第一图书网, tushu007.com

<<深入解析SQL Server 2008>>

图书基本信息

书名：<<深入解析SQL Server 2008>>

13位ISBN编号：9787115211439

10位ISBN编号：7115211434

出版时间：2009..9

出版时间：人民邮电出版社

作者：Kalen Delaney,Paul S. Randal,Kimberly L. Tripp,Conor Cunningham,Adam Machanic,Ben Nevarez

页数：754

版权说明：本站所提供下载的PDF图书仅提供预览和简介，请支持正版图书。

更多资源请访问：http://www.tushu007.com

前言

The developers who create products such as Microsoft SQL Server typically become expertsin one area of the technology， such as access methods or query execution. They live andexperience 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 itdifficult 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 aproduct outside in. Most of these authors acquire a broad， but somewhat shallow， surfaceknowledge of the products they write about and produce valuable books， usually filled withmany screenshots， which help new and intermediate users quickly learn how to get thingsdone with the product. When the curse of knowledge meets surface knowledge， it leaves a gap where many ofthe great capabilities created by product developers dont get communicated in a waythat allows customers， particularly intermediate to advanced users， to use a product toits 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 SQLServer really works. Kalen Delaney has been working with the SQL Server product team forover a decade， spending countless hours with developers breaking through the curse ofknowledge and then capturing the result in an incredibly clear form that allows intermediateto advanced users to wring the most from the capabilities of SQL Server. In Microsoft SO1Server2008 Internals， Kalen isjoined by four SQL Server experts who also share the giftof breaking the curse. Conor Cunningham and Paul Randal have years of experience asSQL Server product developers， and each of them is both a deep technical expert and agifted communicator. Kimberly Tripp and Adam Machanic both combine a passion to reallyunderstand how things work and to then effectively share it with others. Kimberly and Adamare both standing-room-only speakers at SQL Server events. This team has captured andincorporated the details of key architectural changes for SQL Server 2008， resulting in a new， comprehensive internals reference for SQL Server.

内容概要

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。

作者简介

Kalen Delanay世界知名的SQL Server专家。
微软SQLSewer MVP。
从1 987年供职Sybase时与微软合作开发最早的SQL Server版本算起。
她的SQL Sewer研发经验已达20多年。
她本人以对SQL Server底层技术的精湛造诣享誉业内。
所著Inside Microsoft SQL Server系列（本书前身）长期以来被奉为圣经级著作。
Paul S．Randal和KJmberly L．Tripp夫妇世界知名的SQL Server专家。
微软SQL Sewer MVP。
他们都曾长期效力于微软SQL Sewer开发团队。
Randal更是在SQL Sewer多个版本中负责存储引擎的开发。
Conor Cunningham目前担任SQL Sewerver引擎主架构师。
负责下一代引擎的设计和开发。
Adam Machanic世界知名的SQL Server专家。
微软SQL Server MVP。
著名社区SQLblog．com创始人。
名著《SQL Sewer 2005编程艺术》的作者。
Ben Nevarez资深D8A。
有多年SQL Sewer管理经验。

书籍目录

1 SQL Server 2008 Architecture and Configuration SQL Server Editions SQL Server Metadata Compatibility Views Catalog Views Other Metadata Components of the SQL Server Engine Observing Engine Behavior Protocols The Relational Engine The Storage Engine The SQLOS NUMA Architecture The Scheduler SQL Server Workers Binding Schedulers to CPUs The Dedicated Administrator Connection (DAC) Memory The Buffer Pool and the Data Cache Access to In-Memory Data Pages Managing Pages in the Data Cache The Free Buffer List and the Lazywriter Checkpoints Managing Memory in Other Caches Sizing Memory Sizing the Buffer Pool SQL Server Resource Governor Resource Governor Overview Resource Governor Controls Resource Governor Metadata SQL Server 2008 Configuration Using SQL Server Configuration Manager Configuring Network Protocols Default Network Configuration Managing Services SQL Server System Configuration Operating System Configuration Trace Flags SQL Server Configuration Settings The Default Trace Final Words 2 Change Tracking, Tracing, and Extended Events The Basics: Triggers and Event Notifi cations Run-Time Trigger Behavior Change Tracking Change Tracking Configuration Change Tracking Run-Time Behavior Tracing and Profiling SQL Trace Architecture and Terminology Security and Permissions Getting Started: Profi ler Server-Side Tracing and Collection Extended Events Components of the XE Infrastructure Event Sessions Extended Events DDL and Querying Summary 3 Databases and Database Files System Databases master model tempdb The Resource Database msdb Sample Databases AdventureWorks pubs Northwind Database Files Creating a Database A CREATE DATABASE Example Expanding or Shrinking a Database Automatic File Expansion Manual File Expansion Fast File Initialization Automatic Shrinkage Manual Shrinkage Using Database Filegroups The Default Filegroup A FILEGROUP CREATION Example Filestream Filegroups Altering a Database ALTER DATABASE Examples Databases Under the Hood Space Allocation Setting Database Options State Options Cursor Options Auto Options SQL Options Database Recovery Options Other Database Options Database Snapshots Creating a Database Snapshot Space Used by Database Snapshots Managing Your Snapshots The tempdb Database Objects in tempdb Optimizations in tempdb Best Practices tempdb Space Monitoring Database Security Database Access Managing Database Security Databases vs. Schemas Principals and Schemas Default Schemas Moving or Copying a Database Detaching and Reattaching a Database Backing Up and Restoring a Database Moving System Databases Moving the master Database Compatibility Levels Summary 4 Logging and Recovery Transaction Log Basics Phases of Recovery Reading the Log Changes in Log Size Virtual Log Files Observing Virtual Log Files Automatic Truncation of Virtual Log Files Maintaining a Recoverable Log Automatic Shrinking of the Log Log File Size Backing Up and Restoring a Database Types of Backups Recovery Models Choosing a Backup Type Restoring a Database Summary 5 Tables Creating Tables Naming Tables and Columns Reserved Keywords Delimited Identifiers Naming Conventions Data Types Much Ado About NULL User-Defi ned Data Types IDENTITY Property Internal Storage The sys.indexes Catalog View Data Storage Metadata Data Pages Examining Data Pages The Structure of Data Rows Finding a Physical Page Storage of Fixed-Length Rows

Storage of Variable-Length Rows Storage of Date and Time Data Storage of sql_variant Data Constraints Constraint Names and Catalog View Information Constraint Failures in Transactions and Multiple-Row Data Modifi cations Altering a Table Changing a Data Type Adding a New Column Adding, Dropping, Disabling, or Enabling a Constraint Dropping a Column Enabling or Disabling a Trigger Internals of Altering Tables Heap Modifi cation Internals Allocation Structures Inserting Rows Deleting Rows Updating Rows Summary 6 Indexes: Internals and Management Overview SQL Server Index B-trees Tools for Analyzing Indexes Using the dm_db_index_physical_stats DMV Using DBCC IND Understanding Index Structures The Dependency on the Clustering Key Nonclustered Indexes Constraints and Indexes Index Creation Options IGNORE_DUP_KEY STATISTICS_NORECOMPUTE MAXDOP Index Placement Constraints and Indexes Physical Index Structures Index Row Formats Clustered Index Structures The Non-Leaf Level(s) of a Clustered Index Analyzing a Clustered Index Structure Nonclustered Index Structures Special Index Structures Indexes on Computed Columns and Indexed Views Full-Text Indexes Spatial Indexes XML Indexes Data Modifi cation Internals Inserting Rows Splitting Pages Deleting Rows Updating Rows Table-Level vs Index-Level Data Modifi cation Logging Locking Fragmentation Managing Index Structures Dropping Indexes ALTER INDEX Detecting Fragmentation Removing Fragmentation Rebuilding an Index Summary 7 Special Storage Large Object Storage Restricted-Length Large Object Data (Row-Overflow Data) Unrestricted-Length Large Object Data Storage of MAX-Length Data Filestream Data Enabling Filestream Data for SQL Server Creating a Filestream-Enabled Database Creating a Table to Hold Filestream Data Manipulating Filestream Data Metadata for Filestream Data Performance Considerations for Filestream Data Sparse Columns Management of Sparse Columns Column Sets and Sparse Column Manipulation Physical Storage Metadata Storage Savings with Sparse Columns Data Compression Vardecimal Row Compression Page Compression Table and Index Partitioning Partition Functions and Partition Schemes Metadata for Partitioning The Sliding Window Benefits of Partitioning Summary 8 The Query Optimizer Overview Tree Format What Is Optimization? How the Query Optimizer Explores Query Plans Rules Properties Storage of Alternatives—The “ Memo ” Operators Optimizer Architecture Before Optimization Simplifi cation Trivial Plan/Auto-Parameterization Limitations The Memo—Exploring Multiple Plans Effi ciently Statistics, Cardinality Estimation, and Costing Statistics Design Density/Frequency Information Filtered Statistics String Statistics Cardinality Estimation Details Limitations Costing Index Selection Filtered Indexes Indexed Views Partitioned Tables Partition-Aligned Index Views Data Warehousing Updates Halloween Protection Split/Sort/Collapse Merge Wide Update Plans Sparse Column Updates Partitioned Updates Locking Distributed Query Extended Indexes Full-Text Indexes XML Indexes Spatial Indexes Plan Hinting Debugging Plan Issues {HASH | ORDER} GROUP {MERGE | HASH | CONCAT } UNION FORCE ORDER, {LOOP | MERGE | HASH } JOIN INDEX=indexname | indexid FORCESEEK FAST number_rows MAXDOP N OPTIMIZE FOR PARAMETERIZATION {SIMPLE | FORCED} NOEXPAND USE PLAN Summary 9 Plan Caching and Recompilation The Plan Cache Plan Cache Metadata Clearing Plan Cache Caching Mechanisms Adhoc Query Caching Optimizing for Adhoc Workloads

Simple Parameterization Prepared Queries Compiled Objects Causes of Recompilation Plan Cache Internals Cache Stores Compiled Plans
Execution Contexts Plan Cache Metadata Handles sys.dm_exec_sql_text
sys.dm_exec_query_plan sys.dm_exec_text_query_plan sys.dm_exec_cached_plans
sys.dm_exec_cached_plan_dependent_objects sys.dm_exec_requests
sys.dm_exec_query_stats Cache Size Management Costing of Cache Entries Objects in Plan Cache: The Big Picture Multiple Plans in Cache When to Use Stored Procedures and Other Caching Mechanisms Troubleshooting Plan Cache Issues Wait Statistics Indicating Plan Cache Problems Other Caching Issues Handling Problems with Compilation and Recompilation
Plan Guides and Optimization Hints Summary 10 Transactions and Concurrency
Concurrency Models Pessimistic Concurrency Optimistic Concurrency
Transaction Processing ACID Properties Transaction Dependencies Isolation Levels Locking Locking Basics Spinlocks Lock Types for User Data
Lock Modes Lock Granularity Lock Duration Lock Ownership
Viewing Locks Locking Examples Lock Compatibility Internal Locking Architecture
Lock Partitioning Lock Blocks Lock Owner Blocks syslockinfo Table
Row-Level Locking vs Page-Level Locking Lock Escalation Deadlocks
Row Versioning Overview of Row Versioning Row Versioning Details
Snapshot-Based Isolation Levels Choosing a Concurrency Model Controlling Locking
Lock Hints Summary 11 DBCC Internals Getting a Consistent View of the Database
Obtaining a Consistent View Processing the Database Effi ciently Fact Generation
Using the Query Processor Batches Reading the Pages to Process Parallelism
Primitive System Catalog Consistency Checks Allocation Consistency Checks
Collecting Allocation Facts Checking Allocation Facts Per-Table Logical Consistency Checks Metadata Consistency Checks Page Audit Data and Index Page Processing Column Processing Text Page Processing Cross-Page Consistency Checks Cross-Table Consistency Checks Service Broker Consistency Checks
Cross-Catalog Consistency Checks Indexed-View Consistency Checks XML-Index Consistency Checks Spatial-Index Consistency Checks DBCC CHECKDB Output
Regular Output SQL Server Error Log Output Application Event Log Output
Progress Reporting Output DBCC CHECKDB Options NOINDEX Repair Options ALL_ERRORMSGS EXTENDED_LOGICAL_CHECKS
NO_INFOMSGS TABLOCK ESTIMATEONLY PHYSICAL_ONLY
DATA_PURITY Database Repairs Repair Mechanisms Emergency Mode Repair
What Data Was Deleted by Repair? Consistency-Checking Commands Other Than DBCC CHECKDB DBCC CHECKALLOC DBCC CHECKTABLE DBCC CHECKFILEGROUP DBCC CHECKCATALOG DBCC CHECKIDENT
DBCC CHECKCONSTRAINTS Summary Index

章节摘录

The SQLOS is a separate application layer at the lowest level of the SQL Server DatabaseEngine, that both SQL Server and SQL Reporting Services run atop. Earlier versions of SQLServer have a thin layer of interfaces between the storage engine and the actual operatingsystem through which SQL Server makes calls to the operating system for memory allocation,scheduler resources, thread and worker management, and synchronization objects. However,the services in SQL Server that needed to access these interfaces can be in any part of theengine. SQL Server requirements for managing memory, schedulers, synchronization objects,and so forth have become more complex. Rather than each part of the engine growing tosupport the increased functionality, a single application layer has been designed to manageall operating system resources that are specific to SQL Server. The two main functions of SQLOS are scheduling and memory management, both of whichwell talk about in detail later in this section. Other functions of SQLOS include the following:Synchronization Synchronization objects include spinlocks, mutexes, and special reader/writer locks on system resources.

媒体关注与评论

“ Kalen Delaney的书是我的案头必备之作……从中我学到了许多东西。
我将不断重读，加深自己的理解。
”　　——Jim Gray。
已故图灵奖得主，数据库技术大师　　“ 本书是剖析SQL Server底层工作原理的权威之作，值得每一位
微CSQL Server开发团队的成员研读。
”　　——David CamIpbell，微软技术院士。
SQL Server核心引擎主架构师

编辑推荐

由Kalen Delaney编写的微软SQL Server图书一直是同类图书中的佼佼者，是SQL Sewer开发人员、架构师和DBA的案头必备书。
如今。
这本新书纳入微软阵容空前的“深入解析”（Internals）系列，增加了5位SQL Server顶级专家，深入剖析了SQL Server 2008的底层机理及其对应用程序的影响，更具权威性。
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　　让Jim Gray和David Campbell拍案叫绝的圣经级著作

版权说明

本站所提供下载的PDF图书仅提供预览和简介，请支持正版图书。

更多资源请访问:http://www.tushu007.com