




SQL Server 2000 开发宝库影印丛书

Microsoft® Press

数据库与数据仓库实用操作大全 (影印版)



Database Creation, Warehousing, and Optimization

David Iseminger

Series Editor

北京大学出版社

<http://cbs.pku.edu.cn>

SQL Server 2000 开发宝库影印丛书

数据库与数据仓库实用 操作大全 (影印版)

[美] David Iseminger 著

北京大学出版社

内 容 简 介

本书集中讲述了数据库的创建和维护、数据仓库的创建和 SQL Server 2000 的优化。并对 SQL Server 2000 的数据库、表格、索引、视图、固定程序、触发器、用户自定义函数和全文索引等都有全面的阐述。同时也详述了数据仓库的组成、创建和维护，以及如何通过有效的数据库和应用程序设计、查询调试和优化工具、实用程序和性能良好的服务器等来提高数据库的性能。

本书由经验丰富的微软专家编写，内容全面，言简意赅，是从事数据库与数据仓库管理者的理想参考书。

著作权合同登记号：图字 01-2001-0714 号

图书在版编目 (CIP) 数据

数据库与数据仓库实用操作大全/美国微软公司著. —影印本.

北京：北京大学出版社，2001.3

(SQL Server 2000 开发宝库影印丛书)

ISBN 7-301-04875-0

I. 数... II. 美... III. 关系数据库-数据库管理系统, SQL Server 2000-英文
IV. TP311.138

中国版本图书馆 CIP 数据核字 (2001) 第 10138 号

书 名：数据库与数据仓库实用操作大全 (影印版)

责任著作者：[美] David Iseminger 著

标准书号：ISBN 7-301-04875-0/TP·516

出 版 者：北京大学出版社

地 址：北京市海淀区中关村北京大学校内 100871

网 址：<http://cbs.pku.edu.cn>

电 话：出版部 62752015 发行部 62765128 62754140 编辑部 62765126

电 子 信 箱：wzdh@mail.263.net.cn

印 刷 者：北京大学印刷厂印刷

发 行 者：北京大学出版社

经 销 者：新华书店

787 毫米×1092 毫米 16 开本 27.125 印张 660 千字

2001 年 3 月第 1 版 2001 年 3 月第 1 次印刷

定 价：68.00 元

出版前言

如果用一个成语来概括国内计算机图书市场的现状，当谓之“汗牛充栋”。然而，如果您是一位从事计算机应用系统开发或管理的中、高级专业人士，很可能发现这貌似种类齐全的计算机图书中，为您量身定做的并不多见。

依据多年从事计算机图书工作所积累的经验，以及与 IT 领域广泛而深入的接触所获取的信息，我们认识到，具有相当的专业深度和技术前沿性的图书，是计算机专业人员的迫切需要，当然，也是我们从事计算机图书工作、服务专业领域的一大着眼点。

基于这一点，2000 年元月，我们与微软出版社(Microsoft Press)达成合作协议，成立微软图书影印中心，独家代理微软出版社图书影印版在中国大陆的出版、发行，为 IT 业界提供及时的专业技术服务。选题和策划上的匠心独运，使得我们的影印书成为计算机图书中的标新立异者。这里，有四大特色值得读者朋友予以关注：

首先，这是微软出版社第一次授权在中国大陆影印、发行它的版权书。在选题上，可以说独辟蹊径。在内容上，立足技术广度和深度，系统推介微软产品。所有这些，都是目前国内一般计算机图书所无法比拟的。

其次，我们的理念是为国内计算机专业人员学习前沿性的微软技术服务。为此，我们不但与微软公司紧密协作与沟通，及时掌握微软最新技术动向，而且组织了精干的工作人员，倾力于微软影印书的出版和发行。

再者，微软影印书主要面向中、高级专业人员，印量有限。这类书的读者对象有较强的针对性，一般来说，包括 IT 决策人员，中、高级开发人员，以及中、高级系统管理人员。因而，我们将每套书的印数控制在 1000~2000 册之间。

最后，微软图书影印版几乎与原版书保持同步发行，最大限度地满足了国内读者跟踪微软最新技术的需求。软件升级越来越快，新软件令人目不暇接。作为技术载体之一的图书，只有迅速作出反应，把新软件介绍给读者，才能赢得他们的青睐。总之，兵贵神速，这是我们的目标。

正应验了前人的预言，21 世纪是一个信息时代。软件作为信息系统的神经，在我们生活的这个时代里发挥着举足轻重的作用，而微软公司和它推出的各种软件，更是令世人为之瞩目。我们将立足图书，继续并扩大与微软公司的合作，在中国信息产业的发展道路上留下自己的足迹。

出版者
2000 年 10 月

Contents

Acknowledgements	iii
Part 1 Introduction and Overview	1
Chapter 1 Welcome to the SQL Server 2000 Reference Library	3
Chapter 2 Failover Clustering	7
Creating a Failover Cluster	9
Upgrading to a SQL Server 2000 Failover Cluster	15
Handling a Failover Cluster Installation	15
Maintaining a Failover Cluster	21
Using SQL Server Tools with Failover Clustering	22
Failover Cluster Troubleshooting	24
Part 2 Creating and Maintaining Databases	27
Chapter 3 Databases	29
Parts of a Database	29
Files and Filegroups	30
Transaction Logs	35
Database Design Considerations	37
Creating a Database Plan	37
Online Transaction Processing vs. Decision Support	39
Normalization	41
Data Integrity	44
Data Security	45
Database Performance	46
Maintenance	47
Estimating the Size of a Database	47
Creating a Database	56
Using Raw Partitions	58
Modifying a Database	59
Expanding a Database	60
Shrinking a Database	61
Adding and Deleting Data and Transaction Log Files	62
Creating Filegroups	63
Changing the Default Filegroup	64
Setting Database Options	64
Creating a Removable Database	74

Attaching and Detaching a Database.....	75
Renaming a Database	76
Changing the Database Owner.....	77
Transaction Termination for Changing Database States	77
Viewing a Database.....	78
Displaying Database and Transaction Log Space	78
Documenting and Scripting Databases	78
Database Maintenance Plan Wizard.....	79
Deleting a Database	80
Chapter 4 Tables	81
Designing Tables.....	81
Specifying a Column Data Type	82
Text in Row Data.....	87
Autonumbering and Identifier Columns.....	90
Using Constraints, Defaults, and Null Values	90
Creating and Modifying a Table	100
Modifying Column Properties.....	102
Adding and Deleting Columns.....	104
Creating and Modifying PRIMARY KEY Constraints.....	104
Creating and Modifying FOREIGN KEY Constraints.....	105
Creating and Modifying UNIQUE Constraints.....	107
Creating and Modifying CHECK Constraints.....	108
Creating and Modifying DEFAULT Definitions.....	109
Creating and Modifying Identifier Columns	110
Viewing a Table	111
Deleting a Table	112
Chapter 5 Indexes	113
Designing an Index.....	114
Creating an Index	125
Rebuilding an Index	137
Renaming an Index.....	138
Viewing an Index	138
Deleting an Index	138
Chapter 6 Views.....	141
Scenarios for Using Views.....	142
Creating a View.....	144
Designing an Indexed View	146

Creating an Indexed View	150
Creating a Partitioned View	154
Modifying and Renaming a View	160
Modifying Data Through a View	161
Getting Information About a View	163
Deleting a View	163
Chapter 7 Stored Procedures.....	165
Extended Stored Procedures	166
Creating a Stored Procedure	167
Specifying Parameters	169
Programming Stored Procedures	173
Returning Data from a Stored Procedure	181
Executing a Stored Procedure	186
Automatic Execution of Stored Procedures	187
Modifying and Renaming a Stored Procedure	188
Recompiling a Stored Procedure	188
Viewing a Stored Procedure	189
Deleting a Stored Procedure	190
Chapter 8 Enforcing Business Rules With Triggers	191
Designing Triggers	193
Specifying When a Trigger Fires	194
Trigger Execution	194
Designing INSTEAD OF Triggers	195
Creating a Trigger.....	198
Programming Triggers.....	200
Using Triggers that Include ROLLBACK TRANSACTION	202
Using the inserted and deleted Tables.....	203
Multirow Considerations	205
Conditional INSERT Trigger.....	207
Specifying First and Last Triggers.....	209
Using Nested Triggers	211
Using INSTEAD OF Triggers	214
Using text, ntext, and image Data in INSTEAD OF Triggers	223
Activating Triggers with Implicit and Explicit Null Values	224
Modifying and Renaming a Trigger	226
Viewing a Trigger.....	226
Deleting a Trigger.....	227

- Chapter 9 User-Defined Functions 229**
 - User-Defined Functions That Return a Table Data Type 233
 - Inline User-Defined Functions 234
 - Deterministic and Nondeterministic Functions..... 237
 - Rewriting Stored Procedures as Functions 240
 - Using Extended Properties on Database Objects 241

- Chapter 10 Full-Text Indexes 245**
 - Administering Full-Text Features Using SQL Enterprise Manager 247
 - Administering Full-Text Features Using Stored Procedures and Scalar Functions 247
 - Enabling Others to Issue Full-Text Queries 249
 - Performing Investigation and Clean-up Tasks for Full Text Catalogs..... 252
 - Performing Infrequent Tasks 254
 - Maintaining Full-Text Indexes..... 255
 - Filtering Supported File Types..... 257
 - Column-Level Linguistic Analysis 258

- Part 3 Creating and Using Data Warehouses 259**

- Chapter 11 Data Warehousing Overview 261**
 - SQL Server 2000 Tools for Data Warehouses 261

- Chapter 12 Parts of a Data Warehouse 265**
 - Data Marts..... 265
 - Relational Databases 267
 - Data Sources 267
 - Data Preparation Area 268
 - Presentation Services 269
 - End-User Analysis 272

- Chapter 13 Creating a Data Warehouse 273**
 - Designing a Data Warehouse 273
 - Using Dimensional Modeling 275
 - Fact Tables 275
 - Aggregation Tables 277
 - Dimension Tables 278
 - Indexes 281
 - Creating the Data Preparation Area 281
 - Creating the Data Warehouse Database 282
 - Extracting Data from Operational Systems..... 283
 - Cleansing and Transforming Data 284

Loading Data into the Data Warehouse Database	285
Preparing Presentation Information.....	287
Distributing Data to Data Marts	287
Chapter 14 Using a Data Warehouse.....	289
SQL Queries	290
OLAP and Data Mining	291
English Query.....	291
Microsoft Office 2000	292
Web Access and Reporting.....	292
Offline OLAP Cubes	293
Third-Party Applications.....	293
Custom Applications	294
Chapter 15 Maintaining a Data Warehouse.....	295
Updating Data Warehouse Data	295
Scheduling Data Updates.....	296
Synchronizing OLAP Cubes.....	297
Updating Data Marts.....	298
Administering a Data Warehouse	298
Backing Up Data Warehouse Data	299
Automating Data Warehouse Tasks	300
Tuning Data Warehouse Performance	300
Part 4 Optimizing Database Performance	301
Chapter 16 Designing Federated Database Servers	303
Designing Partitions	304
Designing Federated Database Servers for High Availability	307
Backing Up and Restoring Federated Database Servers	308
Chapter 17 Database Design.....	309
Logical Database Design.....	309
Database Design Considerations: Data Types	310
Physical Database Design	311
RAID	311
Partitioning	319
Data Placement Using Filegroups.....	321
Index Tuning Recommendations	323
Optimizing Transaction Log Performance.....	324
Optimizing tempdb Performance.....	325
File Systems.....	326

Chapter 18 Query Tuning	327
Analyzing a Query	329
Graphically Displaying the Execution Plan Using SQL Query Analyzer	330
Logical and Physical Operators.....	335
Query Tuning Recommendations	355
Advanced Query Tuning Concepts	357
Understanding Nested Loops Joins	358
Understanding Merge Joins.....	358
Understanding Hash Joins	359
Chapter 19 Application Design	361
Networking and Performance	364
Named Pipes vs. TCP/IP Sockets.....	364
Optimizing Application Performance Using Efficient Data Retrieval.....	365
Effects of Transactions and Batches on Application Performance	367
Effects of Stored Procedures on Application Performance.....	368
Understanding and Avoiding Blocking.....	369
Optimizing Distributed Queries.....	371
Chapter 20 Optimizing Utility and Tool Performance	375
Optimizing Backup and Restore Performance	375
Optimizing Bulk Copy Performance.....	379
Optimizing DBCC Performance	382
Chapter 21 Optimizing Server Performance	383
Optimizing Server Performance Using Memory Configuration Options	383
Optimizing Server Performance Using I/O Configuration Options.....	385
Optimizing Server Performance Using Windows NT Options	386
Maximizing Throughput	386
Configuring Server Tasking	386
Configuring Virtual Memory.....	387
Index A Topics in the SQL Server 2000 Reference Library	389
Index B Index of Optimizing Database Performance Elements	421

Introduction and Overview

Part 1 of this volume explains the structure of the *Microsoft SQL Server 2000 Reference Library*, briefly introduces the *Windows Programming Reference Series* (WPRS), and provides information about failover clustering, an important part of highly available SQL Server deployments.



Welcome to the SQL Server 2000 Reference Library

Imagine fitting an entire third-grade soccer team into a Volkswagen Beetle—that's how I felt while putting together the *SQL Server 2000 Reference Library*. All told, there was about three times as much SQL Server 2000 Books Online reference material as there was room in the *SQL Server 2000 Reference Library*, so I had to figure out the most widely used, appropriate, and pertinent parts of Books Online, and have included those parts in the library you now have in your possession.

Despite those woes, I believe you'll find the *SQL Server 2000 Reference Library* full of must-have information about SQL Server 2000. From architecture and XML, to Analysis Services (formerly OLAP) and T-SQL, this reference library contains the essential reference information you need to program, administer, deploy, or optimize your SQL Server 2000 solution... without monitor-induced eyestrain.

The SQL Server 2000 Reference Library is part of the Windows Programming Reference Series (WPRS), a series of libraries dedicated to providing printed development and IT material in a timely, intelligently organized, and well-conceived manner. You can find out more about WPRS and other available reference libraries (including the *COM+ Developer's Reference Library* and the *Active Directory Developer's Reference Library*) at www.iseminger.com—a website dedicated to providing additional information about the series, and other books also by yours truly.

How the SQL Server 2000 Reference Library is Structured

The *SQL Server 2000 Reference Library* consists of six volumes, each of which focuses on one or more specific areas of SQL Server 2000. These guides and programming reference volumes have been divided into the following:

Volume 1: SQL Server 2000 Architecture and XML/Internet Support

Volume 2: Database Creation, Warehousing, and Optimization

Volume 3: Analysis Services

Volume 4: Replication and English Query

Volume 5: T-SQL Language Reference

Volume 6: T-SQL Stored Procedures and Tables Reference

Dividing the *SQL Server 2000 Reference Library* into these categories enables you, the reader, to quickly identify the volume you need, based on your task, and facilitates your maintenance of focus for that task. This approach allows you to keep one reference book open and handy, or tucked under your arm while running between server racks.

In addition to the overall library structure, each volume in the *SQL Server 2000 Reference Library* is divided into parts that concentrate on a given subject. In order to provide a quick overview of a part's contents, each begins with a Part Introduction page that outlines what you'll find therein.

Finding Related Topics and Working With the Topic Index

Throughout this library, you're going to see references to related topics; some of them within the text of a given paragraph, others placed in a special section called **Related Topics**. Since the entire body of SQL Server 2000 Books Online constitutes more topics than what you'll find in these volumes, you may occasionally come across a referenced topic that doesn't correspond to a section in this library. Don't worry; you can get to that information through the SQL Server 2000 Books Online. Most of the references *will* pertain to items in these volumes, but there might be some reference you're interested that will lead you online.

To make locating topics as easy as possible, and to enable you to quickly identify which topics are found in this library (versus which are only found in Books Online), Volumes 1 through 4 include a special index called **Topics in the SQL Server 2000 Reference Library**. There's some good information about the topic index (as I'll refer to it from now on) that will help you understand how to get the most use out of it, and out of the *SQL Server 2000 Reference Library*.

For starters, the topic index contains *only topics found in the SQL Server 2000 Reference Library*. So, say you're reading through a chapter in this library, and you come across text that refers you to another topic, such as the following:

Related Topics

Building and Processing Cubes, Updating Cubes and Dimensions

Or something like the following:

...For more information about which editions support which features, see **Features Supported by the Editions of SQL Server 2000**.

You can then look through this alphabetical listing of topics included in this library for the referenced topic, and when you find it, you'll be directed to the Volume and Chapter where that topic can be found, as shown here:

Building and Processing Cubes Volume 3, Chapter 5
Building and Processing Cubes Volume 3, Chapter 8
Features Supported by the Editions of SQL Server 2000 Volume 1, Chapter 11

If two topics with the same title exist (as shown here), each will be listed separately. As previously mentioned, most references refer to topics also found in this library, but for those that don't, you can go to Books Online to get more information.

There are important exceptions to topic references and the topic index: *T-SQL statements and other programming elements are **not** included in the topic index.* Each volume that contains programming elements (such as T-SQL stored procedures, or English Query statements) has its own index of programming elements. Rather than cluttering this index up with programming elements (or vice versa), I've provided separate indexes to help you find the information you need faster.

Since Volumes 5 and 6 are almost entirely dedicated to statement definitions, including the topic index in those volumes didn't seem like a good use of (precious) pages.

The Idea Behind SQL Server 2000 Reference Library

The *SQL Server 2000 Reference Library*, like all libraries in the Windows Programming Reference Series, is designed to deliver the most pertinent information in the most accessible way possible. The *SQL Server 2000 Reference Library* is also designed to integrate seamlessly with SQL Server 2000 Books Online (and with MSDN Online) by providing a look-and-feel that is consistent with the electronic means of disseminating SQL Server 2000 reference information. In other words, the way that a given function reference appears on the pages of this book has been designed specifically to emulate the way that Books Online presents its reference pages.

The reason for maintaining such integration is simple: make it easy for you—the administrator or developer of SQL Server 2000 solutions—to use the tools and get the ongoing information you need to do your job. By providing a “common interface” among reference resources, your familiarity with the *SQL Server 2000 Reference Library* reference material can be immediately applied to Books Online, and vice-versa. In a word, it means *consistency*.

You'll find this philosophy of consistency and simplicity applied throughout Windows Programming Reference Series publications. I've designed the series to go hand-in-hand with online resources. Such consistency lets you leverage your familiarity with electronic reference material, then apply that familiarity to enable you to get away from your computer if you'd like, take a book with you, and—in the absence of keyboards and e-mail and upright chairs—get your reading and research done. Of course, each of the *SQL Server 2000 Reference Library* volumes fits nicely right next to your mouse pad as well, even when opened to a particular reference page.

With any job, the simpler and more consistent your tools are, the more time you can spend doing work rather than figuring out how to use your tools. The structure and design of the *SQL Server 2000 Reference Library* provides you with a comprehensive, pre-sharpened toolset to quickly program, administer, or optimize SQL Server 2000 deployments.

Failover Clustering

In Microsoft SQL Server 2000 Enterprise Edition, SQL Server 2000 failover clustering provides high availability support. For example, during an operating system failure or a planned upgrade, you can configure one failover cluster to fail over to any other node in the failover cluster configuration. In this way, you minimize system downtime, thus providing high server availability.

To install, configure, and maintain a failover cluster, use SQL Server Setup. For information about upgrading to a SQL Server 2000 failover cluster, see *Upgrading to a SQL Server 2000 Failover Cluster*.

Use failover clustering to:

- Install SQL Server on multiple nodes in a failover cluster. You are limited only by the number of nodes supported by the operating system.

Before installing failover clustering, you must install Microsoft Windows NT 4.0, Enterprise Edition; Microsoft Windows 2000 Advanced Server or Windows 2000 Datacenter Server; and the Microsoft Cluster Service (MSCS).

There are specific installation steps that must be followed to use failover clustering. For more information, see *Installing Failover Clustering and Handling a Failover Cluster Installation*.

- Specify multiple IP addresses for each virtual server.
SQL Server 2000 allows you to use all available network IP subnets, thereby providing alternate ways to connect if one subnet fails and increasing network scalability. For example, with a single network adaptor, a network failure can disrupt communications. However, with multiple network cards in the server, each network can be on a different IP subnet. If one subnet fails, at least one connection can continue to function. If a router fails, MSCS continues to function, and all IP addresses still work. However, if the network card on the local computer fails, communication still may be disrupted. For more information, see *Creating a Failover Cluster*.
- Administer a failover cluster from any node in the clustered SQL Server configuration. To perform setup tasks, you must be working from the node in control of the cluster disk resource. For more information, see *Creating a Failover Cluster*.