent operators.	Chapter 10, Lesson 3
Implement the declarative management framework (DMF).	Chapter 8, Lessons 1 and 2
Back up a SQL Server environment.	Chapter 9, Lessons 1, 2 and 3
<b>MANAGING SQL SERVER SECURITY</b>	
Manage logins and server roles.	Chapter 11, Lesson 3
Manage users and database roles.	Chapter 11, Lesson 3
Manage SQL Server instance permissions.	Chapter 11, Lesson 4
Manage database permissions.	Chapter 11, Lesson 4
Manage schema permissions and object permissions.	Chapter 11, Lesson 4
Audit SQL Server instances.	Chapter 11, Lesson 5
Manage transparent data encryption.	Chapter 11, Lesson 6
Configure surface area.	Chapter 8, Lessons 1, 2 and 3 Chapter 11, Lesson 2

OBJECTIVE	LOCATION IN BOOK
<b>MAINTAINING A SQL SERVER DATABASE</b>	
Back up databases.	Chapter 2, Lesson 1 Chapter 9, Lesson 1
Restore databases.	Chapter 9, Lessons 2 and 3
Manage and configure databases.	Chapter 2, Lessons 2, 3 and 4
Manage database snapshots.	Chapter 9, Lesson 3
Maintain database integrity.	Chapter 2, Lesson 4
Maintain a database by using maintenance plans.	Chapter 9, Lesson 1
<b>PERFORMING DATA MANAGEMENT TASKS</b>	
Import and export data.	Chapter 7, Lessons 1, 2, 3 and 4
Manage data partitions.	Chapter 6, Lessons 1, 2, 3 and 4
Implement data compression.	Chapter 3, Lesson 1
Maintain indexes.	Chapter 4, Lesson 3 Chapter 5, Lessons 1, 2 and 3
Manage collations.	Chapter 2, Lesson 3
<b>MONITORING AND TROUBLESHOOTING SQL SERVER</b>	
Identify SQL Server service problems.	Chapter 12, Lesson 4
Identify concurrency problems.	Chapter 12, Lesson 2
Identify SQL Agent job execution problems.	Chapter 10, Lesson 1
Locate error information.	Chapter 12, Lesson 1
<b>OPTIMIZING SQL SERVER PERFORMANCE</b>	
Implement Resource Governor.	Chapter 13, Lesson 6
Use the Database Engine Tuning Advisor.	Chapter 13, Lesson 4
Collect trace data by using SQL Server Profiler.	Chapter 12, Lesson 2
Collect performance data by using Dynamic Management Views (DMVs).	Chapter 13, Lesson 5
Collect performance data by using System Monitor.	Chapter 12, Lesson 1
Use Performance Studio.	Chapter 13, Lesson 7
<b>IMPLEMENTING HIGH AVAILABILITY</b>	
Implement database mirroring.	Chapter 15, Lessons 1, 2 and 3
Implement a SQL Server clustered instance.	Chapter 14, Lessons 1 and 2
Implement log shipping.	Chapter 16, Lessons 1, 2 and 3
Implement replication.	Chapter 17, Lessons 1, 2 and 3

**Exam objectives** The exam objectives listed here are current as of this book's publication date. Exam objectives are subject to change at any time without prior notice and at Microsoft's sole discretion. Please visit the Microsoft Learning Web site for the most current listing of exam objectives: <http://www.microsoft.com/learning/mcp/>.

**MCTS Self-Paced Training  
Kit (Exam 70-432):  
Microsoft® SQL Server®  
2008—Implementation  
and Maintenance**

Training Kit

**Mike Hotek**

**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: 2008940530

Printed and bound in the United States of America.

ISBN: 978-0-7356-2605-8

5 6 7 8 9 10 11 12 13 14 QGT 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](http://www.microsoft.com/mspress). Send comments to [mspinput@microsoft.com](mailto:mspinput@microsoft.com).

Microsoft, Microsoft Press, Excel, IntelliSense, Internet Explorer, MSDN, MSN, SharePoint, Silverlight, SQL Server, Visual Studio, Windows, and Windows Server 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:** Laura Sackerman

**Project Editor:** Denise Bankaitis

**Editorial Production:** S4Carlisle Publishing Services

**Technical Reviewer:** Rozanne Murphy Whalen; Technical Review services provided by Content Master,  
a member of CM Group, Ltd.

**Cover:** Tom Draper Design

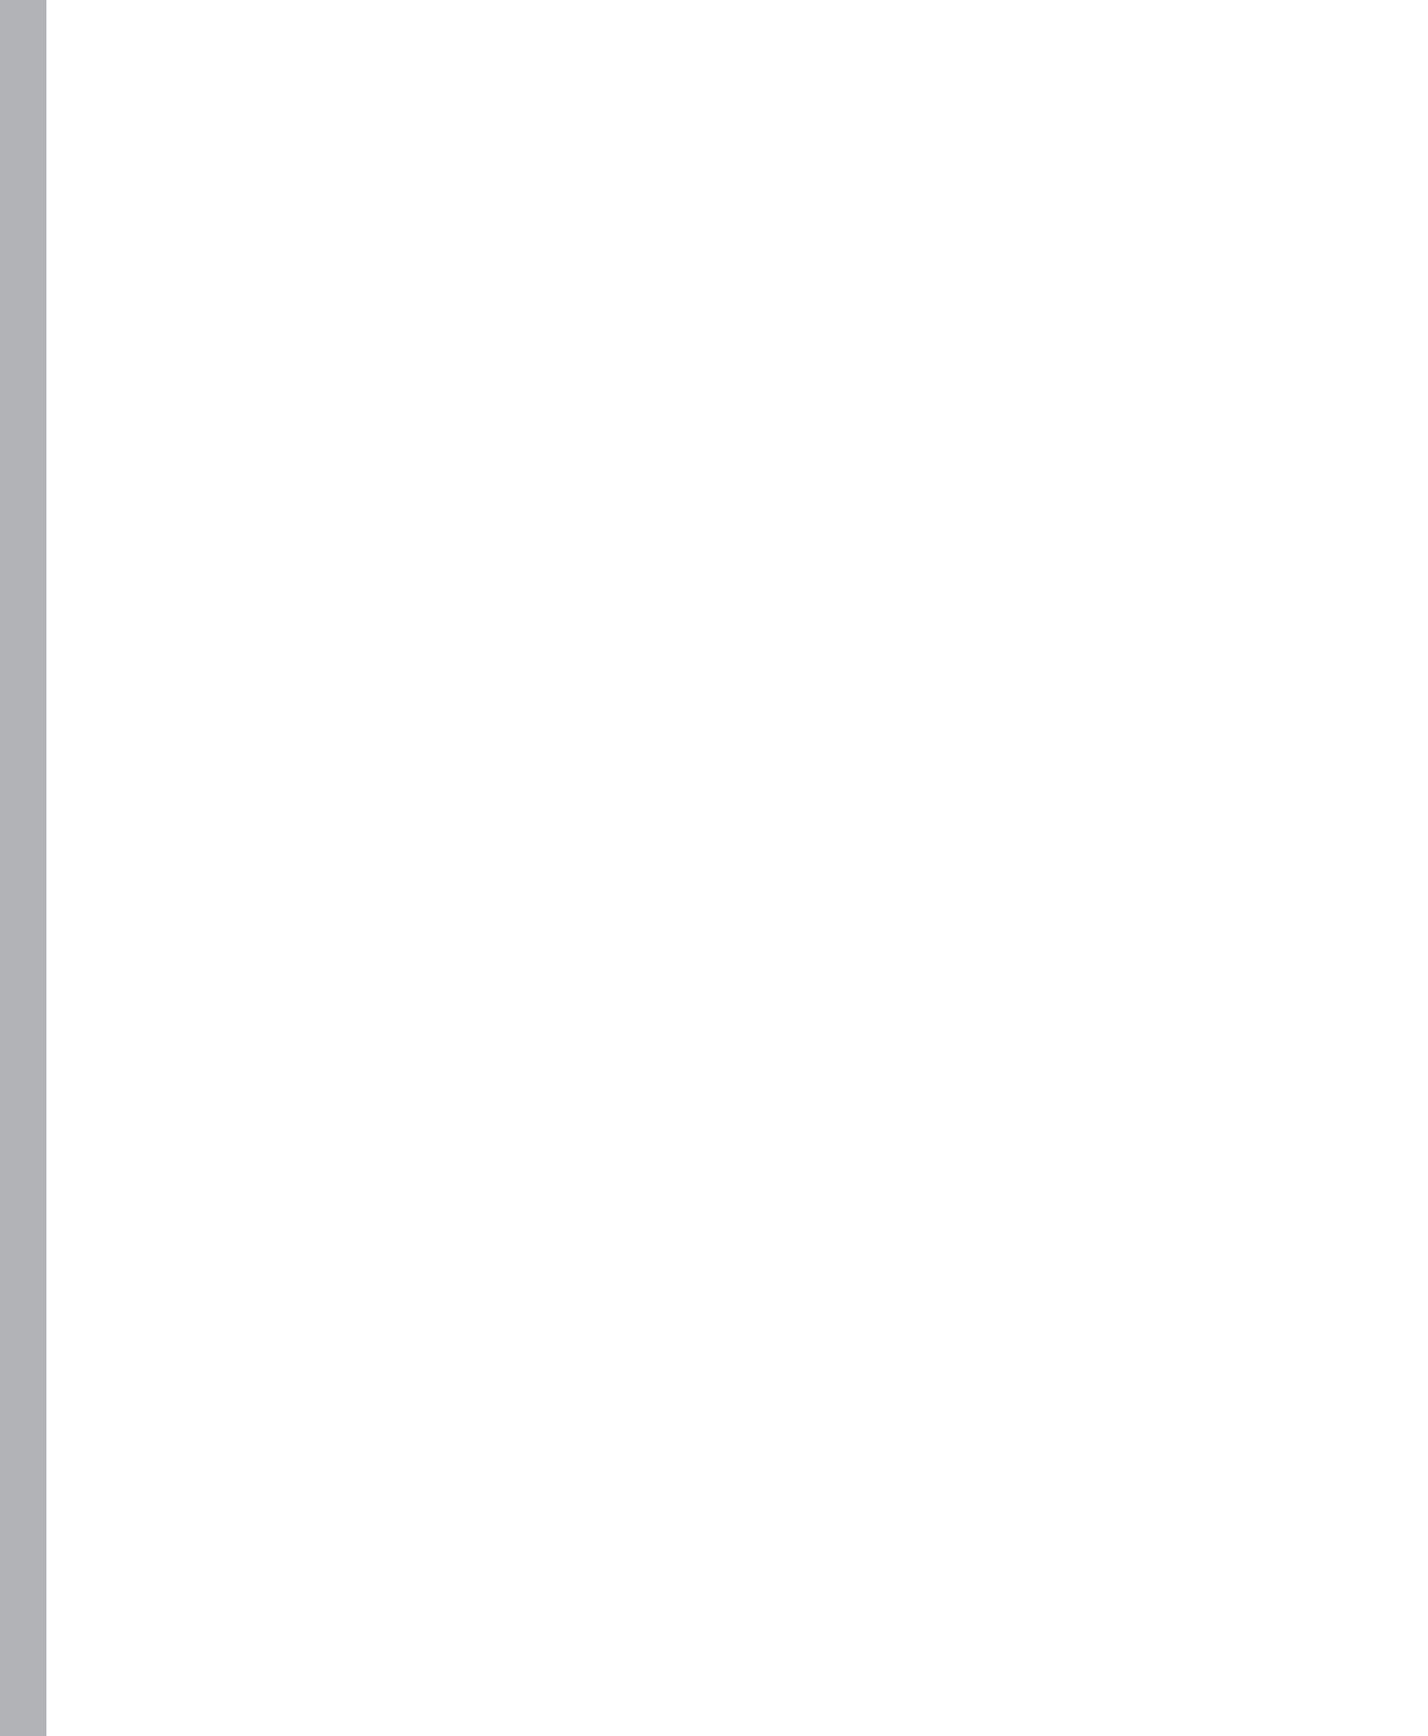
*To Genilyn,*

*My compass in a storm and the light to show me the way home*



# Contents at a Glance

	<i>Introduction</i>	<i>xxvii</i>
<b>CHAPTER 1</b>	<b>Installing and Configuring SQL Server 2008</b>	<b>1</b>
<b>CHAPTER 2</b>	<b>Database Configuration and Maintenance</b>	<b>37</b>
<b>CHAPTER 3</b>	<b>Tables</b>	<b>61</b>
<b>CHAPTER 4</b>	<b>Designing SQL Server Indexes</b>	<b>85</b>
<b>CHAPTER 5</b>	<b>Full Text Indexing</b>	<b>111</b>
<b>CHAPTER 6</b>	<b>Distributing and Partitioning Data</b>	<b>135</b>
<b>CHAPTER 7</b>	<b>Importing and Exporting Data</b>	<b>161</b>
<b>CHAPTER 8</b>	<b>Designing Policy Based Management</b>	<b>177</b>
<b>CHAPTER 9</b>	<b>Backing up and Restoring a Database</b>	<b>197</b>
<b>CHAPTER 10</b>	<b>Automating SQL Server</b>	<b>233</b>
<b>CHAPTER 11</b>	<b>Designing SQL Server Security</b>	<b>251</b>
<b>CHAPTER 12</b>	<b>Monitoring Microsoft SQL Server</b>	<b>307</b>
<b>CHAPTER 13</b>	<b>Optimizing Performance</b>	<b>367</b>
<b>CHAPTER 14</b>	<b>Failover Clustering</b>	<b>407</b>
<b>CHAPTER 15</b>	<b>Database Mirroring</b>	<b>451</b>
<b>CHAPTER 16</b>	<b>Log Shipping</b>	<b>483</b>
<b>CHAPTER 17</b>	<b>Replication</b>	<b>513</b>
	<i>Glossary</i>	<i>553</i>
	<i>Answers</i>	<i>561</i>
	<i>Index</i>	<i>599</i>



# Contents

<b>Introduction</b>	<b>xxvii</b>
<b>Chapter 1 Installing and Configuring SQL Server 2008</b>	<b>1</b>
Before You Begin . . . . .	1
Lesson 1: Determining Hardware and Software Requirements . . . . .	3
Minimum Hardware Requirements	3
Supported Operating Systems	4
Software Requirements	5
Lesson Summary	6
Lesson Review	6
Lesson 2: Selecting SQL Server Editions . . . . .	8
SQL Server Services	8
SQL Server Editions	11
Lesson Summary	15
Lesson Review	15
Lesson 3: Installing and Configuring SQL Server Instances . . . . .	17
Service Accounts	17
Collation Sequences	18
Authentication Modes	18
SQL Server Instances	19
SQL Server Configuration Manager	19
Lesson Summary	27
Lesson Review	27
Lesson 4: Configuring Database Mail . . . . .	28
Database Mail	28

**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/](http://www.microsoft.com/learning/booksurvey/)

Lesson Summary	31
Lesson Review	31
Chapter Review . . . . .	33
Chapter Summary	33
Key Terms	33
Case Scenario	33
Suggested Practices . . . . .	34
Installing SQL Server	34
Managing SQL Server Services	34
Take a Practice Test. . . . .	35
<b>Chapter 2 Database Configuration and Maintenance</b>	<b>37</b>
Before You Begin. . . . .	37
Lesson 1: Configuring Files and Filegroups . . . . .	39
Files and Filegroups . . . . .	39
Transaction Logs . . . . .	42
<b>FILESTREAM</b> data . . . . .	43
<b>tempdb</b> Database. . . . .	44
Lesson Summary	45
Lesson Review	45
Lesson 2: Configuring Database Options . . . . .	46
Database Options . . . . .	46
Recovery Options	46
Auto Options	48
Change Tracking	50
Access	50
Parameterization	51
Collation Sequences . . . . .	52
Lesson Summary	53
Lesson Review	53
Lesson 3: Maintaining Database Integrity . . . . .	54
Database Integrity Checks. . . . .	54
Lesson Summary	56

Lesson Review	56
Chapter Review	57
Chapter Summary	57
Key Terms	57
Case Scenario	57
Suggested Practices	59
Configuring Databases	59
Take a Practice Test	60

## **Chapter 3 Tables 61**

Before You Begin	61
Lesson 1: Creating Tables	63
Schemas	63
Data Types	64
Column Properties	69
Computed Columns	72
Row and Page Compression	72
Creating Tables	73
Lesson Summary	75
Lesson Review	76
Lesson 2: Implementing Constraints	77
Primary Keys	77
Foreign Keys	77
Unique Constraints	78
Default Constraints	78
Check Constraints	78
Lesson Summary	80
Lesson Review	80
Chapter Review	81
Chapter Summary	81
Key Terms	81
Case Scenario	81
Suggested Practices	82
Creating Tables	82

Creating Constraints	83
Take a Practice Test	83
<b>Chapter 4 Designing SQL Server Indexes</b>	<b>85</b>
Before You Begin	85
Lesson 1: Index Architecture	87
Index Structure	87
Balanced Trees (B-Trees)	88
Index Levels	89
Lesson Summary	91
Lesson Review	91
Lesson 2: Designing Indexes	93
Clustered Indexes	93
Nonclustered Indexes	95
Index Options	97
XML Indexes	99
Spatial Indexes	99
Lesson Summary	102
Lesson Review	102
Lesson 3: Maintaining Indexes	104
Index Management and Maintenance	104
Disabling an index	105
Lesson Summary	107
Lesson Review	107
Chapter Review	108
Chapter Summary	108
Key Terms	108
Case Scenario	108
Suggested Practices	110
Creating Indexes	110
Take a Practice Test	110
<b>Chapter 5 Full Text Indexing</b>	<b>111</b>
Before You Begin	111

Lesson 1: Creating and Populating Full Text Indexes . . . . .	113
Full Text Catalogs	113
Full Text Indexes	114
Change Tracking	115
Language, Word Breakers, and Stemmers	116
Lesson Summary	118
Lesson Review	118
Lesson 2: Querying Full Text Data . . . . .	120
FREETEXT	120
CONTAINS	121
Lesson Summary	125
Lesson Review	126
Lesson 3: Managing Full Text Indexes . . . . .	126
Thesaurus	127
Stop Lists	128
Populate Full Text Indexes	128
Lesson Summary	131
Lesson Review	131
Chapter Review . . . . .	131
Chapter Summary	132
Key Terms	132
Case Scenario	132
Suggested Practices . . . . .	133
Create a Full Text Index	133
Query a Full Text Index	133
Create a Thesaurus File	133
Create a Stop List	134
Take a Practice Test. . . . .	134

**Chapter 6 Distributing and Partitioning Data 135**

Before You Begin. . . . .	135
Lesson 1: Creating a Partition Function . . . . .	137
Partition Functions	137
Lesson Summary	141

Lesson Review	141
Lesson 2: Creating a Partition Scheme	142
Partition Schemes	142
Lesson Summary	144
Lesson Review	144
Lesson 3: Creating Partitioned Tables and Indexes	146
Creating a Partitioned Table	146
Creating a Partitioned Index	147
Lesson Summary	149
Lesson Review	149
Lesson 4: Managing Partitions	150
Split and Merge Operators	150
Altering a Partition Scheme	150
Index Alignment	151
Switch Operator	151
Lesson Summary	156
Lesson Review	156
Chapter Review	157
Chapter Summary	157
Key Terms	157
Case Scenario	157
Suggested Practices	160
Partitioning	160
Take a Practice Test	160
<b>Chapter 7 Importing and Exporting Data</b>	<b>161</b>
Before You Begin	161
Lesson 1: Importing and Exporting Data	163
Bulk Copy Program	163
The <b>BULK INSERT</b> command	165
The SQL Server Import and Export Wizard	166
Lesson Summary	172
Lesson Review	172

Chapter Review . . . . .	173
Chapter Summary . . . . .	173
Key Terms . . . . .	173
Case Scenario . . . . .	173
Suggested Practices . . . . .	175
Import and Export Data . . . . .	175
Take a Practice Test. . . . .	176
<b>Chapter 8 Designing Policy Based Management</b>	<b>177</b>
Before You Begin. . . . .	177
Lesson 1: Designing Policies . . . . .	179
Facets . . . . .	179
Conditions . . . . .	180
Policy Targets. . . . .	180
Policies . . . . .	181
Policy Categories. . . . .	181
Policy Compliance. . . . .	181
Central Management Server . . . . .	183
Import and Export Policies . . . . .	183
Lesson Summary . . . . .	191
Lesson Review . . . . .	191
Chapter Review . . . . .	193
Chapter Summary . . . . .	193
Key Terms . . . . .	193
Case Scenario . . . . .	193
Suggested Practices . . . . .	196
Implement Policy Based Management . . . . .	196
Take a Practice Test. . . . .	196
<b>Chapter 9 Backing up and Restoring a Database</b>	<b>197</b>
Before You Begin. . . . .	197
Lesson 1: Backing up Databases. . . . .	199

Backup Types . . . . .	199
Full Backups . . . . .	200
Transaction Log Backups . . . . .	203
Differential Backups . . . . .	204
Filegroup Backups . . . . .	205
Partial Backups . . . . .	205
Page Corruption . . . . .	206
Maintenance Plans . . . . .	206
Certificates and Master Keys . . . . .	207
Validating a Backup . . . . .	209
Lesson Summary . . . . .	211
Lesson Review . . . . .	211
Lesson 2: Restoring Databases . . . . .	212
Transaction Log Internals . . . . .	212
Database Restores . . . . .	214
Restoring a Full Backup . . . . .	214
Restoring a Differential Backup . . . . .	216
Restoring a Transaction Log Backup . . . . .	216
Online Restores . . . . .	217
Restore a Corrupt Page . . . . .	217
Restoring with Media Errors . . . . .	218
Lesson Summary . . . . .	222
Lesson Review . . . . .	222
Lesson 3: Database Snapshots . . . . .	223
Creating a Database Snapshot . . . . .	223
Copy-On-Write Technology . . . . .	224
Reverting Data Using a Database Snapshot . . . . .	225
Lesson Summary . . . . .	227
Lesson Review . . . . .	227
Chapter Review . . . . .	228
Chapter Summary . . . . .	228
Key Terms . . . . .	228
Case Scenario . . . . .	229

Suggested Practices . . . . .	231
Backing up a Database . . . . .	231
Restoring a Database . . . . .	231
Take a Practice Test. . . . .	231
<b>Chapter 10 Automating SQL Server</b>	<b>233</b>
Before You Begin. . . . .	233
Lesson 1: Creating Jobs . . . . .	234
Job Steps . . . . .	234
Job Schedules . . . . .	235
Job History . . . . .	235
Operators . . . . .	236
Lesson Summary . . . . .	240
Lesson Review . . . . .	240
Lesson 2: Creating Alerts . . . . .	242
SQL Server Agent Alerts . . . . .	242
Lesson Summary . . . . .	245
Lesson Review . . . . .	245
Chapter Review . . . . .	246
Chapter Summary . . . . .	246
Key Terms . . . . .	246
Case Scenario . . . . .	246
Suggested Practices . . . . .	248
Create Jobs . . . . .	248
Create Alerts . . . . .	248
Take a Practice Test. . . . .	249
<b>Chapter 11 Designing SQL Server Security</b>	<b>251</b>
Before You Begin. . . . .	251
Lesson 1: TCP Endpoints . . . . .	252
Endpoint Types and Payloads . . . . .	252
Endpoint Access . . . . .	253

TCP Endpoints .....	253
TCP Protocol Arguments	253
Database Mirroring and Service Broker Common Arguments	254
Database Mirroring–Specific Arguments	255
Service Broker–Specific Arguments	255
Lesson Summary	257
Lesson Review	257
Lesson 2: Configuring the SQL Server Surface Area .....	259
Surface Area Configuration	259
Lesson Summary	261
Lesson Review	262
Lesson 3: Creating Principals .....	263
Logins	263
Fixed Server Roles	265
Database Users	266
Loginless Users	266
Fixed Database Roles	267
User Database Roles	268
Lesson Summary	269
Lesson Review	270
Lesson 4: Managing Permissions .....	271
Securables	272
Permissions	272
Metadata Security	273
Ownership Chains	274
Impersonation	275
Master Keys	275
Certificates	276
Signatures	277
Lesson Summary	283
Lesson Review	283
Lesson 5: Auditing SQL Server Instances .....	285
DDL Triggers	285
Audit Specifications	286

C2 Auditing	288
Lesson Summary	290
Lesson Review	291
Lesson 6: Encrypting Data . . . . .	292
Data Encryption	292
Hash Algorithms	293
Symmetric Keys	294
Certificates and Asymmetric Keys	294
Transparent Data Encryption	294
Encryption Key Management	296
Lesson Summary	300
Lesson Review	300
Chapter Review . . . . .	302
Chapter Summary	302
Key Terms	302
Case Scenario: Designing SQL Server Security	303
Suggested Practices . . . . .	304
Manage Logins and Server Roles	304
Manage Users and Database Roles	305
Manage SQL Server Instance Permissions	305
Manage Database Permissions	305
Manage Schema Permissions and Object Permissions	305
Audit SQL Server Instances	305
Manage Transparent Data Encryption	305
Take a Practice Test. . . . .	305

**Chapter 12 Monitoring Microsoft SQL Server 307**

Before You Begin. . . . .	307
Lesson 1: Working with System Monitor . . . . .	309
System Monitor Overview. . . . .	309
Capturing Counter Logs. . . . .	310
Performance Counters. . . . .	312
Lesson Summary	315
Lesson Review	315

Lesson 2: Working with the SQL Server Profiler . . . . .	317
Defining a Trace . . . . .	317
Specifying Trace Events . . . . .	320
Selecting Data Columns. . . . .	322
Applying Filters . . . . .	323
Managing Traces. . . . .	324
Correlating Performance and Monitoring Data . . . . .	325
Lesson Summary	330
Lesson Review	331
Lesson 3: Diagnosing Database Failures . . . . .	332
SQL Server Logs. . . . .	332
Database Space Issues . . . . .	334
Lesson Summary	338
Lesson Review	339
Lesson 4: Diagnosing Service Failures . . . . .	340
Finding Service Startup Failures . . . . .	340
Configuration Manager	340
Lesson Summary	349
Lesson Review	349
Lesson 5: Diagnosing Hardware Failures. . . . .	351
Disk Drives . . . . .	351
Memory and Processors . . . . .	352
Lesson Summary	353
Lesson Review	353
Lesson 6: Resolving Blocking and Deadlocking Issues . . . . .	355
Locks . . . . .	355
Transaction Isolation Levels. . . . .	356
Blocked Processes . . . . .	357
Deadlocks. . . . .	358
Lesson Summary	362
Lesson Review	362

Chapter Review . . . . .	363
Chapter Summary . . . . .	363
Key Terms . . . . .	363
Case Scenario . . . . .	363
Suggested Practices . . . . .	366
Creating a Trace Using SQL Server Profiler to Diagnose Performance and Deadlock Issues . . . . .	366
Create a Counter Log Using System Monitor to Diagnose Performance, Deadlock, and System Issues . . . . .	366
Take a Practice Test . . . . .	366

## **Chapter 13 Optimizing Performance 367**

Before You Begin . . . . .	367
Lesson 1: Using the Database Engine Tuning Advisor . . . . .	369
Database Engine Tuning Advisor . . . . .	369
Lesson Summary . . . . .	375
Lesson Review . . . . .	375
Lesson 2: Working with Resource Governor . . . . .	376
Resource Governor . . . . .	376
Lesson Summary . . . . .	386
Lesson Review . . . . .	386
Lesson 3: Using Dynamic Management Views and Functions . . . . .	387
DMV Categories . . . . .	387
Database Statistics . . . . .	388
Query Statistics . . . . .	389
Disk Subsystem Statistics . . . . .	390
Hardware Resources . . . . .	391
Lesson Summary . . . . .	393
Lesson Review . . . . .	394
Lesson 4: Working with the Performance Data Warehouse . . . . .	395
Performance Data Warehouse . . . . .	395
Lesson Summary . . . . .	400
Lesson Review . . . . .	400

Chapter Review .....	402
Chapter Summary	402
Key Terms	402
Case Scenario	403
Suggested Practices .....	405
Using the Performance Data Warehouse to Gather Data for Performance Optimization	405
Using Database Engine Tuning Advisor to Gather Data for Performance Optimization	405
Using Dynamic Management Views to Gather Data for Performance Optimization	405
Take a Practice Test.....	406

## **Chapter 14 Failover Clustering 407**

Before You Begin.....	407
Lesson 1: Designing Windows Clustering.....	410
Windows Cluster Components	410
Types of Clusters	412
Security Configuration	413
Disk Configuration	413
Network Configuration	414
Cluster Resources	415
Cluster Groups	416
Lesson Summary	428
Lesson Review	429
Lesson 2: Designing SQL Server 2008 Failover Cluster Instances.....	430
Terminology	431
Failover Cluster Instance Components	431
Health Checks	433
Cluster Failover	433
Lesson Summary	444
Lesson Review	444
Chapter Review .....	445
Chapter Summary	445
Key Terms	445

Case Scenario	445
Suggested Practices	448
Windows Clustering	448
SQL Server Failover Clustering	449
Take a Practice Test	449
<b>Chapter 15 Database Mirroring</b>	<b>451</b>
Before You Begin	451
Lesson 1: Overview of Database Mirroring	452
Database Mirroring Roles	452
Principal Role	453
Mirror Role	453
Witness Server	453
Database Mirroring Endpoints	454
Operating Modes	455
Caching	458
Transparent Client Redirect	458
Database Mirroring Threading	459
Database Snapshots	459
Lesson Summary	462
Lesson Review	462
Lesson 2: Initializing Database Mirroring	464
Recovery Model	465
Backup and Restore	465
Copy System Objects	466
Lesson Summary	469
Lesson Review	469
Lesson 3: Designing Failover and Failback Strategies	471
Designing Mirroring Session Failover	471
Designing Mirroring Session Failback	472
Lesson Summary	474
Lesson Review	475

Chapter Review .....	476
Chapter Summary	476
Key Terms	476
Case Scenario	477
Suggested Practices .....	480
Establishing Database Mirroring	480
Creating a Database Snapshot Against a Database Mirror	480
Take a Practice Test.....	481

## **Chapter 16 Log Shipping 483**

Before You Begin.....	483
Lesson 1: Overview of Log Shipping .....	484
Log Shipping Scenarios	484
Log Shipping Components	485
Types of Log Shipping	487
Lesson Summary	487
Lesson Review	488
Lesson 2: Initializing Log Shipping.....	489
Log Shipping Initialization	489
Lesson Summary	499
Lesson Review	499
Lesson 3: Designing Failover and Failback Strategies.....	500
Log Shipping Failover	501
Log Shipping Failback	502
Lesson Summary	505
Lesson Review	505
Chapter Review .....	506
Chapter Summary	506
Key Terms	506
Case Scenario	507
Suggested Practices .....	511
Initiating Log Shipping	511
Failover and Failback Log Shipping	512
Take a Practice Test.....	512

<b>Chapter 17 Replication</b>	<b>513</b>
Before You Begin.....	513
Lesson 1: Overview of Replication.....	514
Replication Components	514
Replication Roles	515
Replication Topologies	516
Replication Agents	517
Agent Profiles	518
Replication Methods	519
Data Conflicts	521
Lesson Summary	524
Lesson Review	525
Lesson 2: Transactional Replication.....	526
Change Tracking	526
Transactional Options	528
Transactional Architectures	530
Monitoring	532
Validation	532
Lesson Summary	536
Lesson Review	537
Lesson 3: Merge Replication.....	538
Change Tracking	538
Validation	541
Lesson Summary	543
Lesson Review	543
Chapter Review.....	545
Chapter Summary	545
Key Terms	545
Case Scenario	546
Suggested Practices.....	550
Transactional Replication	550
Merge Replication	551
Take a Practice Test.....	552

<i>Glossary</i>	553
<i>Answers</i>	561
<i>Index</i>	599

# Acknowledgements

---

Thank you to all my readers over the past decade or so; it's hard to believe that this will be the seventh SQL Server book I've written and it would not be possible without you. I'd like to thank my editorial team at Microsoft Press, Denise Bankaitis and Laura Sackerman. I would especially like to thank Ken Jones, who has gone through five books with me and has proved to be an invaluable asset to Microsoft Press. Thank you to Rozanne Whalen, who has now tech-edited three books for me. I don't know how she does it, but Susan McClung's word wizardry has transformed my writing into the volume you hold in your hands. That all of this content is coherent is a testament to the many hours of hard work put in by Rozanne, Susan, and the rest of the editing team.



# Introduction

---

This training kit is designed for information technology (IT) professionals who plan to take the Microsoft Certified Technology Specialist (MCTS) Exam 70-432, as well as database administrators (DBAs) who need to know how to implement, manage, and troubleshoot Microsoft SQL Server 2008 instances. It's assumed that before using this training kit, you already have a working knowledge of Microsoft Windows and SQL Server 2008, and you have experience with SQL Server or another database platform.

By using this training kit, you learn how to do the following:

- Install and configure SQL Server 2008
- Create and implement database objects
- Implement high availability and disaster recovery
- Secure instances, databases, and database objects
- Monitor and troubleshoot SQL Server instances

## Using the CD and DVD

A companion CD and an evaluation software DVD are included with this training kit.

The companion CD contains the following:

- **Practice tests** You can reinforce your understanding of how to implement and maintain SQL Server 2008 databases by using electronic practice tests that you can customize to meet your needs from the pool of Lesson Review questions in this book. Alternatively, you can practice for the 70-432 certification exam by using tests created from a pool of about 200 realistic exam questions, which will give you enough different practice tests to ensure that you're prepared.
- **Practice files** Not all exercises incorporate code, but for each exercise that has code, you can find one or more files in a folder for the corresponding chapter on the companion CD. You can either type the code from the book or open the corresponding code file in a query window.
- **eBook** An electronic version (eBook) of this training kit is included for use at times when you don't want to carry the printed book with you. The eBook is in Portable Document Format (PDF), and you can view it by using Adobe Acrobat or Adobe Reader. You can use the eBook to cut and paste code as you work through the exercises.
- **Sample chapters** Sample chapters from other Microsoft Press titles on SQL Server 2008. These chapters are in PDF format.

- **Evaluation software** The evaluation software DVD contains a 180-day evaluation edition of SQL Server 2008 in case you want to use it instead of a full version of SQL Server 2008 to complete the exercises in this book.

**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/9780735626058> to get your downloadable content. This content is always up-to-date and available to all readers.

## How to Install the Practice Tests

To install the practice test software from the companion CD to your hard disk, perform the following steps:

1. Insert the companion CD into your CD-ROM drive and accept the license agreement that appears onscreen. A CD menu appears.

### **NOTE ALTERNATIVE INSTALLATION INSTRUCTIONS IF AUTORUN IS DISABLED**

**If the CD menu or the license agreement doesn't appear, AutoRun might be disabled on your computer. Refer to the Readme.txt file on the companion CD for alternative installation instructions.**

2. Click Practice Tests and follow the instructions on the screen.

## How to Use the Practice Tests

To start the practice test software, follow these steps:

1. Click Start and select All Programs, Microsoft Press Training Kit Exam Prep. A window appears that shows all the Microsoft Press training kit exam prep suites that are installed on your computer.
2. Double-click the lesson review or practice test that you want to use.

## Lesson Review Options

When you start a lesson review, the Custom Mode dialog box appears, enabling you to configure your test. You can click OK to accept the defaults, or you can customize the number of questions you want, the way the practice test software works, which exam objectives you want the questions to relate to, and whether you want your lesson review to be timed. If you are retaking a test, you can select whether you want to see all the questions again or only those questions you previously skipped or answered incorrectly.

After you click OK, your lesson review starts. You can take the test by performing the following steps:

1. Answer the questions and use the Next, Previous, and Go To buttons to move from question to question.
2. After you answer an individual question, if you want to see which answers are correct, along with an explanation of each correct answer, click Explanation.
3. If you would rather wait until the end of the test to see how you did, answer all the questions and then click Score Test. You see a summary of the exam objectives that you chose and the percentage of questions you got right overall and per objective. You can print a copy of your test, review your answers, or retake the test.

## Practice Test Options

When you start a practice test, you can choose whether to take the test in Certification Mode, Study Mode, or Custom Mode.

- **Certification Mode** Closely resembles the experience of taking a certification exam. The test has a set number of questions, it is timed, and you cannot pause and restart the timer.
- **Study Mode** Creates an untimed test in which you can review the correct answers and the explanations after you answer each question.
- **Custom Mode** Gives you full control over the test options so that you can customize them as you like.

In all modes, the user interface that you see when taking the test is basically the same, but different options are enabled or disabled, depending on the mode. The main options are discussed in the previous section, "Lesson Review Options."

When you review your answer to an individual practice test question, a "References" section is provided. This section lists the location in the training kit where you can find the information that relates to that question, and it provides links to other sources of information. After you click Test Results to score your entire practice test, you can click the Learning Plan tab to see a list of references for every objective.

## How to Uninstall the Practice Tests

To uninstall the practice test software for a training kit, use the Add Or Remove Programs option (Windows XP or Windows Server 2003) or the Program And Features option (Windows Vista or Windows Server 2008) in Control Panel.

## Microsoft Certified Professional Program

Microsoft certifications provide the best method to prove your command of current Microsoft products and technologies. The exams and corresponding certifications are developed to validate your mastery of critical competencies as you design and develop or implement and support solutions with Microsoft products and technologies. Computer professionals who become Microsoft-certified are recognized as experts and are sought after industry-wide. Certification brings a variety of benefits to the individual and to employers and organizations.

### **MORE INFO LIST OF MICROSOFT CERTIFICATIONS**

For a full list of Microsoft certifications, go to <http://www.microsoft.com/learning/mcp/default.aspx>.

## Technical Support

Every effort has been made to ensure the accuracy of this book and the contents of the companion CD. If you have comments, questions, or ideas regarding this book or the companion CD, please send them to Microsoft Press by using either of the following methods:

### **E-mail**

- [tkinput@microsoft.com](mailto:tkinput@microsoft.com)

### **Postal Mail:**

- *Microsoft Press*

*Attn: MCTS Self-Paced Training Kit (Exam 70-432): Microsoft SQL Server 2008 Implementation and Maintenance Editor*

*One Microsoft Way*

*Redmond, WA, 98052-6399*

For additional support information regarding this book and the companion CD (including answers to commonly asked questions about installation and use), visit the Microsoft Press Technical Support Web site at <http://www.microsoft.com/learning/support/books>. To connect directly to the Microsoft Knowledge Base and enter a query, visit <http://support.microsoft.com/search>. For support information regarding Microsoft software, please connect to <http://support.microsoft.com>.

## Evaluation Edition Software

The 180-day evaluation edition provided with this training kit is not the full retail product and is provided only for the purposes of training and evaluation. Microsoft and Microsoft Technical Support do not support this evaluation edition.

Information about any issues relating to the use of this evaluation edition with this training kit is posted in the Support section of the Microsoft Press Web site (<http://www.microsoft.com/learning/support/books/>). For information about ordering the full version of any Microsoft software, please call Microsoft Sales at (800) 426-9400 or visit <http://www.microsoft.com>.

# Database Configuration and Maintenance

The configuration choices that you make for a database affect its performance, scalability, and management. In this chapter, you learn how to design the file and filegroup storage structures underneath a database. You learn how to configure database options and recovery models. You will also learn how to check and manage the integrity of a database.

## Exam objectives in this chapter:

- Back up databases.
- Manage and configure databases.
- Maintain database integrity.
- Manage collations.

## Lessons in this chapter:

- Lesson 1: Configuring Files and Filegroups **39**
- Lesson 2: Configuring Database Options **46**
- Lesson 3: Maintaining Database Integrity **54**

## Before You Begin

---

To complete the lessons in this chapter, you must have:

- Microsoft SQL Server 2008 installed
- The *AdventureWorks* database installed within the instance



## REAL WORLD

Michael Hotek

I have worked on millions of databases across thousands of customers during the portion of my career where I have worked with SQL Server. In all that time, I have come up with many best practices while at the same time creating many arguments among the “purists.” All my recommendations and approaches to architecting and managing SQL Servers come from a pragmatic, real-world perspective that, although rooted in a deep knowledge of SQL Server, hardware, networking, and many other components, rarely matches up with the perfect world theory.

Designing the disk structures that underlie a database is one of the cases where I deviate from a lot of the theoretical processes and computations that you will find published. Although you can find entire white papers and even sections of training classes devoted to teaching you how to calculate disk transfer and random vs. sequential writes, I have never encountered an environment where I had the time or luxury to run those calculations prior to implementing a system.

It is really nice that there are formulas to calculate the disk transfer of a given disk configuration, and you can also apply statistical methods to further refine those calculations based on the random vs. sequential I/O of a system. However, all the time spent doing the calculations is worthless unless you also know the required read and write capacity of the databases you are going to place on that disk subsystem. Additionally, unless you are buying a new storage system, dedicated to a specific application, you will have a very difficult time architecting the disk storage underneath a database according to all the theories.

The challenge in achieving optimal performance is to separate the transaction logs from data files so that you can isolate disk I/O. The transaction log is the key to high-performance write operations, because the maximum transaction rate is bound by the write capacity to the transaction log file. After taking care of the transaction log, you need to add enough files and filegroups to achieve enough disk throughput to handle the read/write activity. However, the most important component of performance is to write applications with efficient code that accesses only the minimum amount of data necessary to accomplish the business task.

# Lesson 1: Configuring Files and Filegroups

---

Data within a database is stored on disk in one or more data files. Prior to being written to the data file(s), every transaction is written to a transaction log file. In this lesson, you learn how to design the data files underneath a database, group the files into filegroups to link physical storage into a database, and manage the transaction log. You also learn how to configure the *tempdb* database for optimal performance.

**After this lesson, you will be able to:**

- Create filegroups
- Add files to filegroups
- Work with *FILESTREAM* data
- Configure the transaction log

**Estimated lesson time: 20 minutes**

## Files and Filegroups

---

Although storing all your data in memory would provide extremely fast access, you would lose everything after the machine was shut down. To protect your data, it has to be persisted to disk. Underneath each database is one or more files for persisting your data.

SQL Server uses two different types of files—data and transaction log files. Data files are responsible for the long-term storage of all the data within a database. Transaction log files, discussed in more detail later in this lesson, are responsible for storing all the transactions that are executed against a database.

Instead of defining the storage of objects directly to a data file, SQL Server provides an abstraction layer for more flexibility called a *filegroup*. Filegroups are a logical structure, defined within a database, that map a database and the objects contained within a database, to the data files on disk. Filegroups can contain more than one data file.

All objects that contain data, tables, indexes, and indexed views have an *ON* clause that you can use to specify when you create an object that allows you to specify the filegroup where SQL Server stores the object. As data is written to the objects, SQL Server uses the filegroup definition to determine on which file(s) it should store the data.

At the time that a file is added to a database, you specify the initial size of the file. You can also specify a maximum size for the file, as well as whether SQL Server automatically increases the size of the file when it is full of data. If you specify automatic growth, you can specify whether the file size increases based on a percentage of the current size or whether the file size increases at a fixed amount that you define.

Unless a filegroup has only a single file, you do not know in which file a specific row of data is stored. When writing to files, SQL Server uses a proportional fill algorithm. The proportional fill algorithm is designed to ensure that all files within a filegroup reach the maximum defined capacity at the same time. For example, if you had a data file that was 10 gigabytes (GB) and a data file that was 1 GB, SQL Server writes ten rows to the 10 GB file for every one row that is written to the 1 GB file.

The proportional fill algorithm is designed to allow a resize operation to occur at a filegroup level. In other words, all files within a filegroup expand at the same time.

## File Extensions

SQL Server uses three file extensions: `.mdf`, `.ndf`, and `.ldf`. Unfortunately, many people have placed a lot of emphasis and meaning on these three extensions, where no meaning was ever intended. Just like Microsoft Office Word documents have a `.doc` or `.docx` extension, and Microsoft Office Excel files have an `.xls` or `.xlsx` extension, the extension is nothing more than a naming convention. I could just as easily create a Word document with an extension of `.bob`, or even no extension, without changing the fact that it is still a Word document or preventing the ability of Word to open and manipulate the file.

A file with an `.mdf` extension is usually the first data file that is created within a database, generally is associated with the primary filegroup, and usually is considered the primary data file which contains all the system objects necessary to a database. The `.ndf` extension is generally used for all other data files underneath a database, regardless of the filegroup to which the file is associated. The `.ldf` extension generally is used for transaction logs.

The file extensions that you see for SQL Server are nothing more than naming conventions. SQL Server does not care what the file extensions are or even if the files have extensions. If you really wanted to, you could use an `.ldf` extension for the primary data file, just as you could use an `.mdf` extension for a transaction log file. Although the use of file extensions in this way does not affect SQL Server, it generally could cause confusion among the other database administrators (DBAs) in your organization. To avoid this confusion, it is recommended that you use the `.mdf`, `.ndf`, and `.ldf` naming conventions commonly used across the SQL Server industry, but do not forget that this is just a naming convention and has absolutely no effect on SQL Server itself.

All data manipulation within SQL Server occurs in memory within a set of buffers. If you are adding new data to a database, the new data is first written to a memory buffer, then written to the transaction log, and finally persisted to a data file via a background process called *check pointing*. When you modify or delete an existing row, if the row does not already exist

in memory, SQL Server first reads the data off disk before making the modification. Similarly if you are reading data that has not yet been loaded into a memory buffer, SQL Server must read it out of the data files on disk.

If you could always ensure that the machine hosting your databases had enough memory to hold all the data within your databases, SQL Server could simply read all the data off disk into memory buffers upon startup to improve performance. However, databases are almost always much larger than the memory capacity on any machine, so SQL Server retrieves data from disk only on an as-needed basis. If SQL Server does not have enough room in memory for the data being read in, the least recently used buffer pools are emptied to make room for newly requested data.

Because accessing a disk drive is much slower than accessing memory, the data file design underneath a database can have an impact on performance.

The first layer of design is within the disk subsystem. As the number of disk drives within a volume increases, the read and write throughput for SQL Server increases. However, there is an upper limit on the disk input/output (I/O), which is based upon the capacity of the redundant array of independent disks (RAID) controller, host bus adapter (HBA), and disk bus. So you cannot fix a disk I/O bottleneck by continually adding more disk drives. Although entire 200+ page white papers have been written on random vs. sequential writes, transfer speeds, rotational speeds, calculations of raw disk read/write speeds, and other topics, the process of designing the disk subsystem is reduced to ensuring that you have enough disks along with appropriately sized controllers and disk caches to deliver the read/write throughput required by your database.

If it were simply a matter of the number of disks, there would be far fewer disk I/O bottlenecks in systems. But there is a second layer of data file design: determining how many data files you need and the location of each data file.

SQL Server creates a thread for each file underneath a database. As you increase the number of files underneath a database, SQL Server creates more threads that can be used to read and write data. However, you cannot just create a database with thousands of files to increase its number of threads. This is because each thread consumes memory, taking away space for data to be cached, and even if you could write to all the threads at the same time, you would then saturate the physical disks behind the data files. In addition, managing thousands of data files underneath a database is extremely cumbersome, and if a large percentage of the files need to expand at the same time, you could create enough activity to halt the flow of data within the database.

Due to the competing factors and the simple fact that in the real world, few DBAs have the time to spend running complex byte transfer rate calculations or even to design the disk layer based on a precise knowledge of the data throughput required, designing the data layer is an iterative approach.

Designing the data layer of a database begins with the database creation. When you create a database, it should have three files and two filegroups. You should have a file with an .mdf extension within a filegroup named *PRIMARY*, a file with an .ndf extension in a filegroup with any name that you choose, and the transaction log with an .ldf extension.

## **NOTE FILE EXTENSIONS**

**As stated in the sidebar “File Extensions,” earlier in this chapter, file extensions are nothing more than naming conventions. They do not convey any special capabilities.**

Besides being the logical definition for one or more files that defines the storage boundary for an object, filegroups have a property called *DEFAULT*. The purpose of the *DEFAULT* property is to define the filegroup where SQL Server places objects if you do not specify the *ON* clause during object creation.

When the database is created, the primary filegroup is marked as the default filegroup. After you create the database, you should mark the second filegroup as the default filegroup. By changing the default filegroup, you ensure that any objects you create are not accidentally placed on the primary filegroup and that only the system objects for the database reside on the primary filegroup. You change the default filegroup by using the following command:

```
ALTER DATABASE <database name> MODIFY FILEGROUP <filegroup name> DEFAULT
```

The main reason not to place any of your objects on the primary filegroup is to provide as much isolation in the I/O as possible. The data in the system objects does not change as frequently as data in your objects. By minimizing the write activity to the primary data file, you reduce the possibility of introducing corruption due to hardware failures. In addition, because the state of the primary filegroup also determines the state of the database, you can increase the availability of the database by minimizing the changes made to the primary filegroup.

Following the initial creation of the database, you add filegroups as needed to separate the storage of objects within the database. You also add files to filegroups to increase the disk I/O available to the objects stored on the filegroup, thereby reducing disk bottlenecks.

## **Transaction Logs**

---

When SQL Server acknowledges that a transaction has been committed, SQL Server must ensure that the change is hardened to persistent storage. Although all writes occur through memory buffers, persistence is guaranteed by requiring that all changes are written to the transaction log prior to a commit being issued. In addition, the writes to the transaction log must occur directly to disk.

Because every change made to a database must be written directly to disk, the disk storage architecture underneath your transaction log is the most important decision affecting the maximum transaction throughput that you can achieve.

SQL Server writes sequentially to the transaction log but does not read from the log except during a restart recovery. Because SQL Server randomly reads and writes to the data files underneath a database, by isolating the transaction log to a dedicated set of disks you ensure that the disk heads do not have to move all over the disk and move in a mostly linear manner.



---

**EXAM TIP**

The maximum transaction throughput for any database is bound by the amount of data per second that SQL Server can write to the transaction log.

---

## Benchmarks

**B**enchmark disclosures are the best source of information when designing the disk storage for optimal performance. Many organizations and the press place great emphasis on various benchmarks. However, a careful study reveals that, by itself, SQL Server doesn't have as large of an impact on the overall numbers as you are led to believe. The transaction processing engine within SQL Server is extremely efficient and has a fixed contribution to transaction throughput, but the real key to maximizing the transaction rate is in the disk storage. Given the same disk configuration, a 7,200 RPM drive delivers about 50 percent of the SQL Server transaction rate of a 15,000 RPM drive. Having 100 disks underneath a transaction log generally doubles the transaction rate of having only 50 disks. In addition, one of the tricks used in benchmarks is to partition a disk such that all the SQL Server data is written to the outside half or less of the disk platter, because based on physics, as the read/write head of a disk moves toward the edge of a circular object, the velocity increases, thereby spinning a larger segment of the disk platter underneath the drive head per unit of time.

## **FILESTREAM** data

---

Although the volume of data within organizations has been exploding, leading the way in this data explosion is unstructured data. To tackle the problem of storing, managing, and combining the large volumes of unstructured databases with the structured data in your databases, SQL Server 2008 introduced *FILESTREAM*.

The *FILESTREAM* feature allows you to associate files with a database. The files are stored in a folder on the operating system, but are linked directly into a database where the files can be backed up, restored, full-text-indexed, and combined with other structured data.

Although the details of *FILESTREAM* are covered in more detail in Chapter 3, "Tables," and Chapter 5, "Full Text Indexing," to store *FILESTREAM* data within a database, you need to specify where the data will be stored. You define the location for *FILESTREAM* data in a database by designating a filegroup within the database to be used for storage with the *CONTAINS FILESTREAM* property. The *FILENAME* property defined for a *FILESTREAM* filegroup specifies the path to a folder. The initial part of the folder path definition must exist; however, the last folder in the path defined cannot exist and is created automatically. After the *FILESTREAM* folder has been created, a *filestream.hdr* file is created in the folder, which is a system file used to manage the files subsequently written to the folder.

## **tempdb Database**

---

Because the *tempdb* database is much more heavily used than in previous versions, special care needs to be taken in how you design the storage underneath *tempdb*.

In addition to temporary objects, SQL Server uses *tempdb* for worktables used in grouping/sorting operations, worktables to support cursors, the version store supporting snapshot isolation level, and overflow for table variables. You can also cause index build operations to use space in *tempdb*.

Due to the potential for heavy write activity, you should move *tempdb* to a set of disks separated from your databases and any backup files. To spread out the disk I/O, you might consider adding additional files to *tempdb*.

### **NOTE MULTIPLE tempdb FILES**

**A common practice for *tempdb* is to create one file per processor. The one file per processor is with respect to what SQL Server would consider a processor and not the physical processor, which could have multiple cores as well as hyperthreading.**

### **Quick Check**

1. What are the types of files that you create for databases and what are the commonly used file extensions?
2. What is the purpose of the transaction log?

### **Quick Check Answers**

1. You can create data and log files for a database. Data files commonly have either an .mdf or .ndf extension, whereas log files have an .ldf extension.
2. The transaction log records every change that occurs within a database to persist all transactions to disk.

## **PRACTICE Creating Databases**

---

In this practice, you will create a database with multiple files that is enabled for FILESTREAM storage in the c:\test folder

1. Execute the following code to create a database:

```
CREATE DATABASE TK432 ON PRIMARY
( NAME = N'TK432_Data', FILENAME = N'c:\test\TK432.mdf' ,
  SIZE = 8MB , MAXSIZE = UNLIMITED, FILEGROWTH = 16MB ),
FILEGROUP FG1
```

```
( NAME = N'TK432_Data2', FILENAME = N'c:\test\TK432.ndf' ,
    SIZE = 8MB , MAXSIZE = UNLIMITED, FILEGROWTH = 16MB ),
FILEGROUP Documents CONTAINS FILESTREAM DEFAULT
( NAME = N'Documents', FILENAME = N'c:\test\TK432Documents' )
LOG ON
( NAME = N'TK432_Log', FILENAME = N'c:\test\TK432.ldf' ,
    SIZE = 8MB , MAXSIZE = 2048GB , FILEGROWTH = 16MB )
GO
```

2. Execute the following code to change the default filegroup:

```
ALTER DATABASE TK432
MODIFY FILEGROUP FG1
DEFAULT
GO
```

## Lesson Summary

- You can define one or more data and log files for the physical storage of a database.
- Data files are associated to a filegroup within a database.
- Filegroups provide the logical storage container for objects within a database.
- Files can be stored using the new *FILESTREAM* capabilities.

## Lesson Review

The following question is intended to reinforce key information presented in Lesson 1, “Configuring Files and Filegroups.” The question is also available on the companion CD if you prefer to review it in electronic form.

### **NOTE ANSWERS**

**Answers to this question and an explanation of why each answer choice is correct or incorrect is located in the “Answers” section at the end of the book.**

1. You have a reference database named *OrderHistory*, which should not allow any data to be modified. How can you ensure, with the least amount of effort, that users can only read data from the database?
  - A. Add all database users to the `db_datareader` role.
  - B. Create views for all the tables and grant select permission only on the views to database users.
  - C. Set the database to `READ_ONLY`.
  - D. Grant select permission on the database to all users and revoke insert, update, and delete permissions from all users on the database.

## Lesson 2: Configuring Database Options

---

Data within a database is stored on disk in one or more data files. Prior to being written to the data file(s), every transaction is written to a transaction log file. In this lesson, you learn how to design the data files underneath a database, group the files into filegroups to link physical storage into a database, and manage the transaction log.

**After this lesson, you will be able to:**

- Set the database recovery model
- Configure database options
- Manage collation sequences
- Check and maintain database consistency

**Estimated lesson time: 20 minutes**

## Database Options

---

A database has numerous options that control a variety of behaviors. These options are broken down into several categories, including the following:

- Recovery
- Auto options
- Change tracking
- Access
- Parameterization

## Recovery Options

The recovery options determine the behavior of the transaction log and how damaged pages are handled.

### Recovery Models

Every database within a SQL Server instance has a property setting called the *recovery model*. The recovery model determines the types of backups you can perform against a database. The recovery models available in SQL Server 2008 are:

- Full
- Bulk-logged
- Simple

## THE FULL RECOVERY MODEL

When a database is in the Full recovery model, all changes made, using both data manipulation language (DML) and data definition language (DDL), are logged to the transaction log. Because all changes are recorded in the transaction log, it is possible to recover a database in the Full recovery model to a given point in time so that data loss can be minimized or eliminated if you should need to recover from a disaster. Changes are retained in the transaction log indefinitely and are removed only by executing a transaction log backup.

### **BEST PRACTICES RECOVERY MODELS**

**Every production database that accepts transactions should be set to the Full recovery model. By placing the database in the Full recovery model, you can maximize the restore options that are possible.**

## THE BULK-LOGGED RECOVERY MODEL

Certain operations are designed to manipulate large amounts of data. However, the overhead of logging to the transaction log can have a detrimental impact on performance. The Bulk-logged recovery model allows certain operations to be executed with minimal logging. When a minimally logged operation is performed, SQL Server does not log every row changed but instead logs only the extents, thereby reducing the overhead and improving performance. The operations that are performed in a minimally logged manner with the database set in the Bulk-logged recovery model are:

- *BCP*
- *BULK INSERT*
- *SELECT...INTO*
- *CREATE INDEX*
- *ALTER INDEX...REBUILD*

Because the Bulk-logged recovery model does not log every change to the transaction log, you cannot recover a database to a point in time, within the interval that a minimally logged transaction executed, when the Bulk-logged recovery model was enabled.

## THE SIMPLE RECOVERY MODEL

The third recovery model is Simple. A database in the Simple recovery model logs operations to the transaction log exactly as the Full recovery model does. However, each time the database checkpoint process executes, the committed portion of the transaction log is discarded. A database in the Simple recovery model cannot be recovered to a point in time because it is not possible to issue a transaction log backup for a database in the simple recovery model.

Because the recovery model is a property of a database, you set the recovery model by using the *ALTER DATABASE* command as follows:

```
ALTER DATABASE database_name  
SET RECOVERY { FULL | BULK_LOGGED | SIMPLE }
```

The backup types available for each recovery model are shown in Table 2-1.

**TABLE 2-1** Backup Types Available for Each Recovery Model

RECOVERY MODEL	BACKUP TYPE		
	FULL	DIFFERENTIAL	TRAN LOG
Full	Yes	Yes	Yes
Bulk	Yes	Yes	Yes/no minimally logged
Simple	Yes	Yes	No



**EXAM TIP**

You need to know which types of backups are possible for each recovery model.

## Damaged Pages

It is possible to damage data pages during a write to disk if you have a power failure or failures in disk subsystem components during the write operation. If the write operation fails to complete, you can have an incomplete page in the database that cannot be read. Because the damage happens to a page on disk, the only time that you see a result of the damage is when SQL Server attempts to read the page off disk.

The default configuration of SQL Server does not check for damaged pages and could cause the database to go off-line if a damaged page is encountered. The `PAGE_VERIFY CHECKSUM` option can be enabled, which allows you to discover and log damaged pages. When pages are written to disk, a checksum for the page is calculated and stored in the page header. When SQL Server reads a page from disk, a checksum is calculated and compared to the checksum stored in the page header. If a damaged page is encountered, an 824 error is returned to the calling application and logged to the SQL Server error log and Windows Application Event log, and the ID of the damaged page is logged to the `suspect_pages` table in the `msdb` database.

In SQL Server 2005, the only way to fix a damaged page was to execute a page restore, which is discussed in Chapter 9, “Backing Up and Restoring a Database.” In addition to a page restore, if the database is participating in a database mirroring session, SQL Server 2008 automatically replaces the page with a copy of the page from the mirror. When Database Mirroring automatically fixes a corrupt page, an entry is logged and can be reviewed with the `sys.dm_db_mirroring_auto_page_repair` view.

## Auto Options

There are five options for a database that enable certain actions to occur automatically:

- `AUTO_CLOSE`
- `AUTO_SHRINK`

- `AUTO_CREATE_STATISTICS`
- `AUTO_UPDATE_STATISTICS`
- `AUTO_UPDATE_STATISTICS_ASYNC`

Each database within an instance requires a variety of resources, the most significant of which is a set of memory buffers. Each open database requires several bytes of memory and any queries against the database populate the data and query caches. If the `AUTO_CLOSE` option is enabled, when the last connection to a database is closed, SQL Server shuts down the database and releases all resources related to the database. When a new connection is made to the database, SQL Server starts up the database and begins allocating resources.

By default, `AUTO_CLOSE` is disabled. Unless you have severe memory pressure, you should not enable a database for `AUTO_CLOSE`. In addition, a database that is frequently accessed should not be set to `AUTO_CLOSE` because it would cause a severe degradation in performance. This is because you would never be able to use the data and query caches adequately.

Data files can be set to grow automatically when additional space is needed. Although most operations to increase space affect the database on a long-term basis, some space increases are needed only on a temporary basis. If the `AUTO_SHRINK` option is enabled, SQL Server periodically checks the space utilization of data and transaction log files. If the space checking algorithm finds a data file that has more than 25 percent free space, the file automatically shrinks to reclaim disk space.

Expanding a database file is a very expensive operation. Shrinking a database file is also an expensive operation. If the size of a database file increased during normal operations, it is very likely that if the file shrinks, the operation would recur and increase the database file again. The only operations that cause one-time space utilization changes to database files are administrative processes that create and rebuild indexes, archive data, or load data. Because the growth of database files is so expensive, it is recommended to leave the `AUTO_SHRINK` option disabled and manually shrink files only when necessary.

Statistics allow the Query Optimizer to build more efficient query plans. If the `AUTO_CREATE_STATISTICS` option is enabled, SQL Server automatically creates statistics that are missing during the optimization phase of query processing. Although the creation of statistics incurs some overhead, the benefit to query performance is worth the overhead cost for SQL Server to create statistics automatically when necessary.

Statistics capture the relative distribution of values in one or more columns of a table. After the database has been in production for a while, normal database changes do not appreciably change the statistics distribution in general. However, mass changes to the data or dramatic shifts in business processes can suddenly introduce significant skew into the data. If the statistics are not updated to reflect the distribution shift, the Optimizer could select an inefficient query plan.

Databases have two options that allow SQL Server to update out-of-date statistics automatically. The `AUTO_UPDATE_STATISTICS` option updates out-of-date statistics during query optimization. If you choose to enable `AUTO_UPDATE_STATISTICS`, a second

option, `AUTO_UPDATE_STATISTICS_ASYNC`, controls whether statistics are updated during query optimization or if query optimization continues while the statistics are updated asynchronously.

## Change Tracking

One of the challenges for any multiuser system is to ensure that the changes of one user do not accidentally overwrite the changes of another. To prevent the changes of multiple users from overriding each other, applications are usually built within mechanisms to determine whether a row has changed between the time it was read and the time it is written back to the database. The tracking mechanisms usually involve columns with either a datetime or timestamp column and also might include an entire versioning system.

SQL Server 2008 introduces a new feature implemented through the `CHANGE_TRACKING` database option. Change tracking is a lightweight mechanism that associates a version with each row in a table that has been enabled for change tracking. Each time the row is changed, the version number is incremented. Instead of building systems to avoid changes from multiple users overriding each other, applications need only compare the row version to determine if a change has occurred to the row between when the row was read and written.

After change tracking has been enabled for the database, you can choose which tables within a database that change tracking information should be captured for. Over time, change tracking information accumulates in the database, so you can also specify how long tracking information is retained through the `CHANGE_RETENTION` option and whether tracking information should be automatically cleaned up with the `AUTO_CLEANUP` option.

## Access

Access to a database can be controlled through several options.

The status of a database can be explicitly set to `ONLINE`, `OFFLINE`, or `EMERGENCY`. When a database is in an `ONLINE` state, you can perform all operations that would otherwise be possible. A database that is in an `OFFLINE` state is inaccessible. A database in an `EMERGENCY` state can be accessed only by a member of the `db_owner` role, and the only command allowed to be executed is `SELECT`.

You can control the ability to modify data for an online database by setting the database to either `READ_ONLY` or `READ_WRITE`. A database in `READ_ONLY` mode cannot be written to. In addition, when a database is placed in `READ_ONLY` mode, SQL Server removes any transaction log file that is specified for the database. Changing a database from `READ_ONLY` to `READ_WRITE` causes SQL Server to re-create the transaction log file.

User access to a database can be controlled through the `SINGLE_USER`, `RESTRICTED_USER`, and `MULTI_USER` options. When a database is in `SINGLE_USER` mode, only a single user is allowed to access the database. A database set to `RESTRICTED_USER` only allows access to members of the `db_owner`, `dbcreator`, and `sysadmin` roles.

If multiple users are using the database when you change the mode to `SINGLE_USER` or users that conflict with the allowed set for `RESTRICTED_USER`, the `ALTER DATABASE` command is blocked until all the non-allowed users disconnect. Instead of waiting for users to complete operations and disconnect from the database, you can specify a `ROLLBACK` action to terminate connections forcibly. The `ROLLBACK IMMEDIATE` option forcibly rolls back any open transactions, along with disconnecting any nonallowed users. You can allow users to complete transactions and exit the database by using the `ROLLBACK AFTER <number of seconds>` option, which waits for the specified number of seconds before rolling back transactions and disconnecting users.

The normal operational mode for most databases is `ONLINE`, `READ_WRITE`, and `MULTI_USER`.

## Parameterization

One of the “hot button” topics in application development is whether to parameterize calls to the database. When a database call is parameterized, the values are passed as variables. You can find just as many articles advocating for both sides. Unfortunately, applications gain a significant benefit when database calls are parameterized.

SQL Server caches the query plan for every query that is executed. Unless there is pressure on the query cache that forces a query plan from the cache, every query executed since the instance started is in the query cache. When a query is executed, SQL Server parses and compiles the query. The query is then compared to the query cache using a string-matching algorithm. If a match is found, SQL Server retrieves the plan that has already been generated and executes the query.

A query that is parameterized has a much higher probability of being matched because the query string does not change even when the values being used vary. Therefore, parameterized queries can reuse cached query plans more frequently and avoid the time required to build a query plan.

Because not all applications parameterize calls to the database, you can force SQL Server to parameterize every query for a given database by setting the `PARAMETERIZATION FORCED` database option.

The default setting for a database is not to force parameterization. The reuse of query plans provides a benefit so long as the query plan being reused is the most efficient path through the data. For tables where there is significant data skew, one value produces an efficient query plan, whereas another value causes a different query plan to be created. In addition, applications see the effect of parameterization only if the majority of database calls have an extremely short duration.

So long as the majority of your database calls have a very short duration and the query plan generated do not change depending upon the parameters passed, you could see a performance boost by forcing parameterization.

# Collation Sequences

---

SQL Server has the capability to store character data that spans every possible written language. However, not every language follows the same rules for sorting or data comparisons. SQL Server allows you to define the rules for comparison, sorting, case sensitivity, and accent sensitivity through the specification of a collation sequence.

When you install SQL Server, you specify a default collation sequence that is used for all databases, tables, and columns. You can override the default collation sequence at each level. The collation sequence for an instance can be overridden at a database level by specifying the `COLLATE` clause in either the `CREATE DATABASE` or `ALTER DATABASE` command.

## ✓ Quick Check

1. How do you restrict database access to members of the `db_owner` role and terminate all active transactions and connection at the same time?
2. What backups can be executed for a database in each of the recovery models?

### Quick Check Answers

1. You would execute the following command: `ALTER DATABASE <database name> SET RESTRICTED_USER WITH ROLLBACK IMMEDIATE.`
2. You can create full, differential, and file/filegroup backups in the Simple recovery model. The Bulk-logged recovery model allows you to execute types of backups, but you cannot restore a database to a point in time during an interval when a minimally logged transaction is executing. All types of backups can be executed in the Full recovery model.

## **PRACTICE** Changing the Database Recovery Model

---

In this practice, you change the recovery model of the *AdventureWorks* database to *FULL* to ensure that you can recover from a failure to a point in time.

1. Execute the following code:

```
ALTER DATABASE AdventureWorks
    SET RECOVERY FULL
GO
```

2. Right-click the *AdventureWorks* database, select Properties, and select the Options tab to view the recovery model and make sure that it is full.

## Lesson Summary

- You can set the recovery model for a database to Full, Bulk-logged, or Simple.
- You can back up transaction logs for a database in the Full or Bulk-logged recovery model.
- The AUTO\_SHRINK option shrinks a database file when there is more than 25 percent of free space in the file.
- You can track and log damaged pages by enabling the PAGE\_VERIFY CHECKSUM option.

## Lesson Review

The following question is intended to reinforce key information presented in Lesson 2, “Configuring Database Options.” The question is also available on the companion CD if you prefer to review it in electronic form.

### **NOTE ANSWERS**

**Answers to this question and an explanation of why each answer choice is correct or incorrect is located in the “Answers” section at the end of the book.**

1. You are the database administrator at Blue Yonder Airlines and are primarily responsible for the *Reservations* database, which runs on a server running SQL Server 2008. In addition to customers booking flights through the company’s Web site, flights can be booked with several partners. Once an hour, the *Reservations* database receives multiple files from partners, which are then loaded into the database using the Bulk Copy Program (BCP) utility. You need to ensure that you can recover the database to any point in time while also maximizing the performance of import routines. How would you configure the database to meet business requirements?
  - A. Enable AUTO\_SHRINK
  - B. Set PARAMETERIZATION FORCED on the database
  - C. Configure the database in the Bulk-logged recovery model
  - D. Configure the database in the Full recovery model

## Lesson 3: Maintaining Database Integrity

---

In a perfect world, everything that you save to disk storage would always write correctly, read correctly, and never have any problems. Unfortunately, your SQL Server databases live in an imperfect world where things do go wrong. Although this occurs very rarely, data within your database can become corrupted if there is a failure in the disk storage system as SQL Server is writing to a page. Data pages are 8 kilobytes (KB) in size, but SQL Server divides a page into 16 blocks of 512 bytes apiece when performing write operations. If SQL Server begins writing blocks on a page and the disk system fails in the middle of the write process, only a portion of the page is written successfully, producing a problem called a *torn page*. In this lesson, you learn how to detect and correct corruption errors in your database.

**After this lesson, you will be able to:**

- Check a database for integrity
- Use DMVs to diagnose corruption issues

**Estimated lesson time: 20 minutes**

### Database Integrity Checks

---

As you learned in Lesson 2, databases have an option called *PAGE\_VERIFY*. The page verification can be set to either *TORN\_PAGE\_DETECTION* or *CHECKSUM*. The *PAGE\_VERIFY TORN\_PAGE\_DETECTION* option exists for backwards compatibility and should not be used. When the *PAGE\_VERIFY CHECKSUM* option is enabled, SQL Server calculates a checksum for the page prior to the write. Each time a page is read off disk, a checksum is recalculated and compared to the checksum written to the page. If the checksums do not match, the page has been corrupted.

When SQL Server encounters a corrupt page, an error is thrown, the command attempting to access the corrupt page is aborted, and an entry is written into the *suspect\_pages* table in the *msdb* database.

#### **BEST PRACTICES PAGE VERIFICATION**

**You should enable the *PAGE\_VERIFY CHECKSUM* option on every production database.**

Although page verification can detect and log corrupted pages, the page must be read off disk to trigger the verification check. Data is normally read off disk when users and applications access data, but instead of having a user receive an error message, it is much better for you to proactively find corruption and fix the problem by using a backup before the user has a process aborted.

You can force SQL Server to read every page from disk and check the integrity by executing the *DBCC CHECKDB* command. The generic syntax of *DBCC CHECKDB* is:

```
DBCC CHECKDB [( 'database_name' | database_id | 0
  [ , NOINDEX | { REPAIR_ALLOW_DATA_LOSS | REPAIR_FAST
  | REPAIR_REBUILD } ] )]
  [ WITH { [ ALL_ERRORMSGS ] [ , [ NO_INFOMSGS ] ] [ , [ TABLOCK ] ]
    [ , [ ESTIMATEONLY ] ] [ , [ PHYSICAL_ONLY ] ] | [ , [ DATA_PURITY ] ] } ] ]
```

When *DBCC CHECKDB* is executed, SQL Server performs all the following actions:

- Checks page allocation within the database
- Checks the structural integrity of all tables and indexed views
- Calculates a checksum for every data and index page to compare against the stored checksum
- Validates the contents of every indexed view
- Checks the database catalog
- Validates Service Broker data within the database

To accomplish these checks, *DBCC CHECKDB* executes the following commands:

- *DBCC CHECKALLOC*, to check the page allocation of the database
- *DBCC CHECKCATALOG*, to check the database catalog
- *DBCC CHECKTABLE*, for each table and view in the database to check the structural integrity

Any errors encountered are output so that you can fix the problems. If an integrity error is found in an index, you should drop and re-create the index. If an integrity error is found in a table, you need to use your most recent backups to repair the damaged pages.

#### **NOTE DATABASE MIRRORING**

If the database is participating in Database Mirroring, SQL Server attempts to retrieve a copy of the page from the mirror. If the page can be retrieved from the mirror and has the correct page contents, the page is replaced automatically on the principal without requiring any intervention. When SQL Server replaces a corrupt page from the mirror, an entry is written into the *sys.dm\_db\_mirroring\_auto\_page\_repair* view.

#### **Quick Check**

1. Which option should be enabled for all production databases?
2. What checks does *DBCC CHECKDB* perform?

## Quick Check Answers

1. You should set the *PAGE\_VERIFY CHECKSUM* option for all production databases.
2. *DBCC CHECKDB* checks the logical and physical integrity of every table, index, and indexed view within the database, along with the contents of every indexed view, page allocations, Service Broker data, and database catalog.

## PRACTICE Checking Database Integrity

In this practice, you check the integrity of the *AdventureWorks* database.

1. Execute the following code:

```
DBCC CHECKDB ('AdventureWorks') WITH NO_INFOMSGS, ALL_ERRORMSGS  
GO
```

2. Review the results.

## Lesson Summary

- The *PAGE\_VERIFY CHECKSUM* option should be enabled for every production database to detect any structural integrity errors.
- When a corrupt page is encountered, the page is logged to the *suspect\_pages* table in the *msdb* database. If a database is participating in a Database Mirroring session, SQL Server automatically retrieves a copy of the page from the mirror, replaces the page on the principal, and logs an entry in the *sys.dm\_db\_mirroring\_auto\_page\_repair* view.
- *DBCC CHECKDB* is used to check the logical and physical consistency of a database.

## Lesson Review

The following question is intended to reinforce key information presented in Lesson 3, “Maintaining Database Integrity.” The question is also available on the companion CD if you prefer to review it in electronic form.

### NOTE ANSWERS

Answers to this question and an explanation of why each answer choice is correct or incorrect is located in the “Answers” section at the end of the book.

1. Which commands are executed when you run the *DBCC CHECKDB* command? (Check all that apply.)
  - A. *DBCC CHECKTABLE*
  - B. *DBCC CHECKIDENT*
  - C. *DBCC CHECKCATALOG*
  - D. *DBCC FREEPROCCACHE*

# Chapter Review

---

To practice and reinforce the skills you learned in this chapter further, you can perform the following tasks:

- Review the chapter summary.
- Review the list of key terms introduced in this chapter.
- Complete the case scenario. This scenario sets up a real-world situation involving the topics in this chapter and asks you to create a solution.
- Complete the suggested practices.
- Take a practice test.

## Chapter Summary

- Databases can be configured with the Full, Bulk-logged, or Simple recovery model.
- The recovery model of the database determines the backups that can be created, as well as limitations on the recovery options that can be performed.
- You can set a collation sequence for a database that overrides the collation sequence defined for the instance.

## Key Terms

Do you know what these key terms mean? You can check your answers by looking up the terms in the glossary at the end of the book.

- Corrupt page
- Filegroup
- Recovery model

## Case Scenario

In the following case scenario, you apply what you've learned in this chapter. You can find answers to these questions in the "Answers" section at the end of this book.

### Case Scenario: Configuring Databases for Coho Vineyard

#### BACKGROUND

##### Company Overview

Coho Vineyard was founded in 1947 as a local, family-run winery. Due to the award-winning wines it has produced over the last several decades, Coho Vineyards has experienced significant growth. To continue expanding, several existing wineries were acquired over the years. Today, the company owns 16 wineries; 9 wineries are in Washington, Oregon, and California, and the remaining 7 wineries are located in Wisconsin and Michigan. The wineries

employ 532 people, 162 of whom work in the central office that houses servers critical to the business. The company has 122 salespeople who travel around the world and need access to up-to-date inventory availability.

### Planned Changes

Until now, each of the 16 wineries owned by Coho Vineyard has run a separate Web site locally on the premises. Coho Vineyard wants to consolidate the Web presence of these wineries so that Web visitors can purchase products from all 16 wineries from a single online store. All data associated with this Web site be stored in databases in the central office.

When the data is consolidated at the central office, merge replication will be used to deliver data to the salespeople as well as to allow them to enter orders. To meet the needs of the salespeople until the consolidation project is completed, inventory data at each winery is sent to the central office at the end of each day. Merge replication has been implemented to allow salespeople to maintain local copies of customer, inventory, and order data.

### EXISTING DATA ENVIRONMENT

#### Databases

Each winery presently maintains its own database to store all business information. At the end of each month, this information is brought to the central office and transferred into the databases shown in Table 2-2.

**TABLE 2-2** Coho Vineyard Databases

DATABASE	SIZE
<i>Customer</i>	180 megabytes (MB)
<i>Accounting</i>	500 MB
<i>HR</i>	100 MB
<i>Inventory</i>	250 MB
<i>Promotions</i>	80 MB

After the database consolidation project is complete, a new database named *Order* will serve as a data store to the new Web store. As part of their daily work, employees also will connect periodically to the *Order* database using a new in-house Web application.

The *HR* database contains sensitive data and is protected using Transparent Data Encryption (TDE). In addition, data in the Salary table is encrypted using a certificate.

#### Database Servers

A single server named DB1 contains all the databases at the central office. DB1 is running SQL Server 2008 Enterprise on Windows Server 2003 Enterprise.

## Business Requirements

You need to design an archiving solution for the *Customer* and *Order* databases. Your archival strategy should allow the *Customer* data to be saved for six years.

To prepare the *Order* database for archiving procedures, you create a partitioned table named *Order.Sales*. *Order.Sales* includes two partitions. Partition 1 includes sales activity for the current month. Partition 2 is used to store sales activity for the previous month. Orders placed before the previous month should be moved to another partitioned table named *Order.Archive*. Partition 1 of *Order.Archive* includes all archived data. Partition 2 remains empty.

A process needs to be created to load the inventory data from each of the 16 wineries by 4 A.M. daily.

Four large customers submit orders using Coho Vineyards Extensible Markup Language (XML) schema for Electronic Data Interchange (EDI) transactions. The EDI files arrive by 5 P.M. and need to be parsed and loaded into the *Customer*, *Accounting*, and *Inventory* databases, which each contain tables relevant to placing an order. The EDI import routine is currently a single-threaded C++ application that takes between three and six hours to process the files. You need to finish the EDI process by 5:30 P.M. to meet your Service Level Agreement (SLA) with the customers. After the consolidation project has finished, the EDI routine loads all data into the new *Order* database.

You need to back up all databases at all locations. You can lose a maximum of five minutes of data under a worst-case scenario. The *Customer*, *Account*, *Inventory*, *Promotions*, and *Order* databases can be off-line for a maximum of 20 minutes in the event of a disaster. Data older than six months in the *Customer* and *Order* databases can be off-line for up to 12 hours in the event of a disaster.

Answer the following questions.

1. How should you configure the databases for maximum performance?
2. How should the databases be configured to meet recovery obligations?

## Suggested Practices

---

To help you master the exam objectives presented in this chapter, complete the following tasks.

### Configuring Databases

- **Practice 1** Create a database which can store *FILESTREAM* data.
- **Practice 2** Change the recovery model and observe the effects on backup and restore options.

- **Practice 3** Change the database state to READ\_ONLY and observe the effect on the transaction log file.
- **Practice 4** Create multiple connections to a database, change the access to RESTRICTED\_USER, and specify the ROLLBACK IMMEDIATE option. Observe the effects.

## Take a Practice Test

---

The practice tests on this book's companion CD offer many options. For example, you can test yourself on just one exam objective, or you can test yourself on all the 70-432 certification exam content. You can set up the test so that it closely simulates the experience of taking a certification exam, or you can set it up in study mode so that you can look at the correct answers and explanations after you answer each question.

### ***MORE INFO* PRACTICE TESTS**

**For details about all the practice test options available, see the section entitled "How to Use the Practice Tests," in the Introduction to this book.**

# Designing Policy Based Management

Prior to Microsoft SQL Server 2008, you performed configuration management of an environment by using a conglomeration of documents, scripts, and manual checking. The configuration options, naming conventions, and allowed feature set were outlined in one or more documents. To enforce your standards, you would have had to connect to each instance and execute scripts that needed to be maintained and updated with new versions and service packs. In this chapter, you learn about the new Policy Based Management framework that allows you to check and enforce policy compliance across your entire SQL Server infrastructure.

## Exam objectives in this chapter:

- Implement the declarative management framework (DMF).
- Configure surface area.

## Lesson in this chapter:

- Lesson 1: Designing Policies 179

## Before You Begin

---

To complete the lessons in this chapter, you must have:

- SQL Server 2008 installed
- The *AdventureWorks* database installed within the instance



### REAL WORLD

Michael Hotek

Managing a single server running SQL Server or even a small group of them, one at a time, has always been reasonably straightforward. However, when you needed to uniformly manage an entire SQL Server environment or a large group of instances, you had to either write a large amount of custom code or purchase additional products.

One customer I work with has an environment with more than 5,000 SQL Server instances. Prior to the release of SQL Server 2008, two DBAs were required to manage the almost 50,000 lines of code that checked instances for compliance to corporate policies. They devoted more than 70 hours each week to maintaining the code and checking systems.

After deploying SQL Server 2008, they started to convert all their code to policies. After the conversion was completed, they estimate that less than 1,000 lines of custom logic remained. By using the central management features to check and enforce policies across the environment, they should be able to save over 3,000 hours of management and maintenance time per year.

# Lesson 1: Designing Policies

---

SQL Server 2008 has a new feature called Policy Based Management, also known as the declarative management framework (DMF), to tackle the problem of standardizing your SQL Server instances. Although Policy Based Management can be used just to alert an administrator when an object is out of compliance, depending upon the type of policy, you can also enforce compliance by preventing changes that would violate a policy.

Policy Based Management introduces the following new objects that are used to design and check for compliance:

- Facets
- Conditions
- Policies
- Policy targets
- Policy categories

## After this lesson, you will be able to:

- Create conditions
- Define policies
- Specify targets for policy checking
- Configure policy categories
- Check for policy compliance
- Import and export policies

**Estimated lesson time: 30 minutes**

## Facets

---

*Facets* are the core object upon which your standards are built. Facets define the type of object or option to be checked, such as database, Surface Area, and login. SQL Server ships with 74 facets, implemented as .NET assemblies, each with a unique set of properties.

All the objects for Policy Based Management are stored within the *msdb* database. You can get a list of the facets available by querying the `dbo.syspolicy_management_facets` table. Unfortunately, unless you want to write code to interact with Server Management Objects (SMOs), the only way to get a list of facet properties is to open each facet in SQL Server Management Studio (SSMS), one at a time, and view the list of properties.

## Conditions

---

When you define a WHERE clause for a data manipulation language (DML) statement, you set a condition for the DML statement that defines the set of rows that meet your specific inclusion criteria. Within the Policy Based Management framework, *conditions* are the equivalent of a *WHERE* clause that defines the criteria needing to be checked.

You define the conditions that you want to check or enforce for a policy by defining criteria for the properties of a facet. Just like a WHERE clause, a condition can be defined by one or more facet properties, and a single facet property can be checked for multiple criteria. The comparison operators that can be used are restricted by the data type of the property. For example, a property of type *string* can be checked with =, <>, *LIKE*, *NOT LIKE*, *IN*, or *NOT IN*, whereas a boolean type can only be checked for = and <>.

If a condition that you want to check for a facet does not have a specific property that can be used, you can use the advanced editor to define complex conditions that compare multiple properties and incorporate functions. For example, you can check that every table has a primary key and that a table with a single index must be clustered. Unfortunately, if you define a condition using the advanced editor, a policy that incorporates the condition must be executed manually and cannot be scheduled.

Conditions are checked in a single step. You cannot have a condition pull a list of objects, iterate across the list of objects, and then apply subsequent checks. To work within the Policy-Based Management framework, conditions need to return a True or False value. Therefore, when building complex conditions with the advanced editor, you cannot return a list of objects that do not meet your criteria. You have to define the condition such that if any object does not meet your criteria, a value of False is returned.

Although you can check many properties of a facet within a single condition, a single condition can't be defined for multiple facets. For example, you can check all 10 of the properties for the Surface Area Configuration facet in a single condition, but you have to define a second condition to check a property of the Surface Area Configuration for Analysis Services.

## Policy Targets

---

Conditions are the foundation for policies. However, you don't always want to check policies across every object available, such as every database in an instance or every index within every database. Conditions can also be used to specify the objects to compare the condition against, called *policy targeting* or target sets.

You can target a policy at the server level, such as instances that are SQL Server 2005 or SQL Server 2008. You can also target a policy at the database level, such as all user databases or all system databases.

# Policies

---

Policies are created for a single condition and set to either enforce or check compliance. The execution mode can be set as follows:

- **On demand** Evaluates the policy when directly executed by a user
- **On change, prevent** Creates data definition language (DDL) triggers to prevent a change that violates the policy
- **On change, log only** Checks the policy automatically when a change is made using the event notification infrastructure
- **On schedule** Creates a SQL Server Agent job to check the policy on a defined schedule

If a policy contains a condition that was defined using the advanced editor, the only available execution mode is On Demand.

To use the *On change, prevent* and *On change, log only* execution modes, the policy must target instances running SQL Server 2005 and above. The *On change, log only* execution mode uses the event notification infrastructure that is available only for SQL Server 2005 and later. The *On change, prevent* execution mode depends on DDL triggers to prevent a change that is not in compliance with the policy and are available only for SQL Server 2005 and later. In addition, you can set a policy to *On change, prevent* only if it is possible for a DDL trigger to prevent the change. For example, you could prevent the creation of an object that violated your naming conventions, but you could not enforce a policy that all databases have to be in the Full recovery model because the *ALTER DATABASE* command executes outside the context of a transaction.

## Policy Categories

---

*Policy categories* can be used to group one or more policies into a single compliance unit. If not specified, all policies belong to the *DEFAULT* category. To check or enforce policies, you create a subscription to one or more policies.

Subscription occurs at two levels—instance and database. A member of the sysadmin role can subscribe an instance to a policy category. Once subscribed, the owner of each database within the instance can subscribe their database to a policy category.

Each policy category has a *Mandate* property that applies to databases. When a policy category is set to *Mandate* and a sysadmin subscribes the instance to a policy category, all databases that meet the target set are controlled by the policies within the policy category. A policy subscription to a policy category set to *Mandate* cannot be overridden by a database owner.

## Policy Compliance

---

Because you cannot set all policies to enforce compliance, you need to check policies manually that cannot be enforced on a regular basis. You view policies that apply to an instance by right-clicking the name of the instance within Object Explorer and selecting Policies, View.

You can check policies that apply to an instance by right-clicking the name of the instance within Object Explorer and selecting Policies, Evaluate.

You can check all policies within an instance, as shown in Figure 8-1, by right-clicking the Policies node and selecting Evaluate.

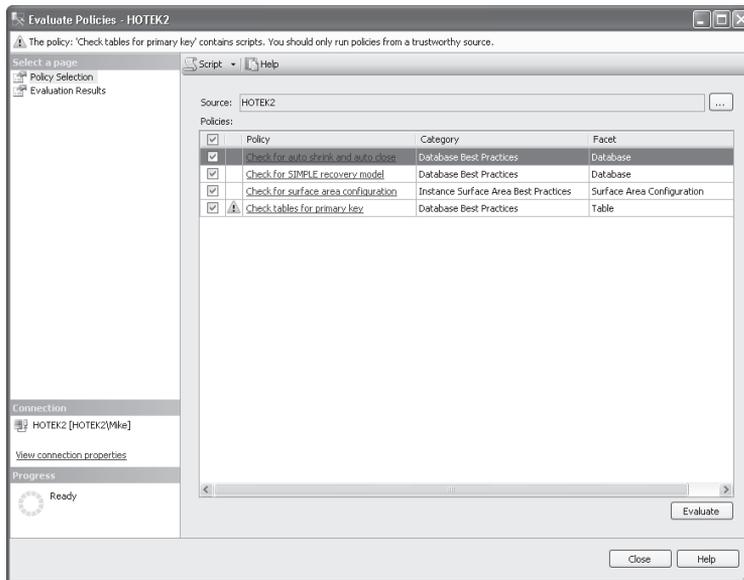


FIGURE 8-1 Evaluate policies

By clicking Evaluate, you execute the policies and review the results, as shown in Figure 8-2.

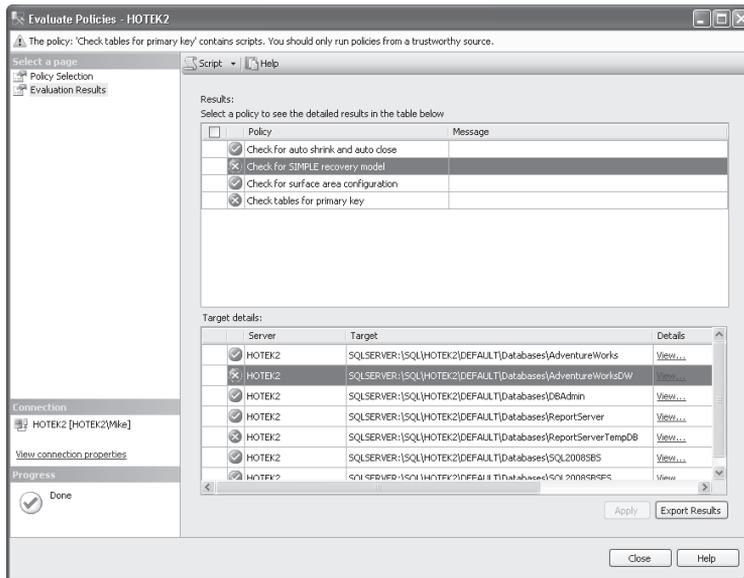


FIGURE 8-2 Policy check results



### EXAM TIP

Defining a condition to be used as a policy target is a critical component to well-defined policies. A policy fails during a check if the object does not conform to the criteria and if the property does not exist. For example, attempting to check that the Web Assistant is disabled against a SQL Server 2008 instance fails because the feature does not exist.

## Central Management Server

Policy Based Management would be limited to SQL Server 2008 and be very tedious if you had to do any of the following:

- Duplicate policies on every instance
- Create subscriptions to each instance in your environment individually
- Check compliance for each instance individually

Within the Registered Servers pane in SSMS, you can configure a Central Management Server. Underneath the Central Management Server, you can create multiple levels of folders, and register instances into the appropriate folder. After you have the Central Management Server structure set up in SSMS, you can evaluate policies against a specific instance, folder, or all instances underneath the Central Management Server. Figure 8-3 shows an example of a Central Management Server.

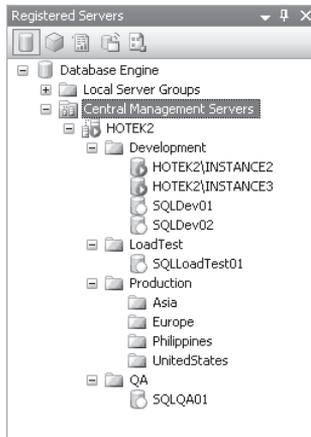


FIGURE 8-3 Central Management Server

## Import and Export Policies

Policies and conditions can be exported to files as well as imported from files. SQL Server ships with 53 policies that are located in the Microsoft SQL Server\100\Tools\Policies folder. There are 50 policies for the database engine, 2 policies for Reporting Services, and 1 policy for Analysis Services. The CodePlex site (<http://www.codeplex.com>) has additional policies that you can download and import.

You can import policies within the Registered Servers pane or the Object Explorer. Within Object Explorer, you can right-click the Policies node underneath Policy Management and select Import Policy. Within Registered Servers, you can right-click the Central Management Server or any folder or instance underneath the Central Management Server and select Import Policies. If you import policies from the Central Management Server, the policies are imported to every instance defined underneath the Central Management Server, but not to the Central Management Server itself. Likewise, right-clicking a folder imports the policies to all instances within the folder hierarchy. To import policies to the Central Management Server, you must connect to the instance within Object Explorer and import from the Policies node.

### ✓ Quick Check

1. What are the five objects that are used within Policy Based Management?
2. What are the allowed execution modes for a policy?
3. Which object has a property that allows you to mandate checking for all databases on an instance?
4. How many facets can be checked within a single condition?
5. How many conditions can be checked within a single policy?

### Quick Check Answers

1. The objects that are used with Policy Based Management are facets, conditions, policies, policy targets, and policy categories.
2. The policy execution modes are *On demand*, *On schedule*, *On change*, *Log only*, and *On change, prevent*.
3. Policy categories allow you to mandate checking of all databases within an instance.
4. A condition can be defined on only one facet.
5. A policy can check only a single condition.

## **PRACTICE** Defining Policies and Checking for Compliance

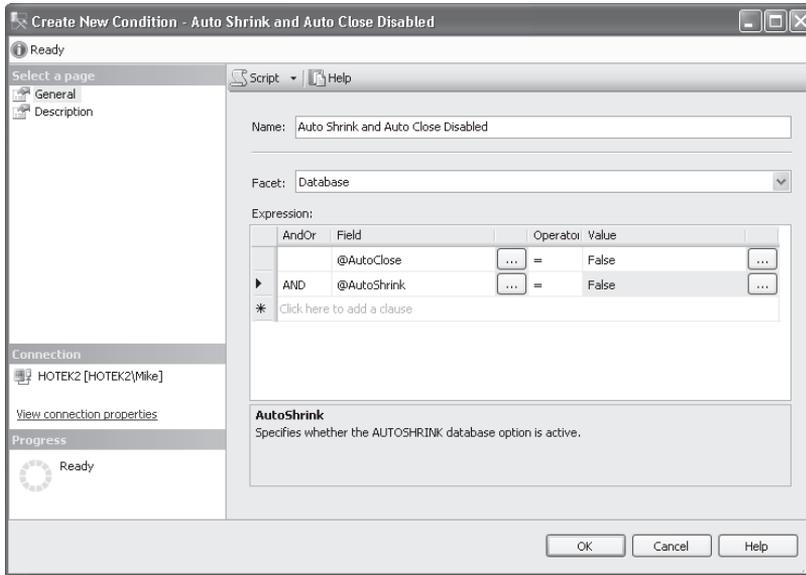
In these practices, you define and check several policies for your environment.

### **PRACTICE 1** Create a Condition

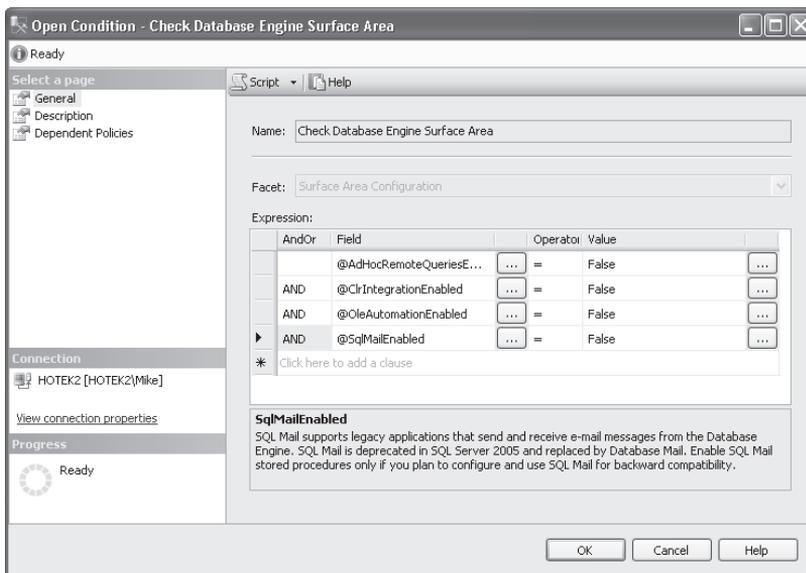
In this practice, you create a condition for the following:

- Check that a database does not have the *auto shrink* or *auto close* properties set.
- Check that CLR, OLE Automation, Ad Hoc Remote Queries, and SQL Mail are all disabled.
- Check that a database is not in the Simple recovery model.
- Check that all tables have a primary key.

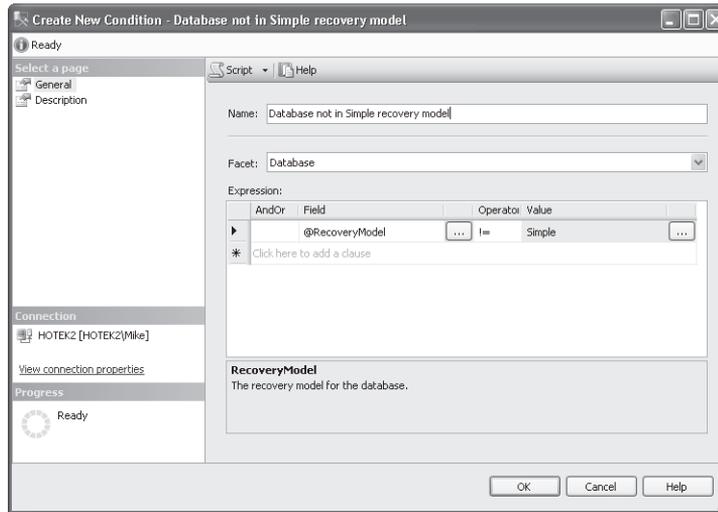
1. In Object Explorer, expand the Policy Management node within the Management node.
2. Right-click the Conditions node and select New Condition.
3. Configure the condition as shown here. Click OK when you are done.



4. Right-click the Conditions node again, select New Condition, and configure the condition as shown here. Click OK.



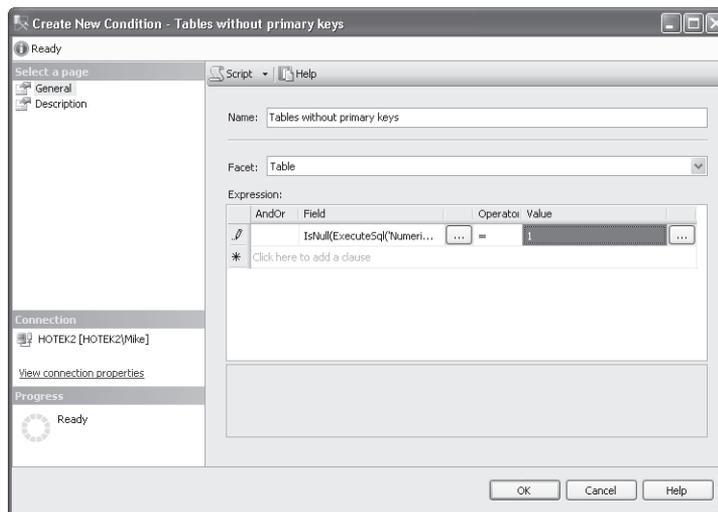
5. Right-click the Conditions node, select New Condition, and configure this third condition as shown here. Click OK when you are finished.



- Right-click the Conditions node and select New Condition. Select the Table facet, click the ellipsis button next to the Field column to display the Advanced Edit dialog box, enter the following code in the Cell Value text box, and click OK:

```
IsNull(ExecuteSql('Numeric', 'SELECT 1 FROM sys.tables a INNER JOIN sys.indexes b ON a.object_id = b.object_id WHERE b.is_primary_key = 1 AND a.name = @@ObjectName AND a.schema_id = SCHEMA_ID(@@SchemaName)'), 0)
```

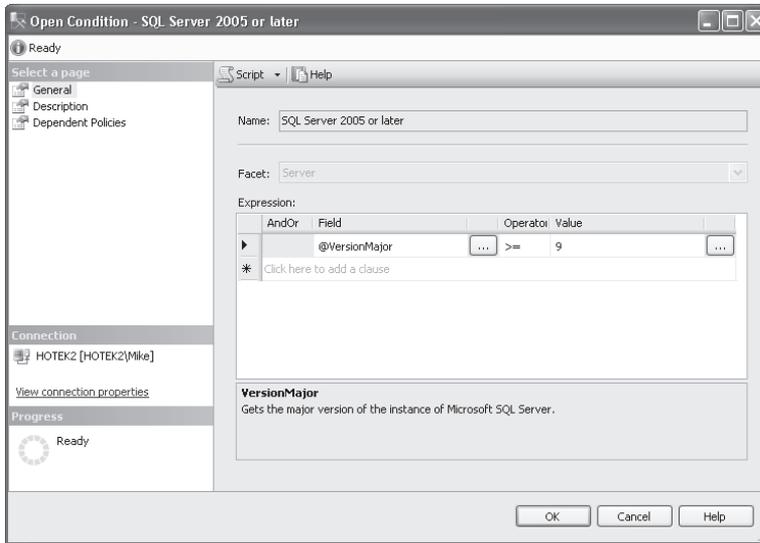
- Configure the Name, Operator, and Value as shown here, and then click OK.



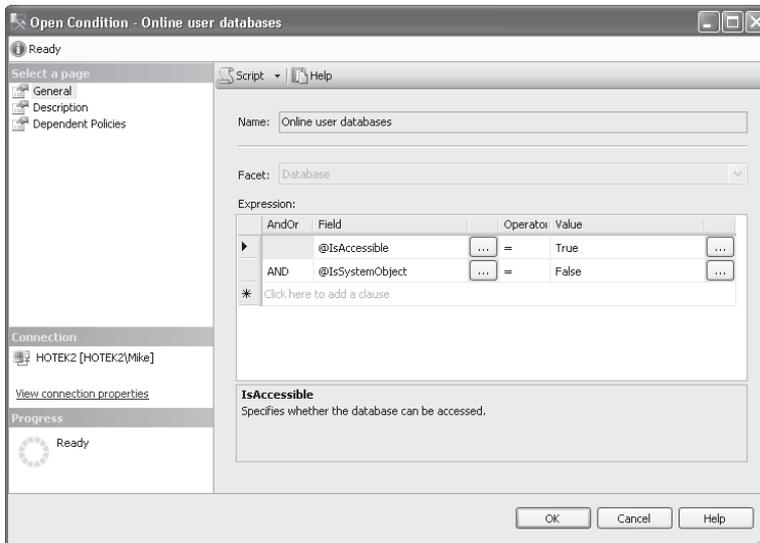
### PRACTICE 2 Create a Condition for a Target Set

In this practice, you create a condition to target all SQL Server 2005 and later instances, along with a condition to target all user databases that are online.

1. Right-click the Conditions node, select New Condition, and configure the condition as shown here. Click OK.



2. Right-click the Conditions node, select New Condition, and configure the condition as shown here. Click OK when you are done.

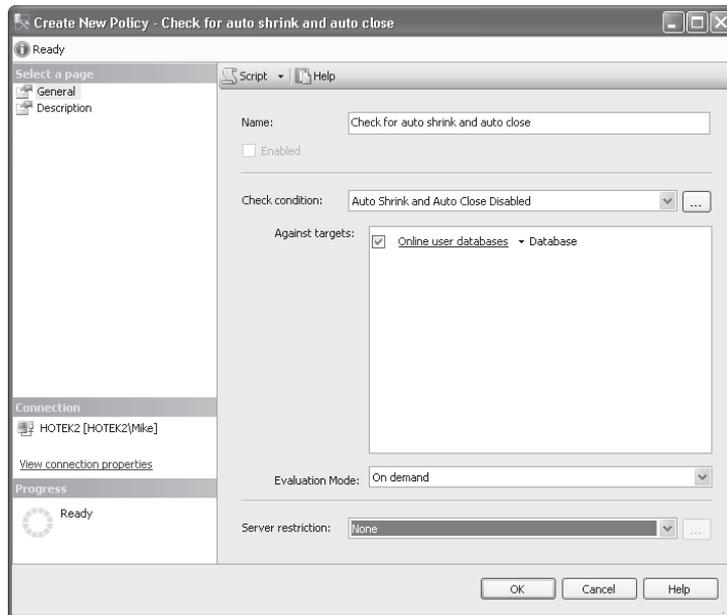


### PRACTICE 3 Create a Policy

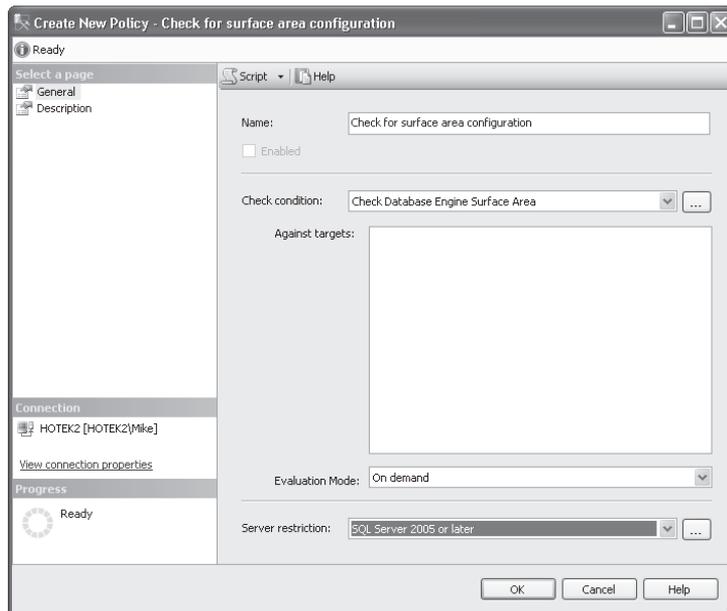
In this practice, you create policies that use the conditions you just created to do the following:

- Check that a database does not have the *auto shrink* or *auto close* properties set.
- Check that CLR, OLE Automation, Ad Hoc Remote Queries, and SQL Mail are all disabled.

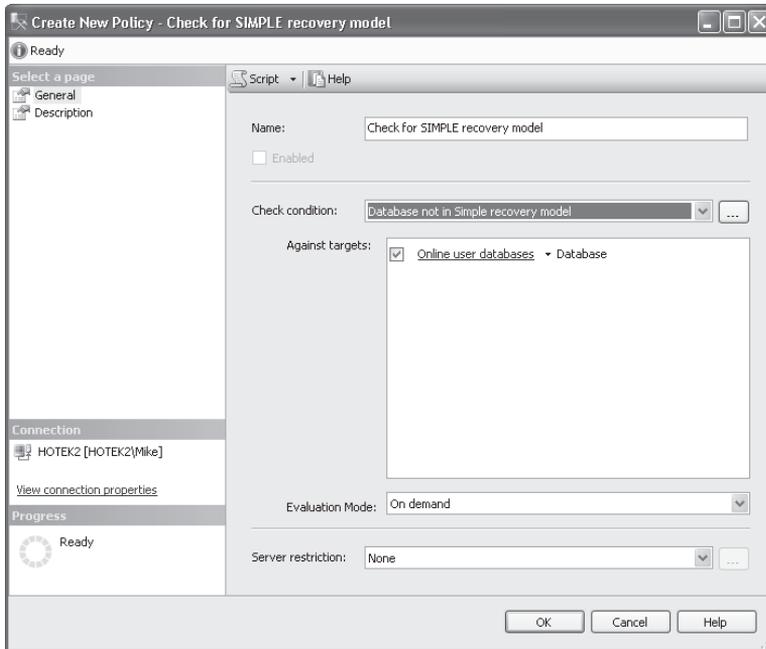
- Check that a database is not in the Simple recovery model.
  - Check that all tables have a primary key.
1. Right-click the Policies node, select New Policy, and configure the policy as shown here. Click OK.



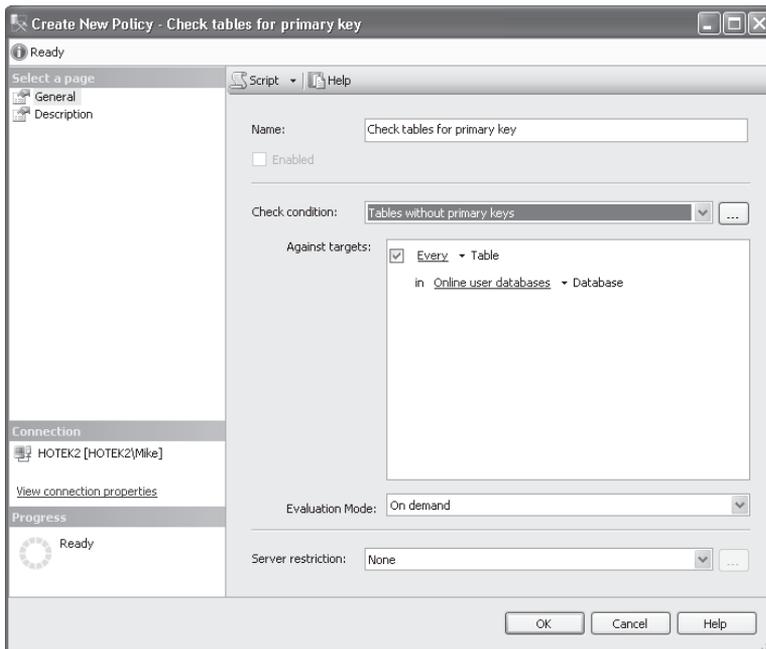
2. Right-click the Policies node, select New Policy, and configure this second policy as shown here. Click OK.



3. Right-click the Policies node, select New Policy, and configure the policy as shown here. Click OK.



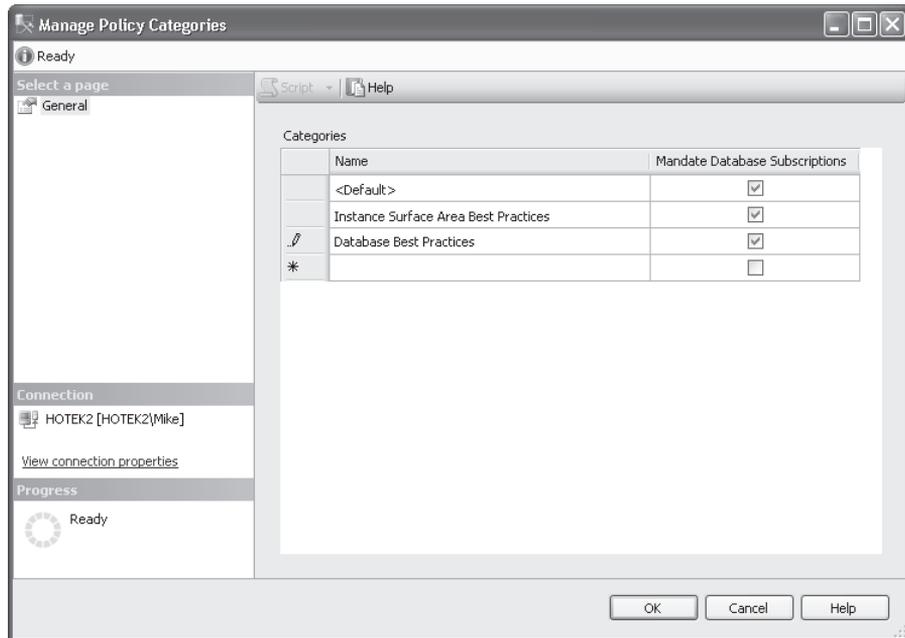
4. Right-click the Policies node, select New Policy, and configure the last policy as shown here. Click OK.



## PRACTICE 4 Create a Policy Category

In this practice, you create two policy categories for the policies that you created.

1. Right-click Policy Management, select Manage Categories, and create the categories as shown here. Click OK.

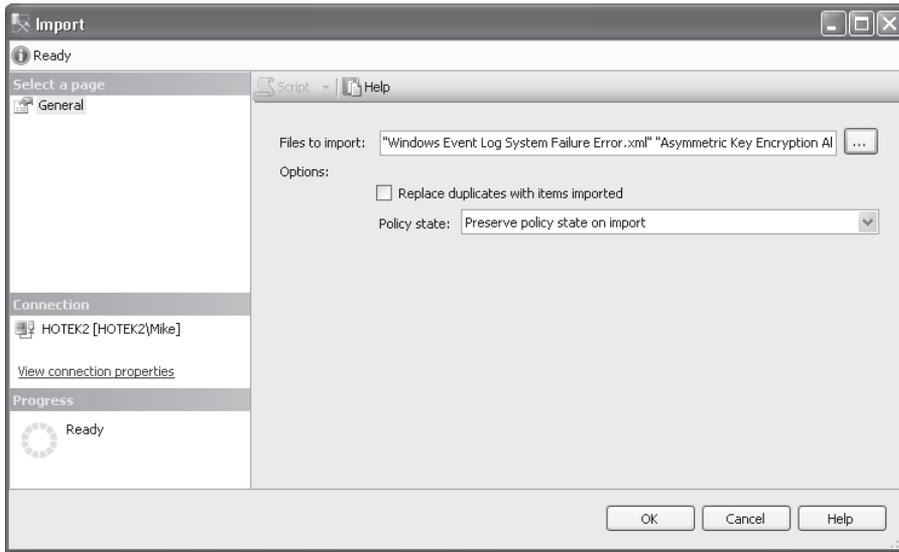


2. In SSMS, in the console tree, expand the Policies folder. Right-click the Check For Auto Shrink And Auto Close Policy, select Properties, click the Description tab, and change the category to Database Best Practices. Click OK.
3. Right-click the Check For Simple Recovery Model Policy, select Properties, select the Description tab, and change the category to Database Best Practices. Click OK.
4. Right-click the Check For Surface Area Configuration Policy, select Properties, click the Description tab, and change the category to Instance Surface Area Best Practices. Click OK.
5. Right-click the Check Tables For Primary Key Policy, select Properties, select the Description tab, and change the category to Database Best Practices. Click OK.

## PRACTICE 5 Import Policies

In this practice, you import the policies that ship with SQL Server.

1. Right-click the Policies node underneath Policy Management and select Import Policy.
2. Click the ellipsis button next to the Files To Import text box, navigate to the Microsoft SQL Server\100\Tools\Policies\DatabaseEngine\1033 folder, select all the files in the folder, as shown here, and click Open.



3. Select the Replace Duplicates With Items Imported check box, select Preserve Policy State On Import, and click OK.
4. Take the time to browse the policies and conditions that were created during the import.

## Lesson Summary

- You can build policies to enforce conditions across any version of SQL Server.
- Policies can enforce a single condition and each condition can be based on a single facet.
- Policy categories allow you to group policies together for compliance checking.
- A policy category can be set with the *Mandate* property, which requires the policy to be checked against all databases within an instance.

## Lesson Review

The following question is intended to reinforce key information presented in this lesson. The question is also available on the companion CD if you prefer to review it in electronic form.

### **NOTE ANSWERS**

**Answers to this question and an explanation of why each answer choice is correct or incorrect is located in the “Answers” section at the end of the book.**

1. You have defined several policies that you want applied to all databases within an instance. How do you ensure that a database owner is not allowed to avoid the policy check with the least amount of administrative effort?
  - A. Create a condition that checks all databases.
  - B. Add the policy to a user-defined policy category and set the *Mandate* property.
  - C. Add the policy to the default policy category.
  - D. Check the policies manually against the instance.

## Chapter Review

---

To practice and reinforce the skills you learned in this chapter further, you can perform the following tasks:

- Review the chapter summary.
- Review the list of key terms introduced in this chapter.
- Complete the case scenario. The scenario sets up a real-world situation involving the topics in this chapter and asks you to create a solution.
- Complete the suggested practices.
- Take a practice test.

## Chapter Summary

- Facets are the .NET assemblies that define the set of properties for an object upon which conditions are built.
- A condition can be defined for a single facet and a policy can be checked for a single instance.
- Policies can be checked manually or automatically. Automatic policy checking can be performed on a scheduled basis or by using the event notification infrastructure.
- A database owner can subscribe a database to one or more policies; however, a policy that belongs to a policy category set with the *Mandate* property requires checking against all databases.

## Key Terms

Do you know what these key terms mean? You can check your answers by looking up the terms in the glossary at the end of the book.

- Condition
- Facet
- Policy category
- Policy target

## Case Scenario

In the following case scenario, you apply what you've learned in this chapter. You can find answers to these questions in the "Answers" section at the end of this book.

# Case Scenario: Designing a Management Strategy for Coho Vineyard

## BACKGROUND

### Company Overview

Coho Vineyard was founded in 1947 as a local, family-run winery. Due to the award-winning wines it has produced over the last several decades, Coho Vineyards has experienced significant growth. To continue expanding, several existing wineries were acquired over the years. Today, the company owns 16 wineries; 9 wineries are in Washington, Oregon, and California, and the remaining 7 wineries are located in Wisconsin and Michigan. The wineries employ 532 people, 162 of whom work in the central office that houses servers critical to the business. The company has 122 salespeople who travel around the world and need access to up-to-date inventory availability.

### Planned Changes

Until now, each of the 16 wineries owned by Coho Vineyard has run a separate Web site locally on the premises. Coho Vineyard wants to consolidate the Web presence of these wineries so that Web visitors can purchase products from all 16 wineries from a single online store. All data associated with this Web site will be stored in databases in the central office.

When the data is consolidated at the central office, merge replication will be used to deliver data to the salespeople as well as to allow them to enter orders. To meet the needs of the salespeople until the consolidation project is completed, inventory data at each winery is sent to the central office at the end of each day.

Management wants to ensure that you cannot execute stored procedures written in C#.NET or use the *OPENROWSET* or *OPENDATASOURCE* command.

## EXISTING DATA ENVIRONMENT

### Databases

Each winery presently maintains its own database to store all business information. At the end of each month, this information is brought to the central office and transferred into the databases shown in Table 8-1.

**TABLE 8-1** Coho Vineyard Databases

DATABASE	SIZE
<i>Customer</i>	180 megabytes (MB)
<i>Accounting</i>	500 MB
<i>HR</i>	100 MB
<i>Inventory</i>	250 MB
<i>Promotions</i>	80 MB

After the database consolidation project is complete, a new database named *Order* will serve as a data store to the new Web store. As part of their daily work, employees also will connect periodically to the *Order* database using a new in-house Web application.

The *HR* database contains sensitive data and is protected using Transparent Data Encryption (TDE). In addition, data in the Salary table is encrypted using a certificate.

### Database Servers

A single server named DB1 contains all the databases at the central office. DB1 is running SQL Server 2008 Enterprise on Windows Server 2003 Enterprise.

### Business Requirements

You need to design an archiving solution for the *Customer* and *Order* databases. Your archival strategy should allow the *Customer* data to be saved for six years.

To prepare the *Order* database for archiving procedures, you create a partitioned table named Order.Sales. Order.Sales includes two partitions. Partition 1 includes sales activity for the current month. Partition 2 is used to store sales activity for the previous month. Orders placed before the previous month will be moved to another partitioned table named Order.Archive. Partition 1 of Order.Archive includes all archived data. Partition 2 remains empty.

A process needs to be created to load the inventory data from each of the 16 wineries by 4 A.M. daily.

Four large customers submit orders using Coho Vineyards Extensible Markup Language (XML) schema for Electronic Data Interchange (EDI) transactions. The EDI files arrive by 5 P.M. and need to be parsed and loaded into the *Customer*, *Accounting*, and *Inventory* databases, which each contain tables relevant to placing an order. The EDI import routine is currently a single threaded C++ application that takes between three and six hours to process the files. You need to finish the EDI process by 5:30 P.M. to meet your Service Level Agreement (SLA) with the customers. After the consolidation project finishes, the EDI routine loads all data into the new *Order* database.

You need to back up all databases at all locations. All production databases are required to be configured with the Full recovery model. You can lose a maximum of five minutes of data under a worst-case scenario. The *Customer*, *Account*, *Inventory*, *Promotions*, and *Order* databases can be off-line for a maximum of 20 minutes in the event of a disaster. Data older than six months in the *Customer* and *Order* databases can be off-line for up to 12 hours in the event of a disaster.

Answer the following question.

- What policies would you implement to check and enforce the business requirements for Coho Vineyard?

## Suggested Practices

---

To help you master the exam objectives presented in this chapter, complete the following tasks.

### Implement Policy Based Management

- **Practice 1** Configure a policy to check the surface area configuration for all your SQL Server instances.
- **Practice 2** Configure a policy to check the last time a database was successfully backed up.
- **Practice 3** Configure policies to check the membership of the sysadmin and db\_owner roles.
- **Practice 4** Configure a policy to ensure that databases are not set to either *auto shrink* or *auto close*.
- **Practice 5** Based on the policies that ship with SQL Server 2008, decide which policies apply to your environment and implement the policy checks.

### Take a Practice Test

---

The practice tests on this book's companion CD offer many options. For example, you can test yourself on just one exam objective, or you can test yourself on all the 70-432 certification exam content. You can set up the test so that it closely simulates the experience of taking a certification exam, or you can set it up in study mode so that you can look at the correct answers and explanations after you answer each question.

#### **MORE INFO PRACTICE TESTS**

**For details about all the practice test options available, see the section entitled "How to Use the Practice Tests," in the Introduction to this book.**

# Index

## Symbols and Numbers

- .ldf files, 40–42
- .mdf files, 40–42
- .ndf files, 40–42
- .NET assemblies, 179
- .NET Framework 2.0, 5
- .NET Framework 3.5, 5
- 32-bit Windows operating systems, 4
- 64-bit Windows operating systems, 4, 413

## A

- Accent sensitivity, 18, 52, 114
- Access statistics, 388–389
- Access, controlling, 50–51, 253, 493
- Accounts
  - Administrative, 265, 271
  - protecting, 265
  - Service Accounts, 17–18, 22, 340–347, 431–433
  - SQL Server Configuration Manager, 19–21
- Active/Active clusters, 431
- Active/Passive clusters, 431
- ActiveX scripts, 234–235
- Ad Hoc Distributed Queries, 260
- ADD SIGNATURE, 277
- Administrative accounts, 265, 271
- AES (Advanced Encryption Standard), 254–255
- Agents, replication, 517–518, 521
- Aggregates, calculating, 9
- Alerts
  - case scenario, designing automation strategy, 246–248
  - creating, 242–244
  - data file alerts, 338
  - Database Mail, 28
  - out-of-space errors, 335
  - practice, creating, 243–244
  - practice, creating failure alerts, 336–338
- Algorithms
  - encryption, 254–255
  - hash algorithms, 293–294, 296–297
  - proportional fill, 40
  - space checking, 49
- Alignment, indexes, 151
- ALL, partition schemes, 142–143
- ALTER, 72
- ALTER DATABASE, 47, 51–52, 163, 285–286
- ALTER FULLTEXT INDEX, 129
- ALTER INDEX, 105
- ALTER INDEX...REBUILD, 47
- ALTER LOGIN, 471, 493, 501–502
- ALTER PARTITION SCHEME, 142–143, 150–151
- ALTER TABLE, 72
- ALTER TABLE...REBUILD, 72–73
- Analysis Services, 11, 14, 234–235, 319
- AND, 79, 324
- ANSI\_NULL\_DEFAULT, 69
- Antivirus software, clustering and, 414
- API (application programming interface), 263–265, 274
- Application Event logs, 332–335
- Application programming interface (API), 263–265, 274
- Application roles, 267
- ApplicationName, 322–323
- Applications, polling, 309–310
- Architecture
  - indexes, 87–91
  - log shipping, 484–487
  - processors, 65
  - replication, 515–517, 526–527, 529–532
  - System Monitor, 309
  - TCP endpoints, 252–256

- Archives, auto options, 49
- Articles, replication, 514, 533–535
- AS DEFAULT, 114
- AS PARTITION, 142–143
- AS SNAPSHOT OF, 223–225
- AS, partitions, 138
- Asymmetric keys, 207–208, 263–266, 276
- Auditing. *See also* Traces
  - C2 auditing, 288
  - database failures, diagnosing, 332–335
  - DDL triggers, 285–286
  - loginless users, 266–267
  - ownership chain, 274–275, 277
  - practice, creating an audit specification, 289–290
  - security audit event group, 320–321, 326–327
  - specifications, 286–288

#### Authentication

- best practices, 254
- database failures, diagnosing, 332–333
- Database Mail, 28–29
- modes, 18
- principals, 263–268
- TCP endpoints, 254–256

#### AUTHORIZATION, 114

- Auto grow, 335
- AUTO options, 48–50
- AUTO\_CLEANUP, 50
- AUTO\_CLOSE, 48–50
- AUTO\_CREATE\_STATISTICS, 48–50, 388
- AUTO\_SHRINK, 48–50
- AUTO\_UPDATE\_STATISTICS, 48–50
- AUTO\_UPDATE\_STATISTICS\_ASYNC, 48–50

#### Automating SQL Server

- alerts, creating, 242–244
- analysis automation, 370
- case scenario, designing automation strategy, 246–248, 364–365
- creating jobs, 234–240

#### Availability. *See* Database Mirroring;

- Failover clustering; Log shipping; Replication

## B

#### BACKUP LOG, 216

#### BACKUP MASTER KEY, 276

#### Backups. *See also* Restore

- case scenario, 229–230
- certificates and master keys, 207–208

- compression, 13

- Database Mirroring, 465–466

- Database Snapshots, 223–226

- differential backups, 204–205, 213–214, 220–221

- filegroups, 205

- FILESTREAM, 70–71

- full backups, 200–203, 213–214, 220

- log shipping, 490–492, 501

- maintenance plans, 206–207

- media errors, 218–219

- page corruption, 206

- partial backups, 205–206

- partitioned tables and indexes, 147

- practice, backing up database, 209–210

- recovery options, 46–48

- replication, 527–528

- security of, 199, 208

- space issues, 334–335

- storage of, 208

- striped backups, 201

- tempdb, 44

- transaction log, 203–204, 213–214, 221

- validating, 209

- Balanced trees (B-trees), 88–89

#### BCP (Bulk Copy Program)

- constraints, 78

- importing/exporting data, 163–165

- page compression, 72–73

- practice, exporting data, 167–170

- query performance, 96

- recovery performance, 47

- replication, 519, 522–523

#### Benchmarks, 43

- Best effort restore, 218–219

#### Best practices

- application APIs, 274

- authentication, 254

- automating analysis, 370

- backup intervals, 498

- balancing security and manageability, 433

- clustering, 414

- counter logs, 311

- deadlocks, handling, 359

- encryption, 255

- failback from a standby, 503

- falling back, 472

- High Safety operating mode, 458

- indexes, determining which to create, 389

- log shipping, 490–491
- logging job steps, 235
- Microsoft Distributed Transaction Coordinator (MS DTC), 428
- out-of-space errors, 335
- page verification, 54
- port numbers, 254
- private network connections, 415
- recovery models, 47
- replication validation, 533
- SQL Browser service, 432
- SQL Server services, managing, 340
- trace files, 319, 325
- Bidirectional topology, replication, 517, 531–532
- BIDS (Business Intelligence Development Studio), 11
- BIGINT, 64
- BINARY, 67
- Binary checksum, 533
- Binary data, 67
- Binary large objects (BLOBs), 70–71
- BIT, 67
- BLOBs (binary large objects), 70–71
- Blocking issues, resolving, 355–359
- blocking\_session\_id, 390
- Boundary points. *See also* Endpoints
  - partitions, 138, 150
  - trace events, 323
- Bounding box, 101
- Broker, event group, 320
- Browser services, best practices, 432
- Buffers, 40–41, 212–214
- Bulk Copy Program (BCP).
  - See* BCP (Bulk Copy Program)
- BULK INSERT, 47, 72–73, 78, 96, 165–166, 522–523
- bulkadmin, 265
- Bulk-logged recovery model, 47, 465, 522–523
- Business Intelligence Development Studio (BIDS), 11

## C

- C2 auditing, 288
- Caches
  - disk subsystem, 41
  - query, 49, 51, 389–390
- Calling stack, 274–275
- CASCADE, 78

- Case scenarios
  - automation strategy, designing, 246–248, 364–365, 403–405
  - backup strategies, creating, 229–230
  - configuring databases, 57–59
  - data management tasks, 81–82
  - distributing and partitioning data, 157–160
  - full text indexing, 133
  - high availability, planning for, 445–448, 477–480, 507–511, 546–550
  - importing data, 173–175
  - indexes, data management, 108–110
  - Policy Based Management strategy, designing, 193–195
  - security, designing, 303–304
  - SQL Server infrastructure, defining, 33–34
- Case sensitivity, 18, 52, 164
- Catalogs, full text, 117–118. *See also* Full text indexing
- CDOC (Cross Database Ownership Chaining), 260
- Central Management Server, 183
- Central publisher topology, 516
- Central subscriber topology, 516–517
- Certificates
  - backups, 207–208
  - endpoints, 254–256
  - log shipping, 492–494
  - logins, 263–266
  - master keys and, 276
  - overview, 276–277, 294
  - practice, encrypting data with, 298–299
- Chains, ownership, 274–275, 277, 281–283
- Change data capture, 13
- Change tracking, 50, 115–116, 129, 526–528, 538–541
- CHANGE\_RETENTION, 50
- CHAR, 66, 114
- Character data, 9, 18, 65–66, 69–70, 164–165
- Check constraints, 78–79, 165, 519
- CHECK\_EXPIRATION, 264
- CHECK\_POLICY, 264
- Checkpoints, 40–41, 212–214, 527–528
- CHECKSUM, 54, 202–203, 206
- Classification, resources, 376, 380–381
- Cleanup agent, 527
- Client connections
  - query statistics, 389–390
  - Resource Governor, 376–377

- SQL Server Configuration Manager, 20
  - TCP endpoints, 252–256
- Client redirect, 458–459
- Client tools, 319
- ClientProcessID, 322–323
- CLOSE MASTER KEY, 276
- CLR (Common Language Runtime), 13, 68–69, 260
- CLR Enables, 260
- CLR event group, 320
- Clustering
  - case scenario, planning for high availability, 445–448
  - cluster analysis warnings, 427
  - cluster failover, 431–434
  - cluster groups, 411, 416
  - cluster name, 411
  - cluster node, 410–411
  - clustering key, 93–95
  - components, 410–411, 431–433
  - disk configuration, 413–414
  - health checks, 433
  - indexes, 93–95, 519
  - network configuration, 414–415
  - practice, creating a Windows cluster, 418–428
  - practice, installing a failover clustered instance, 435–442
  - primary keys, 77
  - resources, 411, 415–416
  - security, 413
  - service failures, 343
  - SQL Server 2008 editions, 13
  - SQL Server instances, 434
  - types of clusters, 412–413
  - Windows clustering, 410–428
- CodePlex, policies, 183
- Collation, 18, 52, 69–70
- Columns. *See also* Tables
  - collation sequences, 18
  - column filter, 515
  - computed, 72
  - full text indexes, 114–115
  - properties of, 69–72
  - trace events, selecting, 322–323
- Command-line switches, Bulk Copy Program (BCP), 163–165
- Common Criteria, auditing, 288
- Common Language Runtime (CLR), 13, 68–69, 260

- Communication protocols
  - global communications thread, 459
  - SQL Server Configuration Manager, 20
  - TCP endpoints, 252–256
- Comparisons
  - collation sequences, 18
  - conditions, 180
- Completion time, 390
- Compliance. *See* Policy-based management
- Compression
  - backups, 202
  - practice, compressed backups, 209
  - SQL Server 2008 editions, 13
  - tables, 72–73, 77
- Conditions, 180, 184–191
- Configuration Manager, 19–21
- Configuring
  - Database Mail, 28–30
  - databases, configuring, 46–52, 57–59
  - files and filegroups, 39–44
  - instances, 17–21
  - Performance Data Warehouse, 397–399
  - surface area, 259–261
- Conflicts, data, 521–523
- CONNECT permission, 253
- Connections
  - query statistics, 389–390
  - Resource Governor, 376–377
  - TCP endpoints, 252–256
- Constraints
  - check constraints, 78–79, 165
  - clustered indexes, 95
  - default, 78
  - foreign keys, 77–78
  - practice, implementing, 79–80
  - precedence, 9
  - primary keys, 77
  - replication, 519
  - tables, 77–79
  - unique, 78, 95, 519
- CONTAINS, 120–123, 127–128
- CONTAINS FILESTREAM, 43
- CONTAINSTABLE, 120, 123, 127–128
- Contention issues, diagnosing, 387–392
- CONTINUE\_AFTER\_ERROR, 219
- CONTINUE\_PAST\_ERROR, 202–203
- Control flow, 235
- Controllers, disk subsystem, 41

COPY\_ONLY, 204–205  
 Copy-On-Write, 224–225  
 Corrupt pages, 54, 206, 217–218.  
   *See also* Validation  
 Corruption, logging, 332–333  
 Counters  
   logs, 310–315  
   performance counters, 312–313  
   Performance Data Warehouse, 395–396  
   practice, creating counter log, 313–315  
   System Monitor, 242, 309–310  
 Covering indexes, 96  
 CREATE DATABASE, 52, 223–225  
 CREATE INDEX, 47, 522–523  
 CREATE PARTITION SCHEME, 142–143  
 CREATE SCHEMA, 63  
 CREATE USER, 266  
 Creating  
   alerts, 242–244, 336–338  
   backups, 209–210  
   classification function, 380–381  
   conditions, 184–191  
   counter logs, 313–315  
   Database Snapshots, 223–225  
   databases, 44  
   failure alerts, 336–338  
   indexes, 98, 102, 114–115, 117–118, 146–147, 389  
   jobs, 234–240  
   log shipping jobs, 490–492  
   logins and database users, 268–269  
   master keys, 281  
   operators, 237–240  
   partition functions, 139–141  
   partition schemes, 142–144  
   partitioned tables and indexes, 146–147  
   principals, 263–268  
   publications, replication, 534–535  
   reports, 10  
   resource pools, 379  
   service accounts, 22  
   subscription, replication, 535–536  
   tables, 63, 73–75  
   traces, 326–330  
   Windows cluster, 418–428  
 Credentials, 28–29, 260, 263–265  
 Cross Database Ownership Chaining (CDOC), 260  
 Cubes, OLAP (Online Analytic Processing), 11  
 Cursors, 44, 320, 335  
 Curves, spatial data, 68–69

## D

d parameters, 346–347  
 Data. *See also* Filegroups; *also* Files  
   compression, 13, 72–73, 77, 202, 209, 212–213  
   conflicts, 521–523  
   distributing  
     case scenario, 157–160  
     partition functions, 137–141  
     partition schemes, creating, 142–144  
     partitions, managing, 150–154  
     practice, partitioning, 140–141, 144, 148, 154–155  
     tables and indexes, partitioning, 146–148  
   encryption. *See* Encryption  
   exporting, 163–170, 183–184  
   file alerts, 338  
   flow tasks, 9  
   importing, 163–166, 173–175, 183–184, 190–191  
   lookups, 9  
   loss exposure, 490  
   movement and manipulation of, 9  
   retrieval. *See* Indexes  
   sources, 166  
   types, tables, 64–69  
 Data Collector, 395–396  
 Data definition language (DDL), 47, 181, 285–286, 492–494  
 Data manipulation language (DML), 47, 180  
 Data Mining, 11, 14  
 Database Engine, 12, 17–18  
 Database Engine Tuning Advisor (DTA), 369–374  
 Database event group, 320  
 Database integrity, 54–56  
 Database Mail, 28–30, 260, 332–335  
 Database master keys (DMKs), 207–208, 275–276, 281  
 Database Mirroring  
   caching, 458  
   case scenario, plan for high availability, 477–480  
   client redirect, 458–459  
   common arguments, 254–256  
   corrupt pages, 55, 206  
   endpoints, 454–455, 460–462, 468  
   failback, designing, 472–473  
   failover sessions, designing, 471  
   FILESTREAM, 71  
   initializing, 464–466  
   log shipping, 484  
   operating modes, 455–458

- practice
  - configuring, 467–469
  - establishing endpoints, 460–462
  - failover, 474
- recovery model, 465
- roles, 452–454
- snapshots, 459
- SQL Server 2008 editions, 13
- threading, 459
- Database reverse, 225–226
- Database roles, fixed, 267
- Database Snapshots, 223–226, 459
  - FILESTREAM, 71
  - practice, creating, 226
  - SQL Server 2008 editions, 13
  - storage space, 335
  - tempdb, 44
  - traces, 396
- Database space issues, 334–335
- DATABASE\_MIRRORING, 252
- DatabaseID, 322
- DatabaseName, 322–323
- Data-driven subscriptions, 14
- Date data, 4, 66, 236
- DATETIME, 66
- Datetime stamps, 50
- DATETIME2, 66
- DATETIMEOFFSET, 66
- db\_accessadmin, 267
- db\_backupoperator, 267
- db\_datareader, 267
- db\_datawriter, 267
- db\_ddladmin, 267
- db\_denydatareader, 267
- db\_denydatawriter, 267
- db\_owner, 50, 267, 271
- db\_securityadmin, 267
- DBCC CHECKALLOC, 55
- DBCC CHECKCATALOG, 55
- DBCC CHECKDB, 55
- DBCC CHECKIDENT, 70
- DBCC CHECKTABLE, 55
- dbcreator, 50, 265
- dbo.sysmail\_log, 333
- DCM (Differential Change Map), 204–205
- DDL (data definition language), 47, 181, 285–286, 492–494
- Deadlocking issues, resolving, 355–361
- Debugging reports, 10. *See also* Troubleshooting
- DECIMAL (P,S), 64
- Decimal number data, 64–65
- Declarative Management Framework.
  - See* Policy-based management
- Decryption, 208, 276, 292–296, 413.
  - See also* Encryption
- DEFAULT, 42, 114
- Default constraints, 78
- Default resource pools, 377
- Defense in depth, 251
- Defragmenting, 97–98, 104–105, 388–389
- DELETE, 74, 78, 96, 225
- DENY, 272
- Deprecation event group, 320
- Device activation errors, 345–349, 351
- Devices, hardware requirements, 4
- Diacritics, 128
- Differential backups, 204–205
  - practice, creating, 210
  - practice, restoring, 220–221
  - restoring, 216–217, 220–221
  - transaction log, 213–214
- Differential Change Map (DCM), 204–205
- Digital signatures, backups, 208
- DISABLED, endpoints, 253
- Disabling, indexes, 105
- Disk Queue Length, 312–313
- diskadmin, 265
- Disks
  - cluster configuration, 413–414, 431–433
  - failures, diagnosing, 351
  - space, 3, 41–43
  - subsystem statistics, 387, 390–391
- distrib.exe, 518
- Distribution Agent, 518–519, 526–528
- Distribution database, 347
- Distribution statistics, 96–97
- Distributor, replication, 515–516
- DMFs (dynamic management functions), 387–393
- DMK (database master keys), 207–208, 275–276, 281
- DML (data manipulation language), 47, 180
- DMVs (dynamic management views), 387–393
- Downloading files, 9.
  - See also* SQL Server Integration Services (SSIS)
- DTA (Database Engine Tuning Advisor), 369–374

dta.exe, 370  
 Dynamic management functions (DMFs), 387–393  
 Dynamic management views (DMVs), 387–393  
 Dynamic row filter, 515

## E

e parameter, 346–347  
 EKM (Extensible Key Management), 296  
 Elapsed time, 390  
 EMERGENCY, 50, 217–219  
 Encryption  
   clusters, 413  
   endpoints, 254–256  
   hierarchy, 276  
   log shipping, 492–494  
   master keys, 207–208, 276  
   overview, 292–296  
   practice, 296–300  
   SQL Server 2008 editions, 13  
   SQL Server Configuration Manager, 20  
 Endpoints  
   Database Mirroring, 454–455, 468  
   objects, copying, 492–494  
   practice, establishing Database Mirroring endpoints,  
     460–462  
   TCP, 252–257  
 EndTime, 323  
 Enumeration, 20  
 Error codes  
   damaged pages, 48  
   deadlock, 359  
   disk space, 334–335  
 Error logs  
   corrupt pages, backup, 206  
   damaged pages, 48  
   database failures, diagnosing, 332–335  
   disk failure, diagnosing, 351  
   parameters, 346–347  
   service failures, diagnosing, 342  
 errorlog, 332  
 Errors. *See also* Error codes; *also* SQL Server Profiler  
   database failures, diagnosing, 332–335  
   device activation errors, 345–347  
   diagnosing, 317  
   import/export of data, 163, 166  
   severity level, 242  
 Errors and Warnings event group, 320  
 Estimated completion time, 390  
 ETL (Extract, Transform and Load) applications, 8  
 Event groups, 286  
 Event handlers, 9  
 Event logs, 48, 332–335  
 EVENTDATA ( ), 286  
 Events  
   alerts and, 242–244  
   Analysis Services, 319  
   notification policies, 181  
   trace events  
     applying filters, 323–324  
     data columns, 322–323  
     managing, 324–325  
     managing event groups, 324–325  
     specifying, 320–322  
 Exclusive locks, 355–356  
 Executable files, job steps, 234–235  
 EXECUTE AS, 275, 322–323  
 Execute job, 234–235, 242  
 Execution, query statistics, 389–390  
 Exporting data, 9, 163–170, 183–184.  
   *See also* SQL Server Integration Services (SSIS)  
 Extensible Key Management (EKM), 296  
 Extensible Markup Language.  
   *See* XML (Extensible Markup Language)  
 Extensions, file, 40–42  
 External Key Management, 260  
 Extract, Transform and Load (ETL) applications, 8  
 Extractions, 14

## F

Facets, 179–180  
 Failback, 472–473, 484–487, 502–503  
 Failed login attempts, 265. *See also* Logins  
 Failover. *See also* Database Mirroring  
   forced, 473  
   graceful, 472–473  
   log shipping, 484–487, 501–502  
   practice, Database Mirroring, 474  
   practice, log shipping failover, 504  
 Failover clustering. *See also* Database Mirroring  
   case scenario, planning for high availability, 445–448  
   cluster analysis warnings, 427  
   cluster failover, 433–434

- cluster groups, 416
- cluster resources, 415–416
- components, 410–411, 431–433
- disk configuration, 413–414
- health checks, 433
- network configuration, 414–415
- practice, creating a Windows cluster, 418–428
- practice, installing a failover clustered instance, 435–442
- resources, 411
- security, 413
- service failures, diagnosing, 343
- SQL Server instances, 434
- types of clusters, 412–413
- windows clustering, 410–428
- Fibre drives, 413, 416
- File extensions, 40–42
- File Transfer Protocol (FTP), 9.
  - See also* SQL Server Integration Services (SSIS)
- FILEGROUP, 113
- FILENAME, 43
- Files/filegroups. *See also* Data
  - backups, 205
  - configuring, 39–44
  - default, 42
  - file growth, 335
  - partition schemes, 142–143
  - PRIMARY, 41
  - size of, 39
- FILESTREAM, 70–71
  - compression, 73
  - configuring, 43
  - Database Snapshots, 223–226
  - Filestream Access Level, 260
  - indexes full text, 114
  - Window Server 2008 Server Core, 4
- FILLFACTOR, 97–98, 104
- Filters
  - indexes, 97, 114–115
  - noise words, 128
  - replication, 514–515
  - TCP endpoints, 252
  - trace events, 323–324
- Firewalls, TCP endpoints, 252–256
- Fixed database roles, 267
- FLOAT (N), 64–65, 70
- fn\_trace\_geteventinfo, 324
- fn\_trace\_getfilterinfo, 324
- fn\_trace\_getinfo, 324
- fn\_trace\_gettable, 324, 396
- Folder paths, 43
- Folders, Snapshot, 519
- FOR PATH, 99
- FOR PROPERTY, 99
- FOR VALUE, 99, 138
- Forced failover, 473
- Foreign keys, 77–78, 519
- FORMSOF, 122–123
- FORMSOF THESAURUS, 127–128
- Forwarding messages, endpoints, 255–256
- Forwarding pointers, 94–95
- Fragmentation, 97–98, 104–105, 388–389
- FREETEXT, 120–122, 127–128
- FREETEXTTABLE, 120–121, 127–128
- FTDATA, 127–128
- FTP (File Transfer Protocol), 9. *See also* SQL Server Integration Services (SSIS)
- FULL, 129
- Full backups, 200–203. *See also* Backups
  - creating, 209
  - restoring, 214–216, 220
  - transaction log, 213–214
- Full recovery model, 47, 465
- Full text catalogs, 117–118, 225
- Full Text event group, 320
- Full text indexing
  - case scenario, full text indexing, 133
  - catalogs, full text, 113–114
  - cluster failover, 434
  - creating, 114–115
  - Database Snapshots, 224
  - managing, 127–129
  - populating, 128–129
  - practice, full text queries, 123–125
  - practice, managing full text indexes, 129–131
- Full text search, 12
- Fuzzy groupings and lookups, 9, 14

## G

- Generations, merge replication, 540
- Geography, spatial indexes, 68–69, 99–101
- Geometry, spatial indexes, 68–69, 99–101
- GeometryCollection, 69
- Global communications thread, 459

Graceful failover, 472–473  
 GRANT, 272  
 GRANT/REVOKE, 454  
 GROUP BY, 335  
 Grouping
 

- cluster groups, 416
- event groups, 286
- failover, 432–433, 471
- fixed server roles, 265
- SSIS (SQL Server Information Services) and, 9
- tempdb, 44
- trace event groups, 320–322, 324–325
- user database roles, 268
- Windows group logins, 263–265

## H

Hacking, prevention of, 253–254, 454.  
*See also* Security  
 Hardware. *See also* Failover clustering
 

- cluster compatible, 410
- dynamic management views (DMVs), 387, 391–392
- failures, diagnosing, 332–333, 351–352
- requirements, 3–6, 12

 Hash algorithms, 293–294, 296–297, 335  
 Health checks, 415, 433, 484–487  
 Heap, 94  
 Hierarchy, counters, 309–310  
 Hierarchy, encryption, 276  
 HIERARCHYID, 69  
 High Availability operating mode, 455–457  
 High Performance operating mode, 457, 473  
 High Safety operating mode, 457–458  
 Histograms, 96–97  
 History, jobs, 235–236  
 HostName, 322–323  
 Hot add memory/CPU, 13  
 HTML (Hypertext Markup Language), 114  
 HTTP transports, 252

## I

I/O (input/output), 41, 139, 390–391  
 IDE (Integrated Development Environment), 413  
 IDENTITY, 70  
 IGNORE\_DUP\_KEY, 98

Immediate Updating Subscriber option, 528–529  
 Impersonation, 275  
 Import/Export Wizard, 14, 166  
 Importing data, 9. *See also* SQL Server Integration Services (SSIS)
 

- Bulk Copy Program (BCP), 163–166
- BULK INSERT, 165–166
- case scenario, 173–175
- policies, 183–184, 190–191
- practice, importing policies, 190–191
- trace files, 329

 IN PATH, 113  
 INCLUDE, 96  
 INCREMENTAL, 129  
 Indexed views
 

- case scenario, 157–160
- creating partition functions, 137–141
- Database Engine Tuning Analysis (DTA), 372
- managing partitions, 150–154
- partition schemes, creating, 142–144
- practice, partitioning, 140–141, 144, 148, 154–155
- SQL Server 2008 editions, 13
- tables and indexes, 146–148

 Indexes
 

- architecture, 87–91
- auto options, 49
- backup maintenance, 206–207
- case scenario
  - data management, 108–110
  - full text indexing, 133
- change tracking, full text, 115–116
- cluster failover, 434
- clustered indexes, 93–95
- covering indexes, 96
- designing, 93
- dynamic management views (DMVs), 387
- filters, 97
- full text catalogs, 113–114
- full text indexes, 114–115
- included columns, 96
- language, word breakers, and stemmers, 116
- maintenance of, 95–96, 104–105
- managing full text indexes, 127–129
- nonclustered indexes, 95
- online index creation, 98
- options, 97–98
- partitioning, 146–148
  - case scenario, 157–160
  - creating partition functions, 137–141

- managing partitions, 150–154
- partition schemes, creating, 142–144
- practice, partitioning, 140–141, 144, 148, 154–155
- populating full text indexes, 128–129
- practice
  - creating full text indexes, 117–118
  - creating indexes, 102
  - evaluating missing indexes, 392–393
  - full text queries, 123–125
  - maintaining, 106
  - managing full text indexes, 129–131
  - querying full text data, 120–123
  - spatial, 99–101
  - statistics on, 96–97, 388–389
  - tempdb, 44
  - work table space, 335
  - XML indexes, 99
- INFLECTIONAL, 122–123
- INIT/NOINT, 202
- Input/output (I/O), 41, 139, 390–391
- INSERT, 72–74, 78, 96, 163, 225
- INSERT INTO...WITH (TABLOCK), 72–73
- Installing
  - Database Engine, 17–18
  - failover clustered instance, 435–442
  - hardware requirements, 3–4
  - instances, 17–26
  - Server Agent, 17–18
  - Service Accounts, 17–18
- Instances
  - auditing, 285–290
  - clustering, 431–434
  - collation sequences, 18
  - installing and configuring, 17–21, 332–333
  - Performance Data Warehouse, 395–396
  - practice, installing, 22–26
  - standardizing. *See* Policy-based management
- Instance-to-disk ratio, 432
- INT, 64
- Integer data, 64–65
- Integrated Development Environment (IDE), 413
- Integration Services, 234–235
- Integrity, maintaining, 54–56. *See also* Validation
- Internal networks, clustering, 414–415
- Internet Explorer, requirements, 5
- IOStall, 390–391
- IP addresses, 416, 431–433
- IsAlive test, 433
- Isolation levels, transactions, 356

**J**

- Jobs
  - automating analysis with, 370
  - case scenario, 246–248
  - creating, 234–236
  - history, 235–236
  - log shipping, 490–492
  - operators, 236
  - practice, creating jobs and operators, 237–240
  - schedules, 235
  - steps, 234–235
- JOIN, 74
- Join filter, 515

**K**

- Kana sensitivity, 18
- Kerberos, 254–256, 413
- KEY INDEX, 115
- Key Performance Indicators (KPI), 11
- Keys
  - asymmetric, 263–266, 276
  - backups, 207–208
  - clustering key, 93–95
  - database master keys (DMK), 207–208, 275–276, 281
  - EKM (Extensible Key Management), 296
  - External Key Management, 260
  - foreign keys, 77–78, 519
  - partitioning, 146–147
  - practice, creating and managing master keys, 281
  - primary keys, 77, 95, 519, 521
  - service master, 276
  - symmetric, 276, 294, 297
- KPI (Key Performance Indicators), 11

**L**

- l parameter, 346–347
- LANGUAGE, 115, 120
- Languages
  - COLLATE, 69–70
  - collation sequences, 18, 52
  - indexes, full text, 114, 116
- Latching, statistics on, 388–389
- Latitude, spatial data, 68–69

Leaf level pages, 88–89, 97–98. *See also* Pages  
 LIKE, 323  
 Lines, spatial data, 68–69  
 LineString, 69  
 Linked servers, 492–494  
 LISTENER\_IP, 253–254  
 LISTENER\_PORT, 253–254  
 Local quorum, 412–413  
 Lock Manager, 355–356  
 Locks  
   escalation, 356  
   event group, 320–321  
   service failures, diagnosing, 344  
   statistics on, 388–389  
   troubleshooting, 355–356  
 Log File Viewer, 333  
 Log Reader Agent, 518, 526–528  
 Log Sequence Number (LSN), 203, 212–214, 217,  
   526–527  
 Log shipping  
   case scenarios, 507–511  
   components, 485–486  
   failback, 502–503  
   failover, 458, 501–502, 504  
   initializing, 489–494  
   practice, failover, 504  
   practice, setting up, 494–498  
   scenarios for, 484–485  
   types of, 487  
 Loginless users, 266–267  
 LoginName, 322–323  
 Logins, 253  
   database failures, diagnosing, 332–335  
   Database Mirroring, 466  
   database users, 266  
   DDL triggers, 285–286  
   failover, migrating, 471  
   log shipping, 492–494  
   practice, creating, 268–269  
   principals, 263–265  
   Resource Governor, 377  
   surface area configuration, 259  
   timeout, 518  
   trace events, 322–323  
 LoginSID, 322–323  
 logread.exe, 518  
 Logs  
   alerts, 337–338  
   audit trail, 286

  compression, 13  
   counter logs, 310–312  
   database failures, diagnosing, 332–335  
   files and filegroups, 39–44  
   jobs, 234–235  
   minimally logged transactions, 522–523  
   practice, creating counter log, 313–315  
   space issues, 334–335  
   trace files, 319  
 Longitude, spatial data, 68–69  
 LooksAlive test, 415, 433  
 Lookups, 9  
 LSN (Log Sequence Number), 203, 212–214, 217,  
   526–527

## M

Mail  
   database failures, diagnosing, 332–335  
   Database Mail, 28–30, 260  
   SQL Mail, 260  
 Maintenance  
   backups, 206–207  
   indexes, 95–96, 104–105  
   log shipping and, 491–492  
   practice, maintaining indexes, 106  
 Majority node set cluster, 412–413  
 Manage Schedules, 235  
 Management, Policy Based  
   case scenario, designing a strategy, 193–195  
   categories, policies, 181  
   Central Management Server, 183  
   compliance, policies, 181–183  
   conditions, 180  
   DDL triggers, 286  
   facets, 179–180  
   import and export policies, 183–184  
   overview, 177–179  
   policies, overview, 181  
   policy targets, 180  
   practice, defining policies, 184–191  
   SQL Server editions, 13  
 Mandate, 181  
 Mantissa, 65  
 MARK, 217  
 MARS (Multiple Active Result Sets), 335  
 Master data file, 346–347  
 Master database, device activation errors, 345–347

Master keys, 207–208, 275–276, 281  
 Master transaction log file, 346–347  
 Matches, exact, 122  
 MAX, 66–67  
 MD2, 293  
 MD4, 293  
 MD5, 293  
 MDAC (Microsoft Data Access Components), 5  
 Media errors, restore, 218–219  
 Memory  
   auto options, 49  
   buffers, 40–41, 212–214  
   hardware failures, diagnosing, 352  
   requirements, 3  
   Resource Governor and, 376–378  
   SQL Server 2008 editions, 12  
   tables, 64  
 MERGE, 74, 150, 154–155, 225  
 Merge Agent, 518  
 Merge replication, 13, 518, 520–521, 538–542  
 Message forwarding, 255–256  
 Messages, 9  
   configuring Database Mail, 28–29  
   database failures, diagnosing, 332–335  
   practice  
     Database Mail, configuring, 29–30  
     queues, 29  
 Metadata operations, 152–154,  
   273–274, 458, 540  
 Microsoft Data Access Components (MDAC), 5  
 Microsoft Distributed Transaction Coordinator  
   (MS DTC), 414, 428  
 Microsoft Internet Explorer, requirements, 5  
 Microsoft Message Queue (MSMQ), 9.  
   *See also* SQL Server Integration Services (SSIS)  
 Microsoft Office Integration, 14  
 Microsoft Windows, supported versions, 4  
 Migrating logins, 471  
 Migrating, log shipping, 484  
 Minimally logged transactions, 522–523  
 Mirror role, Database Mirroring, 453.  
   *See also* Database Mirroring  
 MIRROR TO, 201–202  
 Mirroring, Database. *See* Database Mirroring  
 Mobile users, merge replication, 538–541  
 Monetary data, 64–65  
 MONEY, 64–65  
 Monitor server, 486

Monitoring SQL Server  
   blocking and deadlocking, 355–359  
   case scenario, designing automation strategy,  
     364–365  
   correlating performance and monitoring data, 325,  
     329–330  
   counter logs, 310–312  
   database failures, diagnosing, 332–335  
   dynamic management views (DMVs), 387–392  
   hardware failures, diagnosing, 351–352  
   overview, 308  
   overview, System Monitor, 309–310  
   performance counters, 312–313  
   practice  
     creating counter log, 313–315  
     creating failure alerts, 336–338  
     creating traces, 326–330  
     deadlocks, troubleshooting, 359–361  
     service startup errors, troubleshooting, 348–349  
 Replication Monitor, 532  
 Server Profiler  
   data columns, selecting, 322–323  
   filters, 323–324  
   overview, 317  
   trace events, specifying, 320–322  
   traces, defining, 317–319  
   traces, managing, 324–325  
   service failures, diagnosing, 340–347  
   space issues, 334–335  
 MS DTC (Microsoft Distributed Transaction  
   Coordinator), 414, 428  
 MSCS directory, 414  
 msdb database, 28, 179  
 MSlogreader\_history, 526–527  
 MSmerge\_contents, 538–540  
 MSmerge\_genhistory, 539–540  
 MSmerge\_tombstone, 538–540  
 MSMQ (Microsoft Message Queue), 9  
 MSrepl\_commands, 526  
 MSrepl\_transactions, 526  
 mssqlsystemresource, 347  
 MULTI\_USER, 50  
 MultiLineString, 69  
 Multiple Active Result Sets (MARS), 335  
 Multiple instances, 13, 431.  
   *See also* Instances  
 Multipoint, 69  
 MultiPolygon, 69

## N

Named Pipes, 5  
 Names, login, 263–265  
 NAS (Network Attached Storage), 351  
 NCHAR (n), 66  
 NEAR, 123  
 NEGOTIATE, 254–256  
 Network Attached Storage (NAS), 351  
 Network configuration, clusters, 414–415, 431–432  
 Network interface card (NIC), 434  
 Network names, 416, 431–433  
 NIC (network interface card), 434  
 NO REVERT, 275  
 NODE, 99  
 Noise word filters, 128  
 Nonclustered indexes, 95  
 Non-unicode data, 65–66  
 NORECOVERY, 215, 465  
 NOT, 79  
 NOT FOR REPLICATION, 71  
 NOT LIKE, 323  
 Notifications, 242–244, 320  
 Notify operator, 242  
 NOTNULL, 69, 77–78  
 NTDomainName, 322–323  
 NTLM, 254–256  
 NTUserName, 322–323  
 NULL, 69, 71–72, 78, 138  
 Nullability, 69. *See also* NOTNULL; *also* NULL  
 NUMERIC (P,S), 64  
 Numeric data, 64–65, 323  
 NVARCHAR (MAX), 74  
 NVARCHAR (n), 66, 73

## O

Object Linking and Embedding (OLE), 259–260  
 Object Linking and Embedding Database (OLE DB), 166, 493  
 ObjectName, 323  
 Objects  
   change alerts, 28  
   copying, 492–494  
   event group, 320  
   system, copying, 466  
 ODBC (Open Database Connectivity), 493

OFFLINE, 50  
 Offsite storage, backups, 208  
 OLAP (Online Analytic Processing), 11, 14, 96  
 OLE Automation, 259–260  
 OLE DB (Object Linking and Embedding Database), 166, 493  
 OLEDB event group, 320  
 ON, 146–147  
 ON\_FAILURE, 287  
 On change, log only, 181  
 On change, prevent, 181  
 On Demand, 181  
 On schedule, 181  
 ONLINE, 50, 217  
 ONLINE = ON, 98  
 Online Analytical Processing (OLAP), 11, 14, 96  
 Online indexes, 98  
 Online operations, 13  
 Online restores, 217  
 Online Transaction Processing (OLAP), 11, 14, 96  
 Open Database Connectivity (ODBC), 493  
 OPEN MASTER KEY, 276  
 OPENROWSET/OPENDATASOURCE, 259–260  
 Operating systems  
   clustering, 413–414  
   supported versions, 4  
   tasks, job steps, 234–235  
 Operational workflows, 9  
 Operators, jobs, 236–240  
 OR, 79, 324  
 ORDER BY, 335  
 Out-of-space errors, 335  
 Output Queue Length, 312–313  
 Overflow, tempdb, 44  
 Ownership chains, 274–275, 277, 281–283

## P

Package configurations, 9  
 Package designer, 14  
 PAD\_INDEX, 97–98  
 PAGE\_VERIFY, 54  
 PAGE\_VERIFY CHECKSUM, 48  
 Pages  
   B-trees, 88–89  
   compression, 72–73, 77  
   corrupt, 206, 217–218

## Parallel operations

- damaged, recovery, 46–48
  - index levels, 89–91
  - index options, 97–98
  - integrity, maintaining, 54–55
  - locking, 355–356
  - partitioning and, 151–154
  - size calculations, 89–91
  - sorting, 93–95
  - splits, 90, 95, 97–98
  - verification, 54–55
  - Parallel operations, 12
  - Parameterization, 51
  - Parameters, startup, 345–347
  - Partial backups, 147, 205–206.
    - See also* Backups
  - Partial restore, 217. *See also* Restore
  - Partitioning
    - benchmarks, 43
    - case scenario, 157–160
    - Database Engine Tuning Analysis (DTA), 372
    - managing partitions, 150–154
    - partition functions, creating, 137–141
    - partition schemes, creating, 142–144
    - practice, partitioning, 140–141, 144, 148, 154–155
    - SQL Server 2008 editions, 12
    - tables and indexes, 146–148
    - tables, compression of, 73
  - Partitioning key, 146–147
  - PARTNER, 255
  - Passphrase, 297
  - Passwords, 19–21, 263–265, 340–347
  - PATH, 99
  - PATH XML, 99
  - Payloads, endpoints, 252
  - Peer-to-peer replication, 530–531
  - Pending requests, 391
  - PerfMon. *See* System Monitor
  - Performance. *See also* Performance, monitoring;
    - also* Performance, optimizing;
    - also* Query performance
  - alert messages, 28
  - auto options, 48–50
  - backups, 202
  - baseline, establishing, 327–329
  - bulk-logged recovery model, 47
  - Database Mirroring, 455–456
  - disk subsystem, 41
  - forwarding pointers, 94
  - indexes
    - architecture, 87–91
    - case scenario, data management, 108–110
    - case scenario, full text indexing, 133
    - change tracking, full text, 115–116
    - clustered indexes, 93–95
    - covering indexes, 96
    - designing indexes, 93
    - distribution statistics, 96–97
    - filters, 97
    - full text catalogs, 113–114
    - full text indexes, 114–115, 127–129
    - included columns, 96
    - language, word breakers and stemmers, 116
    - maintenance of, 95–96, 104–105
    - nonclustered indexes, 95
    - online index creation, 98
    - options, 97–98
    - practice, creating, 102, 117–118
    - practice, full text queries, 123–125
    - practice, maintaining, 106
    - practice, managing full text indexes, 129–131
    - querying full text data, 120–123
    - spatial indexes, 99–101
    - XML indexes, 99
  - parameterization, 51
  - replication, 527
  - SQL Server 2008 editions, 13
  - tables, 64
- Performance Conditions, 242–244
  - Performance Counter, 395
  - Performance Data Warehouse, 395–399
  - Performance event group, 320–321
  - Performance Studio, 395–396
  - Performance, monitoring. *See also* Performance;
    - also* Performance, optimizing;
    - also* Query performance
  - blocking and deadlocking, 355–359
  - case scenario, designing automation strategy, 364–365
  - correlating performance and monitoring data, 325, 329–330
  - counter logs, 310–312
  - database failures, diagnosing, 332–335
  - dynamic management views (DMVs), 387–392
  - hardware failures, diagnosing, 351–352
  - overview, 308

- overview, System Monitor, 309–310
- performance counters, 312–313
- practice
  - creating counter log, 313–315
  - creating failure alerts, 336–338
  - creating traces, 326–330
  - deadlocks, troubleshooting, 359–361
  - service startup errors, troubleshooting, 348–349
- Replication Monitor, 532
- Server Profiler
  - data columns, selecting, 322–323
  - filters, 323–324
  - overview, 317
  - trace events, specifying, 320–322
  - traces, defining, 317–319
  - traces, managing, 324–325
- service failures, diagnosing, 340–347
- space issues, 334–335
- Performance, optimizing. *See also* Performance;
  - also* Performance, monitoring;
  - also* Query performance
- case scenario, designing an automation strategy, 403–405
- Database Engine Tuning Advisor (DTA), 369–373
- dynamic management views (DMVs), 387–392
- Performance Data Warehouse, 395–396
- practice
  - analyzing query workload, 374
  - Performance Data Warehouse, configuring, 397–399
  - Resource Governor, implementing, 379–385
- Resource Governor, 376–385
- Permissions
  - CREATE SCHEMA, 63
  - device activation errors, 345–347
  - failover clustering, 432–433
  - managing, 271–277
  - practice managing, 278–283
  - replication agents, 535
  - roles and, 268
  - service failures, diagnosing, 343–345
- PERSISTED, 72, 147
- Ping, 415, 433, 456
- Platforms, migrating, 484
- Pointers, forwarding, 94–95, 151–154
- Point-in-time restore, 217
- Points, spatial data, 68–69
- Policy-based management
  - case scenario, designing a strategy, 193–195
  - categories, policies, 181
  - Central Management Server, 183
  - compliance, policies, 181–183
  - conditions, 180
  - DDL triggers, 286
  - facets, 179–180
  - import and export policies, 183–184
  - overview, 177–179
  - policies, overview, 181
  - policy targets, 180
  - practice, defining policies, 184–191
  - SQL Server 2008 editions, 13
- Polling
  - intervals, 518
  - System Monitor, 309–310
- Polygons, spatial data, 68–69
- Ports, 252–254, 454
- Practice
  - alerts, creating, 243–244, 336–338
  - audit specification, creating, 289–290
  - backing up database, 209–210
  - clusters, creating, 418–428
  - constraints, implementing, 79–80
  - counter log, creating, 313–315
  - database integrity, checking, 56
  - Database Mail, configuring, 29–30
  - Database Mirroring, 460–462, 467–469
  - database recovery model, changing, 52
  - databases, creating, 44
  - deadlocks, troubleshooting, 359–361
  - encrypting data, 296–300
  - endpoints, inspecting, 257
  - exporting data, 167–170
  - failover, 435–442, 474
  - indexes, 102, 106, 117–118, 123–125, 129–131
  - jobs and operators, creating, 237–240
  - log shipping, 494–498, 504
  - logins and database users, creating, 268–269
  - merge replication, implementing, 541–542
  - missing indexes, evaluating, 392–393
  - partitioning, 144, 148
  - Performance Data Warehouse, configuring, 397–399
  - permissions, managing, 278–283

- policies, defining and checking compliance, 184–191
- replication, configuring publishing, 523–524
- Resource Governor, implementing, 379–385
- restoring a database, 219–221
- service startup errors, troubleshooting, 348–349
- SPLIT, MERGE, and SWITCH, 154–155
- surface area, configuring, 261
- tables, creating, 75
- traces, creating, 326–330
- transactional replication, implementing, 534–536
  - verify minimum requirements, 5–6

Precedence, constraints, 9

PRIMARY, 41, 217

Primary database, 486

Primary keys, 77, 95, 519, 521

Principal role, database mirroring, 453

Principals, 263–269

Private certificates, 276–277

Private keys, backups, 208

Private networks, best practices, 415

processadmin, 265

Processes, blocking, 357–358

Processor Queue Length, 312–313

Processors
 

- database failures, diagnosing, 332–333
- hardware failures, diagnosing, 352
- numeric precision and, 65
- requirements, 3
- Resource Governor and, 376–378

Profiles, Database Mail, 28–29

Progress Report, event group, 320

Properties, facets, 180

PROPERTY, 99

Proximity, searching, 121–123

Public certificate, 276–277

Public keys, backups, 208

Public networks, clustering, 414–415

Public, role, 267

Publications, replication, 514, 533–535

Publisher, replication
 

- agents, 517–518
- data conflicts, 521–523
- merge replication, 538–541
- overview, 515–516
- practice, configuring publishing, 523–524
- replication methods, 519–521

- topology, 516
- transactional replication
  - architecture, 530–532
  - change tracking, 526–528
  - options, 528–530
  - validation, 532–533

## Q

qrdrsvc.exe, 518

Quarantine, corrupt pages, 206

Queries. *See also* Query performance
 

- Bulk Copy Program (BCP), 164
- dynamic management views (DMVs), 387
- notification event group, 320
- Performance Data Warehouse, 395–396
- performance optimization, 317, 321
- practice, analyzing workload, 374
- Resource Governor, 376
- statistics on, 389–390
- timeout, 518

Query Activity, 395–396

Query Optimizer, 49.
 

- See also* Database Engine Tuning Advisor (DTA)

Query performance. *See also* Queries
 

- distribution statistics, 96–97
- indexes, 96, 99
  - case scenario, full text indexing, 133
  - change tracking, full text, 115–116
  - full text catalogs, 113–114
  - full text data, 120–123
  - full text indexes, 114–115, 128–129
  - language, word breakers and stemmers, 116
  - managing full text indexes, 127–129
  - practice, creating full text indexes, 117–118
  - practice, managing full text indexes, 129–131
- practice, full text queries, 123–125
- query caches, 49, 51, 389–390
- query plans, 49, 51
- XML indexes, 99

queryout, 164

QUEUE\_DELAY, 287

Queues, 391–392
 

- performance counters, 312–313
- Queue Reader Agent, 518, 529–530
- Queued Updating Subscriber, 529–530

Quorum database, 411–413

## R

- RAISERROR, 333
- RANGE LEFT, 138
- RANGE RIGHT, 138
- RANK, 123
- RC4, 254–255
- READ COMMITTED, 356
- Read operations
  - backups, 199
  - Copy-On-Write, 224–225
  - indexes, 97–98
  - isolation levels, 356
  - locking, 355–356
  - query statistics, 389–390
- READ SERIALIZABLE, 357
- READ UNCOMMITTED, 356
- Read/write access, backups, 199
- READ\_ONLY, 50
- READ\_WRITE, 50
- READ\_WRITE\_FILEGROUPS, 205–206
- Read-only files, partial backups, 205–206
- REAL, 64–65, 70
- REBUILD, 105
- Receiving messages, 9
- Recovery. *See also* Backups
  - models for, 46–48, 465, 522–523
  - options, configuring, 46–48
  - practice, changing recovery model, 52
  - restart, 434
  - transaction logs, 42
- RECOVERY, 215
- Redirect, client, 458–459
- REDO, 213
- Redundancy. *See* Database Mirroring; *also* Replication
- Relational data, full text catalogs, 114
- Remapping users, 493
- Remote Admin connections, 260
- Remote connections, 259–260, 341–343, 538–541
- Remote procedure calls (RPC), 415
- REORGANIZE, 105
- REPEATABLE READ, 357
- REPLACE, 215
- Replication
  - agents, 517–518, 521
  - case scenario, planning for high availability, 546–550
  - components, 514–515
  - data conflicts, 521–523
  - failover, 458
  - jobs, 234–235
  - merge replication, 518, 520–521, 538–542
  - methods, 519–521
  - objects, copying, 492–494
  - overview, 514
  - practice
    - configuring publishing, 523–524
    - merge replication, implementing, 541–542
    - transactional replication, implementing, 534–536
  - roles, 515–516
  - SQL Server 2008 editions, 13
  - topology, 516–517
  - transactional replication
    - agents, 518
    - architecture, 530–532
    - change tracking, 526–528
    - minimally logged transactions, 522–523
    - monitoring, 532
    - options, 528–530
    - overview, 520, 526
    - practice, implementing, 534–536
    - validation, 532–533
- Replication Monitor, 532, 536
- Replication watermark, 526–527
- replmerg.exe, 518
- Report Builder, 14
- Report designer, 10
- Report server, 10
- Reports
  - log shipping, 484
  - progress report event group, 320
  - Report Builder, 14
  - SSRS (SQL Server Reporting Services), 10, 14
- Requirements
  - hardware, 3–4
  - practice, 5–6
  - software, 3–5
- Resolution, conflict, 522
- Resources
  - cluster, 411, 415–416
  - contention, diagnosing, 387–392
  - mssqlsystemresource, 347
  - Resource Governor, 13, 376–385
  - resource pools, 376–379, 385

Restart recovery, 42, 213, 434

## Restore

corrupt pages, 217–218

Database Mirroring, 465–466

Database Snapshots, 223–226

differential backups, 216–217

full backups, 214–216

log shipping, 490–492, 501–502

media errors, 218–219

online restores, 217

partitioned tables and indexes, 147

practice, restoring database, 219–221

striped backups, 201

transaction log, 212–214, 216–217

RESTORE, 215

RESTORE DATABASE, 219

RESTORE LOG, 219

RESTORE MASTER KEY, 276

RESTRICTED\_USER, 50

Retry settings, 234

REVERT, 275

Reverting data, snapshots, 225–226

REVOKE, 272

## Roles

Database Mirroring, 452–454

database-level vs server level, 454

fixed, 265, 267

replication, 515–516

user, 267–268

ROLLBACK, 51

ROLLBACK TRANSACTION, 285

Rolling forward, 213

Root pages, 88–89, 97–98. *See also* Pages

Routing, static, 9

Row filter, 515

ROWGUIDCOL, 70

Rows. *See also* Tables

compression, 72–73

data conflicts, 521

locking, 355–356

replication validation, 533

row filter, 515

## RPC

Completed, 322, 369

Starting, 322, 369

RPC (remote procedure calls), 415

Runnable queues, 391–392

Running queues, 391–392

## S

sa account, 265, 271.

*See also* Administrative accounts

Salting a hash, 294

SAN (storage area network), 351, 414, 416

Scalable shared databases, 13

Scale-out reporting, 14

Scans event group, 320

Schedules, 181, 235

## Schemas

EVENTDATA (), 286

permissions, 272–273

replication, 519

sys schema, 387

tables, creating, 63

XML data, 67

SCSI/iSCSI, 413

## Search arguments

full text data, 120

stop lists, 128, 130–131

thesaurus, 122–123, 127–130

wildcards, 122

word forms and proximity, 121–123

Secondary database, 486

Securables, 272

## Security

access, controlling, 50–51, 253, 493

auditing, 286–290, 320–321

authentication modes, 18

backups, 199, 208

best practices, 433

case scenario, designing, 303–304

clusters, 413, 432–433

counter logs, 311

Database Mail, 28–29

Database Mirroring, 454–455

database users, 266

DDL triggers, 285–286

encrypting data, 292–300.

*See also* Encryption

failover, login migration, 471

fixed roles, 265, 267

hacking, prevention of, 253–254, 454

impersonation, 275

jobs, 234

master keys, 275–276

metadata, 273–274

- objects, copying, 492–494
- permissions, managing, 268, 271–283
- practice
  - configuring surface area, 261
  - creating an audit specification, 289–290
  - encrypting data, 296–300
  - inspecting existing endpoints, 257
  - managing permissions, 278–283
  - signatures, adding, 281–283
  - surface area, configuring, 261
- principals, creating, 263–268
- replication agents, 535
- reports, 10
- Security Event logs, 332–335
- service account, 17
- service failures, diagnosing, 340–347
- signatures, 277, 281–283
- stored procedures, 274–275
- surface area, configuring, 259–261
- TCP endpoints, 252–256
- users, loginless, 266–267
- Security Event logs, 332–335
- Security identifier (SID), 263–266, 492
- securityadmin, 265
- SELECT, 50, 163, 224
- SELECT...INTO, 47, 522–523
- Selectivity, indexes, 96–97
- Sending messages, 9
- Server Agent. *See* SQL Server Agent
- Server event group, 320
- Server Profiler. *See* SQL Server Profiler
- Server roles, 265, 454
- serveradmin, 266
- ServerName, 323
- Service Accounts, 431–433
  - installing and configuring, 17–18
  - practice, creating, 22
  - service failures, diagnosing, 340–347
- Service Broker, 8, 252, 254–256
- Service master key, 207, 276
- Services. *See* Service Broker; SQL Server
  - Analysis Services (SSAS); SQL Server Integration Services (SSIS); SQL Server Reporting Services (SSRS); SQL Server services
- SessionLoginName, 322–323
- Sessions, event group, 321
- SET IDENTITY\_INSERT, 70
- Setup, 347
- setupadmin, 266
- SHA, 293
- SHA1, 293
- Shared drive arrays, 431–433
- Shared locks, 355–356
- Shared memory, requirements, 5
- Showplan Statistics Profile, 321
- Showplan Text, 321
- Showplan XML, 321
- SHUTDOWN, 287
- SID (security identifier), 263–266, 492
- signal wait, 391
- Signatures, digital, 208, 277, 281–283
- Simple Mail Transfer Protocol (SMTP), 28
- Simple recovery model, 47–48, 465, 522–523, 527–528
- SINGLE\_USER, 50
- Single-instance cluster, 431
- Size considerations. *See also* Space management
  - backups, 205–206
  - binary data, 67
  - character data, 65–66
  - date and time data, 66
  - decimal data, 65
  - FILESTREAM, 70–71
  - index columns, 96
  - index key, 93
  - nonclustered indexes, 95
  - numeric data, 64
  - page compression, 72–73
  - pages, 89–91
  - statistics on, 388–389
  - tables, 64
  - trace files, 319
- Slammer Trojan, 259
- Small Computer System Interface/Internet
  - Small Computer System Interface (SCSI/iSCSI), 413
- SMALLDATETIME, 66
- SMALLINT, 64
- SMALLMONEY, 64–65
- SMTP (Simple Mail Transfer Protocol), 28
- SNAPSHOT, 357
- Snapshot agent, 518–519
- Snapshot folder, 519
- Snapshot replication, 518–519
- Snapshots, 223–226, 459
  - FILESTREAM, 71
  - practice, creating, 226

- SQL Server 2008 editions, 13
- storage space, 335
- tempdb, 44
- snapshots.trace\_data, 396
- snapshots.trace\_info, 396
- SOAP, 252
- Software requirements, 3–6
- SORT\_IN\_TEMPDB, 97–98
- Sorting
  - clustered indexes, 93–95
  - collation sequences, 18, 52
  - index options, 97–98
  - tempdb, 44
- sp\_article\_validation, 533
- sp\_browsereplcmds, 532
- sp\_configure, 260
- SP\_Counts, 318
- sp\_publication\_validation, 533
- sp\_trace\_create, 324
- sp\_trace\_generateevent, 324
- sp\_trace\_setevent, 324
- sp\_trace\_setfilter, 324
- sp\_trace\_setstatus, 324
- sp\_validatemergepublication, 541
- sp\_validatemergesubscription, 541
- Space management.
  - See also* Size considerations
  - AUTO options, 49
  - backup management, 205–207
  - database space issues, 334–335
  - dynamic management views (DMVs), 387
  - out-of-space errors, 335
  - space checking algorithms, 49
- SPARSE, 71–72
- Sparse files, 224
- SPATIAL, 4
- Spatial data, 13, 68–69
- Spatial indexes, 99–101
- SPID, 322–323
- SPLIT, 150, 154–155
- Split-brain problem, 456
- Spool operators, work table space, 335
- SQL
  - BatchCompleted, 322, 369
  - BatchStarting, 322, 369
- SQL Browser service, 432
- SQL Mail extended stored procedures (XPs), 260
- SQL Server 2005, 260

- SQL Server 2008
  - Compact edition, 12
  - Developer edition, 12
  - editions, 8–14
  - Enterprise edition, 19, 98, 217–218
  - Enterprise editions, 11
  - Evaluation edition, 12
  - Express edition, 4, 11, 255
  - service failures, diagnosing, 340–347
  - Standard edition, 11
  - word breakers and stemmers, 116
  - Workgroup edition, 11
- SQL Server Agent
  - database failures, diagnosing, 332–335
  - failover clustering, 433–434
  - installing and configuring, 17–18
  - jobs, creating, 234–236
  - log shipping, 490
  - objects, copying, 492–494
  - security, 535
  - service failures, diagnosing, 340–347
  - SMTP (Simple Mail Transfer Protocol), 28
- SQL Server Analysis Services (SSAS), 11, 14, 234–235, 319
- SQL Server Configuration Manager, 19–21, 260, 340
- SQL Server Full Text Search Daemon accounts, 433
- SQL Server Integration Services (SSIS), 8–9
  - Database Mirroring, 466
  - failover, login migration, 471
  - import/export of data, 163, 166
  - packages, copying, 492–494
  - SQL Server 2008 editions, 14
- SQL Server Management Studio (SSMS), 12, 163–166, 395–396
- SQL Server Profiler
  - data columns, selecting, 322–323
  - filters, 323–324
  - overview, 317
  - trace events, specifying, 320–322
  - traces, defining, 317–319
- SQL Server Reporting Services (SSRS), 10, 14
- SQL Server services, 8–11, 431–434
- SQL Trace. *See also* Database Engine
  - Tuning Advisor (DTA)
  - correlating performance and monitoring data, 325
  - data columns, selecting, 322–323
  - defining a trace, 317–319
  - filters, applying, 323–324

- managing traces, 324–325
- overview, 317
- Performance Data Warehouse, 395–396
- trace events, specifying, 320–322
- Sqlagent.out, 333
- SSAS (SQL Server Analysis Services), 11, 14, 234–235, 319
- SSIS (SQL Server Integration Services), 8–9
  - Database Mirroring, 466
  - failover, login migration, 471
  - import/export of data, 163, 166
  - packages, copying, 492–494
  - SQL Server 2008 editions, 14
- SSMS (SQL Server Management Service), 12, 163–166, 395–396
- SSRS (SQL Server Reporting Services), 10, 14
- Stack dumps, 332–333, 352
- Stack, calling, 274–275
- Standard cluster, 412
- Standard SQL Server login, 263–265
- Standard workweek, 236
- Standards, facets, 179–180
- STANDBY, 215
- Standby Mode, 484
- Standby server, 457, 502–503
- STARTED, endpoints, 253
- StartTime, 323, 390
- Startup, 19–21, 332–333, 340–347
- Startup folder, changing, 348
- Static routing, 9
- Static row filter, 515
- Statistics
  - auto options, 48–50, 388
  - disk subsystem, 387, 390–391
  - distribution, 96–97
  - dynamic management views, 388–389
  - queries, 389–390
  - Replication Monitor, 532
  - Showplan Statistics Profile, 321
- Stemmers, 114, 116, 120
- StmtCompleted, 322
- StmtStarting, 322
- Stop word lists, 115, 128, 130–131
- STOP\_ON\_ERROR, 202–203
- STOPAT, 217
- STOPATMARK, 217
- STOPBEFOREMARK, 217
- STOPPED, endpoints, 253
- Storage. *See also* Size considerations; *also* Space Management
  - binary data, 67
  - character data types, 65–66
  - collation sequences, 18
  - database space issues, 334–335
  - date and time data, 66
  - decimal data, 65
  - disk failure, diagnosing, 351
  - files and filegroups, 39–44
  - FILESTREAM, 70–71
  - indexes
    - architecture, 87–91
    - case scenario, data management, 108–110
    - case scenario, full text indexing, 133
    - change tracking, full text, 115–116
    - clustered indexes, 93–95
    - covering indexes, 96
    - designing, 93
    - distribution statistics, 96–97
    - filters, 97
    - full text catalogs, 113–114
    - full text indexes, 114–115, 127–129
    - included columns, 96
    - language, word breakers and stemmers, 116
    - maintenance of, 95–96, 104–105
    - nonclustered indexes, 95
    - online index creation, 98
    - options, 97–98
    - practice, creating indexes, 102, 117–118
    - practice, full text queries, 123–125
    - practice, maintaining, 106
    - practice, managing full text indexes, 129–131
    - querying full text data, 120–123
    - spatial indexes, 99–101
    - XML indexes, 99
  - integrity, maintaining, 54–55
  - numeric data, 64
  - offsite, backups, 208
  - page compression, 72–73
  - partitions
    - case scenario, 157–160
    - creating partition functions, 137–141
    - managing partitions, 150–154
    - partition schemes, creating, 142–144
    - practice, partitioning, 140–141, 144, 148, 154–155
    - tables and indexes, 146–148

## Storage Area Network (SAN)

- sparse files, 224
- tables, 64
- tempdb, 44
- transaction logs, 42–43
- Storage Area Network (SAN), 351, 414, 416
- Stored procedures
  - event group, 321
  - merge replication, 541
  - security, 274–275
  - SQL Server 2008 editions, 13
  - trace events, managing, 324–325
  - transactional replication, validation, 532–533
- Striped backups, 201
- Subscriber, replication
  - agents, 517–518
  - data conflicts, 521–523
  - merge replication, 538–541
  - overview, 515–516
  - practice, creating the subscription, 535–536
  - replication methods, 519–521
  - topology, 516–517
  - transactional replication
    - architecture, 530–532
    - change tracking, 526–528
    - options, 528–530
    - validation, 532–533
- Subscriptions, 10, 181
- Surface Area Configuration Manager, 260
- Surface area, configuring, 259–261
- Suspect pages table, 48
- SWITCH, 151–155
- Switches, command-line, 163–165
- Symmetric keys, 207–208, 276, 294, 297
- Synonym searches, 122–123
- sys schema, 387
- sys.database\_mirroring\_witnesses, 453
- sys.dm\_db\_index\_operational\_stats, 388–389
- sys.dm\_db\_index\_physical\_stats, 388–389
- sys.dm\_db\_index\_usage\_stats, 388–389
- sys.dm\_db\_missing\_index\_details, 388–389
- sys.dm\_db\_missing\_index\_group\_stats, 388–389
- sys.dm\_db\_missing\_index\_groups, 388–389
- sys.dm\_exec\_query\_stats, 396
- sys.dm\_exec\_requests, 396
- sys.dm\_exec\_sessions, 396
- sys.dm\_io\_pending\_requests, 391
- sys.dm\_io\_virtual\_file\_stats, 390–391
- sys.dm\_os\_wait\_stats, 391–392

- sysadmin role, 50, 266, 271
- System databases, device activation errors, 347
- System Event log, 332–335
- System Monitor
  - counter logs, 310–312
  - counters, 242
  - overview, 309–310
  - performance counters, 312–313
  - practice, creating counter log, 313–315
- System objects, copying, 466

## T

- T parameter, 346–347
- Tables. *See also* Columns; *also* Indexes;  
*also* Pages; *also* Rows
  - Bulk Copy Program (BCP), 164
  - case scenarios, data management, 81–82
  - collation sequences, 18
  - column properties, 69–72
  - compression, row and page, 72–73, 77
  - computed columns, 72
  - constraints, implementing, 77–79
    - creating, 63, 73–74
  - data types, 64–69
  - locking, 355–356
  - locking online indexes, 98
  - partitioning, 146–147
  - practice
    - constraints, 79–80
    - creating, 75
    - exporting data, 167–170
    - partitioning a table, 148
  - schemas, 63
- TAKE OWNERSHIP, 114
- TCP endpoints, 252–257
- TCP transports, 252
- TCP/IP networking, 5, 259
- TDE (transparent data encryption), 294–296,  
299–300
- tempdb, 44
  - device activation errors, 345
  - index options, 98
  - online indexes, 98
  - service failures, diagnosing, 345
  - space issues, 335
  - temporary tables, 74

- Templates, Profiler traces, 318
- Temporary tables, 74
- Term extraction, 14
- Tessellation, 99–101
- Test environments, 204–205
- Text mining, 9
- THESAURUS, 122–123, 127–130
- Time, 236
  - data, 66
  - elapsed, 390
  - estimated completion time, 390
  - trace events, 323
- TIME, 66
- Timeout, logins, 518
- Timeout, queries, 518
- Timestamps, 50, 129
- TINYINT, 64
- Tokens, thesaurus files, 127–128
- Tokens, tracer, 532
- Tools, client, 319
- Torn pages, 54–55. *See also* Pages
- TORN\_PAGE\_VERIFY, 54
- Tracer token, 532
- Traces
  - database failures, diagnosing, 332–333
  - defining, 317–319
  - events, selecting, 320–323
  - filters, 323–324
  - managing, 324–325
- Transaction isolation levels, 356
- Transaction log. *See also* Transactional replication
  - access control, 50
  - backups, 203–204, 210, 213–214
  - Database Mirroring, 453, 455–457
  - Database Snapshots, 225
  - database space issues, 334–335
  - files and filegroups, 39–44
  - log chain, 203
  - overview, 42–43, 212–214
  - practice, backups, 210
  - recovery options, 46–48
  - restoring, 212–214, 216–217, 221
- Transactional replication
  - architecture, 530–532
  - change tracking, 526–528
  - minimally logged transactions, 522–523
  - monitoring, 532
  - options, 528–530
    - overview, 520, 526
    - practice, implementing, 534–536
    - replication agents, 518
    - SQL Server 2008 editions, 13
    - validation, 532–533
- Transactions event group, 321
- Transact-SQL (T-SQL) commands, 165–166, 234–235
- Transformation, import/export of data, 163, 166
- Transient operating states, 453
- Transmission Control Protocol (TCP), 252–257
- Transmission Control Protocol/Internet Protocol (TCP/IP), 259
- Transparent data encryption (TDE), 294–296, 299–300
- Transports, endpoints, 252
- Triggers
  - Bulk Copy Program (BCP), 165
  - DDL triggers, 285–286
  - replication, 522–523, 529–530, 538–540
- Troubleshooting
  - blocking and deadlocking issues, 355–359
  - case scenario, designing automation strategy, 364–365
  - correlating performance and monitoring data, 325, 329–330
  - database failures, diagnosing, 332–335
  - dynamic management views (DMVs), 387–392
  - hardware failures, diagnosing, 351–352
  - performance counters, 312–313
  - practice
    - creating failure alerts, 336–338
    - creating traces, 326–330
    - deadlocks, troubleshooting, 359–361
    - startup errors, troubleshooting, 348–349
  - Profiler traces, 317
  - remote desktop startup failures, 341–343
  - Replication Monitor, 532
  - service failures, diagnosing, 340–347
  - trace events
    - applying filters, 323–324
    - defining, 317–319
    - managing, 324–325
    - selecting data columns, 322–323
    - specifying, 320–322
- TRUNCATE TABLE, 522–523
- TSQL, 252, 318
- T-SQL (Transact-SQL) commands, 165–166, 234–235
- TSQL event group, 321
- T-SQL Query, 395

TSQL\_Duration, 318  
 TSQL\_Grouped, 318  
 TSQL\_Locks, 318  
 TSQL\_Replay, 318  
 TSQL\_SPs, 318  
 Tuning, 318, 371–373  
 TYPE COLUMN, 115

## U

UMS (User Mode Scheduler), 391  
 UNDO, 213  
 Unicode data, 65–66  
 UNION, 335  
 Unique constraints, 78, 95, 519  
 UNIQUEIDENTIFIER, 70  
 UPDATE, 74, 96, 129, 225  
 Update locks, 355–356  
 Updating, replication, 521, 528–530  
 Upgrades, log shipping, 484  
 Uploading files, 9. *See also* SQL Server Integration Services (SSIS)  
 Usage statistics, 388–389  
 User Configurable event group, 321  
 User Mode Scheduler (UMS), 391  
 Users  
   loginless, 266–267  
   logins, 266  
   mobile/remote, 259–260, 341–343, 538–541  
   practice creating, 268–269  
   remapping, 493  
 Utilities, startup failures, 340–347

## V

Validation, 209, 532–533, 541  
 VALUE, 99  
 VARBINARY, 67, 73, 114  
 VARBINARY (MAX), 71, 74, 114  
 VARCHAR, 114  
 VARCHAR (MAX), 74  
 VARCHAR (n), 66, 73  
 Version store alerts, 336  
 Version upgrades, 484  
 Versioning, 50, 335  
 VIEW ANY DATABASE, 274

VIEW ANY DEFINITION, 274  
 Views  
   Bulk Copy Program (BCP), 164  
   CREATE SCHEMA, 63  
   Database Engine Tuning Analysis (DTA), 372  
 Views, indexed  
   Database Engine Tuning Analysis (DTA), 372  
   partitions  
     case scenario, 157–160  
     creating partition functions, 137–141  
     managing partitions, 150–154  
     partition schemes, creating, 142–144  
     practice, partitioning, 140–141, 144, 148, 154–155  
     tables and indexes, 146–148  
   SQL Server 2008 editions, 13  
 Viruses, 259, 414  
 Visual Studio Tools for Applications, 9

## W

Wait time, 391  
 Wait type, 391  
 Waiting queues, 391–392  
 Warnings, cluster analysis, 427  
 Warnings, event group, 320  
 Watermark, replication, 526–527  
 Web services, 9. *See also* SQL Server Integration Services (SSIS)  
 Web sites, for more information  
   CodePlex, policies, 183  
   event groups, 286  
   EVENTDATA ( ) schemas, 286  
   installation requirements, 3–4  
   wait types, 392  
   Windows clustering information, 410  
 WHERE, 78, 180  
 Wide tables, 74  
 Wildcards, 79, 122  
 Windows Application Event log, 48, 242  
   database space issues, 334–335  
   device activation errors, 345–347  
   disk failure, diagnosing, 351  
   service failures, diagnosing, 342, 345  
 Windows Datacenter, 413  
 Windows Event log, 332–335  
 Windows group logins, 263–265

Windows login, 263–265

Windows Management Instrumentation (WMI),  
9, 242–244. *See also* SQL Server Integration  
Services (SSIS)

Windows process ID, 332–333

Windows Server 2003, 4, 410  
Small Business Server, 4

Windows Server 2008, 4, 410  
Server Core edition, 4

Windows Service Control applet, 20–21

Windows Services, 340–347

Windows System Event log, 351

Windows Vista Home Basic, 4

Windows XP  
Home Edition, 4  
Home Reduced Media Edition, 4  
Media Center 2002, 4  
Professional, 4  
Professional Embedded Edition, 4  
Professional Reduced Media Edition, 4  
Tablet Edition SP2, 4

Windows, supported versions, 4

WITH FORMAT, 202

WITH ONLINE = OFF, 98

WITH RECOVERY, 502–503

WITNESS, 255

Witness server, database mirroring, 453

WMI (Windows Management Instrumentation),  
9, 242–244. *See also* SQL Server Integration  
Services (SSIS)

Word breakers, 114, 116, 127–128

Word filters, 128

Word forms, searching, 121–123

Word proximity, searching, 121–123

Work tables, 44, 97–98, 335

Workflows, operational, 9

Workload  
Database Engine Tuning Advisor (DTA),  
370–373  
practice assigning, 380  
practice, analyzing query workload, 374  
Resource Governor and, 376–378

Workweek, standard, 236

Write operations  
backups, 199  
Copy-On-Write, 224–225  
database space and, 334–335  
incomplete, 206  
indexes, 95, 97–98  
integrity, maintaining, 54–55  
query statistics, 389–390

## X

X.509 standard, 294

XML (Extensible Markup Language)  
data types, 67  
indexes, full text, 114  
Showplan events, 321  
SQL Server 2008 editions, 13  
thesaurus files, 127–128

XML indexes, 99

xp\_cmdshell, 260

