

让Jim Gray和David Campbell拍案叫绝的圣经级著作

Microsoft SQL Server 2008 Internals

深入解析

SQL Server 2008

(英文版)

[美] Kalen Delaney  
Paul S. Randal  
Kimberly L. Tripp  
Conor Cunningham  
Adam Machanic  
Ben Nevarez 著

- 微软SQL Server开发团队必读之作
- 六位SQL Server专家巨献
- 深入剖析SQL Server 2008技术内幕



人民邮电出版社  
POSTS & TELECOM PRESS



TURING

图灵程序设计丛书

微软技术系列

Microsoft SQL Server 2008 Internals

深入解析

SQL Server 2008

(英文版)

江苏工业学院图书馆

藏书章

Kalen Delaney

Paul S. Randal

Kimberly L. Tripp

[美]

Conor Cunningham

Adam Machanic

Ben Nevarez

人民邮电出版社

北京

## 图书在版编目 (CIP) 数据

深入解析 SQL Server 2008: 英文 / (美) 德莱尼  
(Delaney, K.) 等著. —北京: 人民邮电出版社, 2009.9  
(图灵程序设计丛书)  
书名原文: Microsoft SQL Server 2008 Internals  
ISBN 978-7-115-21143-9

I. 深… II. 德… III. 关系数据库-数据库管理系统,  
SQL Server 2008-英文 IV. TP311.138

中国版本图书馆CIP数据核字 (2009) 第123112号

## 内 容 提 要

本书是讲述 SQL Server 关系数据库引擎内部机理和架构的权威指南。书中详细阐述了 SQL Server 处理查询、管理数据的相关内容, 包括 SQL Server 架构和配置、跟踪/扩展事件、日志和恢复、索引、表格、查询优化、事务/并发以及 DBCC。

本书适合中高级数据库开发人员阅读。

图灵程序设计丛书

### 深入解析SQL Server 2008 (英文版)

◆ 著 [美] Kalen Delaney Paul S. Randal  
Kimberly L. Tripp Conor Cunningham  
Adam Machanic Ben Nevarez

责任编辑 傅志红

◆ 人民邮电出版社出版发行 北京市崇文区夕照寺街14号

邮编 100061 电子函件 315@ptpress.com.cn

网址 <http://www.ptpress.com.cn>

北京艺辉印刷有限公司印刷

◆ 开本: 800×1000 1/16

印张: 48.5

字数: 931千字

2009年9月第1版

印数: 1-2 000册

2009年9月北京第1次印刷

著作权合同登记号 图字: 01-2009-3799号

ISBN 978-7-115-21143-9/TP

定价: 99.00元

读者服务热线: (010)51095186 印装质量热线: (010)67129223

反盗版热线: (010)67171154

# 版 权 声 明

© 2009 by Microsoft Corporation. All rights reserved. Original edition, entitled *Microsoft SQL Server 2008 Internals* by Kalen Delaney, Paul S. Randal, Kimberly L. Tripp, Conor Cunningham, Adam Machanic, and Ben Nevarez, ISBN 978-0-7356-2624-9, published by Microsoft Press in 2009.

This reprint edition is published with the permission of the Syndicate of the Microsoft Press.

Copyright © 2009 by Kalen Delaney.

THIS EDITION IS LICENSED FOR DISTRIBUTION AND SALE IN THE PEOPLE'S REPUBLIC OF CHINA ONLY, EXCLUDING HONG KONG, MACAO AND TAIWAN, AND MAY NOT BE DISTRIBUTED AND SOLD ELSEWHERE.

本书原版由微软出版社出版。

本书英文影印版由微软出版社授权人民邮电出版社独家出版。未经出版者书面许可，不得以任何方式复制或抄袭本书内容。

此版本仅限在中华人民共和国（香港、澳门特别行政区和台湾地区除外）境内销售发行。

版权所有，侵权必究。

# Foreword

The developers who create products such as Microsoft SQL Server typically become experts in one area of the technology, such as access methods or query execution. They live and experience the product inside out and often know their component so deeply they acquire a “curse of knowledge”: they possess so much detail about their particular domain, they find it difficult to describe their work in a way that helps customers get the most out of the product.

Technical writers who create product-focused books, on the other hand, experience a product outside in. Most of these authors acquire a broad, but somewhat shallow, surface knowledge of the products they write about and produce valuable books, usually filled with many screenshots, which help new and intermediate users quickly learn how to get things done with the product.

When the curse of knowledge meets surface knowledge, it leaves a gap where many of the great capabilities created by product developers don’t get communicated in a way that allows customers, particularly intermediate to advanced users, to use a product to its full potential. This is where *Microsoft SQL Server 2008 Internals* comes in. This book, like those in the earlier “Inside SQL Server” series, is the definitive reference for how SQL Server really works. Kalen Delaney has been working with the SQL Server product team for over a decade, spending countless hours with developers breaking through the curse of knowledge and then capturing the result in an incredibly clear form that allows intermediate to advanced users to wring the most from the capabilities of SQL Server. In *Microsoft SQL Server 2008 Internals*, Kalen is joined by four SQL Server experts who also share the gift of breaking the curse. Conor Cunningham and Paul Randal have years of experience as SQL Server product developers, and each of them is both a deep technical expert and a gifted communicator. Kimberly Tripp and Adam Machanic both combine a passion to really understand how things work and to then effectively share it with others. Kimberly and Adam are both standing-room-only speakers at SQL Server events. This team has captured and incorporated the details of key architectural changes for SQL Server 2008, resulting in a new, comprehensive internals reference for SQL Server.

There is a litmus test you can use to determine if a technical product title deserves a “definitive reference” classification. It’s a relatively easy test but a hard one for everybody to conduct. The test, quite simply, is to look at how many of the developers who created the product in question have a copy of the book on their shelves—and reference it. I can assure you that each version of *Inside Microsoft SQL Server* that Kalen has produced has met this test. *Microsoft SQL Server 2008 Internals* will, too.

*Dave Campbell*

*Technical Fellow  
Microsoft SQL Server*

# Introduction

The book you are now holding is the evolutionary successor to the *Inside SQL Server* series, which included *Inside SQL Server 6.5*, *Inside SQL Server 7*, *Inside SQL Server 2000*, and *Inside SQL Server 2005* (in four volumes). The *Inside* series was becoming too unfocused, and the name “Inside” had been usurped by other authors and even other publishers. I needed a title that was much more indicative of what this book is really about.

*SQL Server 2008 Internals* tells you how SQL Server, Microsoft’s flagship relational database product, works. Along with that, I explain how you can use the knowledge of how it works to help you get better performance from the product, but that is a side effect, not the goal. There are dozens of other books on the market that describe tuning and best practices for SQL Server. This one helps you understand why certain tuning practices work the way they do, and it helps you determine your own best practices as you continue to work with SQL Server as a developer, data architect, or DBA.

## Who This Book Is For

This book is intended to be read by anyone who wants a deeper understanding of what SQL Server does behind the scenes. The focus of this book is on the core SQL Server engine—in particular, the query processor and the storage engine. I expect that you have some experience with both the SQL Server engine and with the T-SQL language. You don’t have to be an expert in either, but it helps if you aspire to become an expert and would like to find out all you can about what SQL Server is actually doing when you submit a query for execution.

This series doesn’t discuss client programming interfaces, heterogeneous queries, business intelligence, or replication. In fact, most of the high-availability features are not covered, but a few, such as mirroring, are mentioned at a high level when we discuss database property settings. I don’t drill into the details of some internal operations, such as security, because that’s such a big topic it deserves a whole volume of its own.

My hope is that you’ll look at the cup as half full instead of half empty and appreciate this book for what it does include. As for the topics that aren’t included, I hope you’ll find the information you need in other sources.

## What This Book Is About

*SQL Server Internals* provides detailed information on the way that SQL Server processes your queries and manages your data. It starts with an overview of the architecture of the SQL Server relational database system and then continues looking at aspects of query processing and data storage in 10 additional chapters, as follows:

- Chapter 1 SQL Server 2008 Architecture and Configuration
- Chapter 2 Change Tracking, Tracing, and Extended Events
- Chapter 3 Databases and Database Files
- Chapter 4 Logging and Recovery
- Chapter 5 Tables
- Chapter 6 Indexes: Internals and Management
- Chapter 7 Special Storage
- Chapter 8 The Query Optimizer
- Chapter 9 Plan Caching and Recompilation
- Chapter 10 Transactions and Concurrency
- Chapter 11 DBCC Internals

A twelfth chapter covering the details of reading query plans is available in the companion content (which is described in the next section). This chapter, called “Query Execution,” was part of my previous book, *Inside SQL Server 2005: Query Tuning and Optimization*. Because 99 percent of the chapter is still valid for SQL Server 2008, we have included it “as is” for your additional reference.

## Companion Content

This book features a companion Web site that makes available to you all the code used in the book, organized by chapter. The companion content also includes an extra chapter from my previous book, as well as the “History of SQL Server” chapter from my book *SQL Server 2000*. The site also provides extra scripts and tools to enhance your experience and understanding of SQL Server internals. As errors are found and reported, they will also be posted online. You can access this content from the companion site at this address: <http://www.SQLServerInternals.com/companion>.

## System Requirements

To use the code samples, you’ll need Internet access and a system capable of running SQL Server 2008 Enterprise or Developer edition. To get system requirements for SQL Server 2008 and to obtain a trial version, go to <http://www.microsoft.com/downloads>.

## Support for This Book

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

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

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

## Questions and Comments

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

[mspinput@microsoft.com](mailto:mspinput@microsoft.com)

Or via postal mail to

Microsoft Press  
Attn: *Microsoft SQL Server 2008 Internals* Editor  
One Microsoft Way  
Redmond, WA 98052-6399

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

## Acknowledgments

As always, a work like this is not an individual effort, and for this current volume, it is truer than ever. I was honored to have four other SQL Server experts join me in writing *SQL Server 2008 Internals*, and I truly could not have written this book alone. I am grateful to Adam Machanic, Paul Randal, Conor Cunningham, and Kimberly Tripp for helping to make this book a reality. In addition to my brilliant co-authors, this book could never have seen the light of day with help and encouragement from many other people.

First on my list is you, the readers. Thank you to all of you for reading what I have written. Thank you to those who have taken the time to write to me about what you thought of the book and what else you want to learn about SQL Server. I wish I could answer every question in detail. I appreciate all your input, even when I'm unable to send you a complete reply. One particular reader of one of my previous books, *Inside SQL Server 2005: The Storage Engine*, deserves particular thanks. I came to know Ben Nevarez as a very astute reader who found some uncaught errors and subtle inconsistencies and politely and succinctly reported them to me through my Web site. After a few dozen e-mails, I started to look forward to Ben's e-mails and was delighted when I finally got the chance to meet him. Ben is now my most valued technical reviewer, and I am deeply indebted to him for his extremely careful reading of every one of the chapters.



As usual, the SQL Server team at Microsoft has been awesome. Although Lubor Kollar and Sunil Agarwal were not directly involved in much of the research for this book, I always knew they were there in spirit, and both of them always had an encouraging word whenever I saw them.

Boris Baryshnikov, Kevin Farlee, Marcel van der Holst, Peter Byrne, Sangeetha Shekar, Robin Dhamankar, Artem Oks, Srini Acharya, and Ryan Stonecipher met with me and responded to my (sometimes seemingly endless) e-mails. Jerome Halmans, Joanna Omel, Nikunj Koolar, Tres London, Mike Purtell, Lin Chan, and Dipti Sangani also offered valuable technical insights and information when responding to my e-mails. I hope they all know how much I appreciated every piece of information I received.

I am also indebted to Bob Ward, Bob Dorr, and Keith Elmore of the SQL Server Product Support team, not just for answering occasional questions but for making so much information about SQL Server available through white papers, conference presentations, and Knowledge Base articles. I am grateful to Alan Brewer and Gail Erickson for the great job they and their User Education team did putting together the SQL Server documentation in *SQL Server Books Online*.

And, of course, Buck Woody deserves my gratitude many times over. First from his job in the User Education group, then as a member of the SQL Server development team, he was always there when I had an unanswered question. His presentations and blog posts are always educational as well as entertaining, and his generosity and unflagging good spirits are a true inspiration.

I would also like to thank Leona Lowry and Cheryl Walter for finding me office space in the same building as most of the SQL Server team. The welcome they gave me was much appreciated.

I would like to extend my heartfelt thanks to all of the SQL Server MVPs, but most especially Erland Sommarskog. Erland wrote the section in Chapter 5 on collations just because he thought it was needed, and that someone who has to deal with only the 26 letters of the English alphabet could never do it justice. Also deserving of special mention are Tibor Karaszi and Roy Harvey, for all the personal support and encouragement they gave me. Other MVPs who inspired me during the writing of this volume are Tony Rogerson, John Paul Cook, Steve Kass, Paul Nielsen, Hugo Kornelis, Tom Moreau, and Linchi Shea. Being a part of the SQL Server MVP team continues to be one of the greatest honors and privileges of my professional life.

I am deeply indebted to my students in my “SQL Server Internals” classes, not only for their enthusiasm for the SQL Server product and for what I have to teach and share with them, but for all they have to share with me. Much of what I have learned has been inspired by questions from my curious students. Some of my students, such as Cindy Gross and Lara Rubbelke, have become friends (in addition to becoming Microsoft employees) and continue to provide ongoing inspiration.

Most important of all, my family continues to provide the rock-solid foundation I need to do the work that I do. My husband, Dan, continues to be the guiding light of my life after 24 years of marriage. My daughter, Melissa, and my three sons, Brendan, Rickey, and Connor,

are now for the most part all grown, and are all generous, loving, and compassionate people. I feel truly blessed to have them in my life.

*Kalen Delaney*

## Paul Randal

I've been itching to write a complete description of what DBCC CHECKDB does for many years now—not least to get it all out of my head and make room for something else! When Kalen asked me to write the “Consistency Checking” chapter for this book, I jumped at the chance, and for that my sincere thanks go to Kalen. I'd like to give special thanks to two people from Microsoft, among the many great folks I worked with there (and in many cases still do). The first is Ryan Stonecipher, who I hired away from being an Escalation Engineer in SQL Product Support in late 2003 to work with me on DBCC, and who was suddenly thrust into complete ownership of 100,000+ lines of DBCC code when I became the team manager two months later. I couldn't have asked for more capable hands to take over my precious DBCC.... The second is Bob Ward, who heads up the SQL Product Support team and has been a great friend since my early days at Microsoft. We must have collaborated on hundreds of cases of corruption over the years, and I've yet to meet someone with more drive for solving customer problems and improving Microsoft SQL Server. Thanks must also go to Steve Lindell, the author of the original online consistency checking code for SQL Server 2000, who spent many hours patiently explaining how it worked in 1999. Finally, I'd like to thank my wife, Kimberly, who is, along with Katelyn and Kiera, the other passions in my life apart from SQL Server.

## Kimberly Tripp

First, I want to thank my good friend Kalen, for inviting me to participate in this title. After working together in various capacities—even having formed a company together back in 1996—it's great to finally have our ideas and content together in a book as deep and technical as this. In terms of performance tuning, indexes are critical; there's no better way to improve a system than by creating the *right* indexes. However, knowing what's right takes multiple components, some of which is only known after experience, after testing, and after seeing something in action. For this, I want to thank many of you—readers, students, conference attendees, customers—those of you who have asked the questions, shown me interesting scenarios, and stayed late to “play” and/or just figure it out. It's the deep desire to know why something is working the way that it is that keeps this product interesting to me and has always made me want to dive deeper and deeper into understanding what's really going on. For that, I thank the SQL team in general—the folks that I've met and worked with over the years have been inspiring, intelligent, and insightful. Specifically, I want to thank a few folks on the SQL team who have patiently, quickly, and thoroughly responded to questions about what's really going on and often, why: Conor Cunningham,

Cesar Galindo-Legaria, and from my early days with SQL Server, Dave Campbell, Nigel Ellis, and Rande Blackman. Gert E. R. Drapers requires special mention due to the many hours spent together over the years where we talked, argued, and figured it out. And, to Paul, my best friend and husband, who before that was also a good source of SQL information. We just don't talk about it anymore ... at home. OK, maybe a little.

## Conor Cunningham

I'd like to thank Bob Beauchemin and Milind Joshi for their efforts to review my chapter, "The Query Optimizer," in this book for technical correctness. I'd also like to thank Kimberly Tripp and Paul Randal for their encouragement and support while I wrote this chapter. Finally, I'd like to thank all the members of the SQL Server Query Processor team who answered many technical questions for me.

## Adam Machanic

I would like to, first and foremost, extend my thanks to Kalen Delaney for leading the effort of this book from conception through reality. Kalen did a great job of keeping us focused and on task, as well as helping to find those hidden nuggets of information that make a book like this one great. A few Microsoft SQL Server team members dedicated their time to helping review my work: Jerome Halmans and Fabricio Voznika from the Extended Events team, and Mark Scurrell from the Change Tracking team. I would like to thank each of you for keeping me honest, answering my questions, and improving the quality of my chapter. Finally, I would like to thank Kate and Aura, my wife and daughter, who always understand when I disappear into the office for a day or two around deadline time.

# Table of Contents

|          |   |          |
|----------|---|----------|
| <b>1</b> | <b>SQL Server 2008 Architecture and Configuration</b> | <b>1</b> |
|          | SQL Server Editions                                   | 1        |
|          | SQL Server Metadata                                   | 2        |
|          | Compatibility Views                                   | 3        |
|          | Catalog Views   | 4        |
|          | Other Metadata  | 6        |
|          | Components of the SQL Server Engine                   | 8        |
|          | Observing Engine Behavior                             | 9        |
|          | Protocols   | 11       |
|          | The Relational Engine                                 | 12       |
|          | The Storage Engine                                    | 14       |
|          | The SQLOS   | 18       |
|          | NUMA Architecture                                     | 19       |
|          | The Scheduler   | 20       |
|          | SQL Server Workers                                    | 21       |
|          | Binding Schedulers to CPUs                            | 24       |
|          | The Dedicated Administrator Connection (DAC)          | 27       |
|          | Memory  | 29       |
|          | The Buffer Pool and the Data Cache                    | 29       |
|          | Access to In-Memory Data Pages                        | 30       |
|          | Managing Pages in the Data Cache                      | 30       |
|          | The Free Buffer List and the Lazywriter               | 31       |
|          | Checkpoints   | 32       |
|          | Managing Memory in Other Caches                       | 34       |
|          | Sizing Memory   | 35       |
|          | Sizing the Buffer Pool                                | 36       |

|  |            |
|--|------------|
| SQL Server Resource Governor .....                           | 42         |
| Resource Governor Overview .....                             | 42         |
| Resource Governor Controls .....                             | 51         |
| Resource Governor Metadata .....                             | 52         |
| SQL Server 2008 Configuration .....                          | 54         |
| Using SQL Server Configuration Manager .....                 | 54         |
| Configuring Network Protocols .....                          | 54         |
| Default Network Configuration .....                          | 55         |
| Managing Services .....                                      | 55         |
| SQL Server System Configuration .....                        | 57         |
| Operating System Configuration .....                         | 57         |
| Trace Flags .....  | 60         |
| SQL Server Configuration Settings .....                      | 60         |
| The Default Trace .....                                      | 71         |
| Final Words .....  | 73         |
| <b>2 Change Tracking, Tracing, and Extended Events .....</b> | <b>75</b>  |
| The Basics: Triggers and Event Notifications .....           | 75         |
| Run-Time Trigger Behavior .....                              | 76         |
| Change Tracking .....  | 76         |
| Change Tracking Configuration .....                          | 77         |
| Change Tracking Run-Time Behavior .....                      | 82         |
| Tracing and Profiling .....                                  | 86         |
| SQL Trace Architecture and Terminology .....                 | 86         |
| Security and Permissions .....                               | 88         |
| Getting Started: Profiler .....                              | 89         |
| Server-Side Tracing and Collection .....                     | 97         |
| Extended Events .....  | 108        |
| Components of the XE Infrastructure .....                    | 108        |
| Event Sessions .....   | 118        |
| Extended Events DDL and Querying .....                       | 121        |
| Summary .....  | 124        |
| <b>3 Databases and Database Files .....</b>                  | <b>125</b> |
| System Databases .....                                       | 126        |
| <i>master</i> .....  | 126        |
| <i>model</i> .....   | 126        |
| <i>tempdb</i> .....  | 126        |
| The Resource Database .....                                  | 127        |
| <i>msdb</i> .....  | 128        |



|   |     |
|---|-----|
| Sample Databases .....                    | 128 |
| <i>AdventureWorks</i> .....               | 128 |
| <i>pubs</i> .....                         | 129 |
| <i>Northwind</i> .....                    | 129 |
| Database Files .....                      | 130 |
| Creating a Database .....                 | 132 |
| A <i>CREATE DATABASE</i> Example .....    | 134 |
| Expanding or Shrinking a Database .....   | 135 |
| Automatic File Expansion .....            | 135 |
| Manual File Expansion .....               | 136 |
| Fast File Initialization .....            | 136 |
| Automatic Shrinkage .....                 | 136 |
| Manual Shrinkage .....                    | 137 |
| Using Database Filegroups .....           | 138 |
| The Default Filegroup .....               | 139 |
| A <i>FILEGROUP CREATION</i> Example ..... | 140 |
| Filestream Filegroups .....               | 141 |
| Altering a Database .....                 | 142 |
| <i>ALTER DATABASE</i> Examples .....      | 143 |
| Databases Under the Hood .....            | 144 |
| Space Allocation .....                    | 145 |
| Setting Database Options .....            | 148 |
| State Options .....                       | 151 |
| Cursor Options .....                      | 155 |
| Auto Options .....                        | 155 |
| SQL Options .....                         | 156 |
| Database Recovery Options .....           | 158 |
| Other Database Options .....              | 159 |
| Database Snapshots .....                  | 159 |
| Creating a Database Snapshot .....        | 160 |
| Space Used by Database Snapshots .....    | 162 |
| Managing Your Snapshots .....             | 164 |
| The <i>tempdb</i> Database .....          | 164 |
| Objects in <i>tempdb</i> .....            | 165 |
| Optimizations in <i>tempdb</i> .....      | 166 |
| Best Practices .....                      | 168 |
| <i>tempdb</i> Space Monitoring .....      | 169 |
| Database Security .....                   | 170 |
| Database Access .....                     | 170 |
| Managing Database Security .....          | 172 |

|   |            |
|---|------------|
| Databases vs. Schemas .....                     | 173        |
| Principals and Schemas .....                    | 173        |
| Default Schemas .....                           | 174        |
| Moving or Copying a Database .....              | 175        |
| Detaching and Reattaching a Database .....      | 175        |
| Backing Up and Restoring a Database .....       | 177        |
| Moving System Databases .....                   | 177        |
| Moving the <i>master</i> Database .....         | 179        |
| Compatibility Levels .....                      | 179        |
| Summary .....                                   | 180        |
| <b>4 Logging and Recovery .....</b>             | <b>181</b> |
| Transaction Log Basics .....                    | 181        |
| Phases of Recovery .....                        | 184        |
| Reading the Log .....                           | 186        |
| Changes in Log Size .....                       | 187        |
| Virtual Log Files .....                         | 187        |
| Observing Virtual Log Files .....               | 188        |
| Automatic Truncation of Virtual Log Files ..... | 192        |
| Maintaining a Recoverable Log .....             | 193        |
| Automatic Shrinking of the Log .....            | 196        |
| Log File Size .....                             | 196        |
| Backing Up and Restoring a Database .....       | 197        |
| Types of Backups .....                          | 197        |
| Recovery Models .....                           | 198        |
| Choosing a Backup Type .....                    | 203        |
| Restoring a Database .....                      | 203        |
| Summary .....                                   | 209        |
| <b>5 Tables .....</b>                           | <b>211</b> |
| Creating Tables .....                           | 211        |
| Naming Tables and Columns .....                 | 212        |
| Reserved Keywords .....                         | 213        |
| Delimited Identifiers .....                     | 214        |
| Naming Conventions .....                        | 215        |
| Data Types .....                                | 215        |
| Much Ado About NULL .....                       | 241        |

|  |            |
|--|------------|
| User-Defined Data Types .....  | 244        |
| <i>IDENTITY</i> Property .....   | 245        |
| Internal Storage .....   | 249        |
| The <i>sys.indexes</i> Catalog View .....  | 250        |
| Data Storage Metadata .....  | 251        |
| Data Pages .....   | 254        |
| Examining Data Pages .....   | 256        |
| The Structure of Data Rows .....   | 260        |
| Finding a Physical Page .....  | 262        |
| Storage of Fixed-Length Rows .....   | 265        |
| Storage of Variable-Length Rows .....  | 267        |
| Storage of Date and Time Data .....  | 272        |
| Storage of <i>sql_variant</i> Data .....   | 275        |
| Constraints .....  | 279        |
| Constraint Names and Catalog View Information .....                              | 280        |
| Constraint Failures in Transactions and Multiple-Row<br>Data Modifications ..... | 281        |
| Altering a Table .....   | 282        |
| Changing a Data Type .....   | 283        |
| Adding a New Column .....  | 284        |
| Adding, Dropping, Disabling, or Enabling a Constraint .....                      | 284        |
| Dropping a Column .....  | 285        |
| Enabling or Disabling a Trigger .....  | 286        |
| Internals of Altering Tables .....   | 286        |
| Heap Modification Internals .....  | 289        |
| Allocation Structures .....  | 289        |
| Inserting Rows .....   | 290        |
| Deleting Rows .....  | 291        |
| Updating Rows .....  | 294        |
| Summary .....  | 297        |
| <b>6 Indexes: Internals and Management .....</b>                                 | <b>299</b> |
| Overview .....   | 299        |
| SQL Server Index B-trees .....   | 300        |
| Tools for Analyzing Indexes .....  | 304        |
| Using the <i>dm_db_index_physical_stats</i> DMV .....                            | 304        |
| Using <i>DBCC IND</i> .....  | 308        |

|   |     |
|---|-----|
| Understanding Index Structures .....                | 310 |
| The Dependency on the Clustering Key .....          | 311 |
| Nonclustered Indexes .....                          | 314 |
| Constraints and Indexes .....                       | 315 |
| Index Creation Options .....                        | 316 |
| IGNORE_DUP_KEY .....                                | 316 |
| STATISTICS_NORECOMPUTE .....                        | 317 |
| MAXDOP .....  | 317 |
| Index Placement .....                               | 317 |
| Constraints and Indexes .....                       | 318 |
| Physical Index Structures .....                     | 318 |
| Index Row Formats .....                             | 318 |
| Clustered Index Structures .....                    | 319 |
| The Non-Leaf Level(s) of a Clustered Index .....    | 320 |
| Analyzing a Clustered Index Structure .....         | 321 |
| Nonclustered Index Structures .....                 | 326 |
| Special Index Structures .....                      | 337 |
| Indexes on Computed Columns and Indexed Views ..... | 337 |
| Full-Text Indexes .....                             | 345 |
| Spatial Indexes .....                               | 346 |
| XML Indexes .....                                   | 346 |
| Data Modification Internals .....                   | 347 |
| Inserting Rows .....                                | 347 |
| Splitting Pages .....                               | 348 |
| Deleting Rows .....                                 | 352 |
| Updating Rows .....                                 | 358 |
| Table-Level vs. Index-Level Data Modification ..... | 362 |
| Logging .....                                       | 363 |
| Locking .....                                       | 363 |
| Fragmentation .....                                 | 363 |
| Managing Index Structures .....                     | 364 |
| Dropping Indexes .....                              | 365 |
| ALTER INDEX .....                                   | 365 |
| Detecting Fragmentation .....                       | 368 |
| Removing Fragmentation .....                        | 369 |
| Rebuilding an Index .....                           | 371 |
| Summary .....                                       | 374 |