SQL Server 2000 开发宝库影印丛书

Microsoft Press

SQL Server 2000 分析服务（影印版）

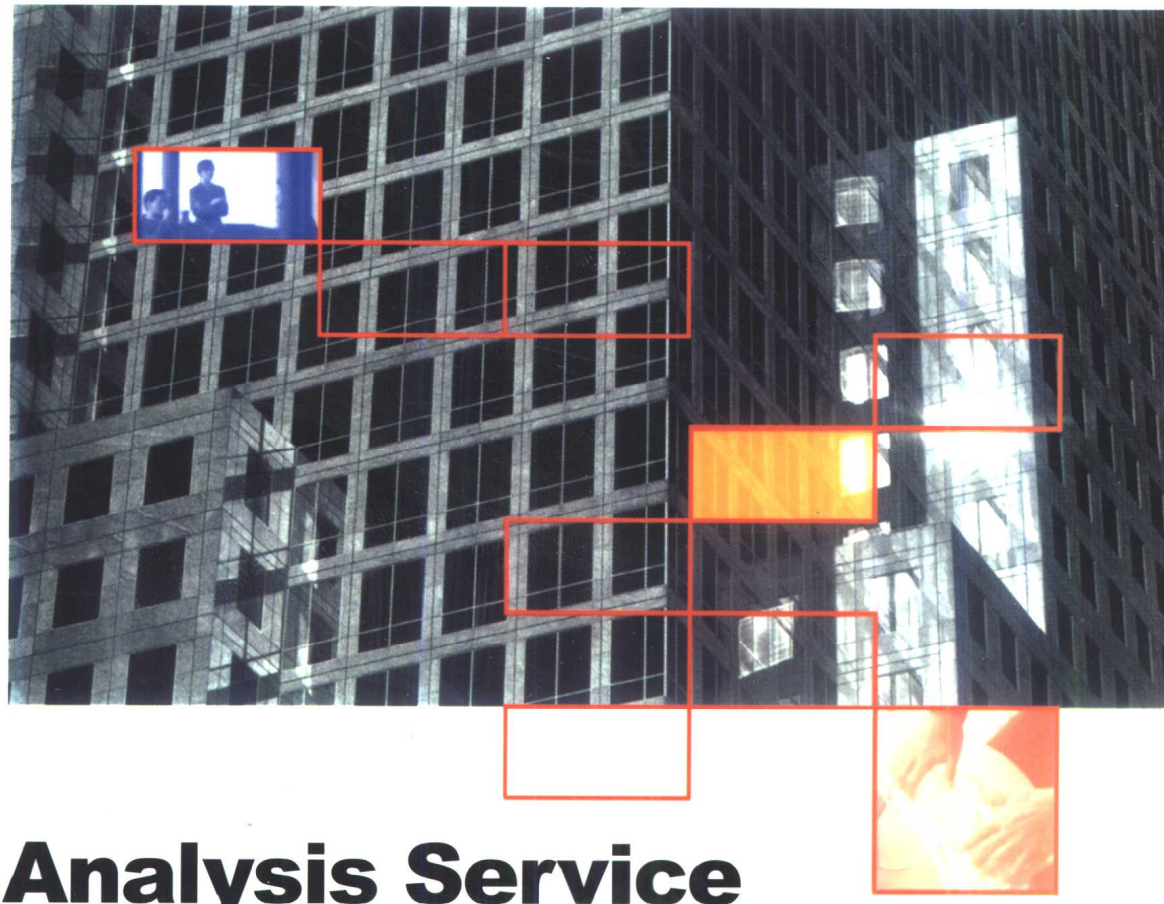

Analysis Service

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

SQL Server 2000 分析服务

（影印版）

[美] David Iseminger 著

北京大学出版社

内 容 简 介

本书包含了 SQL Server 2000 分析服务的提纲和详细内容。包括 SQL Server 2000 分析服务新增内容，以及如何安装这些服务。论述了 SQL Server 2000 分析服务结构、数据仓库、数据采集和 OLAP 的详细内容。为您提供了如何管理分析服务、MXD 和相关决策工具的全面信息。

本书由微软公司专家编写，实用性强，是系统管理员和程序开发人员的必备参考书。

著作权合同登记号：图字 01-2001-0714 号

图书在版编目（CIP）数据

SQL Server 2000 分析服务/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.S... 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 10136 号

书　　名：SQL Server 2000 分析服务（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：wdzh@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　42.125 印张　1025 千字
2001 年 3 月第 1 版　2001 年 3 月第 1 次印刷
定　　价：108.00 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提供服务。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

Introduction 1

Chapter 1 What's New In Analysis Services 5

Cube Enhancements 6
Dimension Enhancements 8
Data Mining Enhancements 10
Security Enhancements 11
Client Connectivity Enhancements in PivotTable Service 12
Other Enhancements 12

Chapter 2 Installing Analysis Services 15

Hardware and Software Requirements for Installing Analysis Services 16
Running Setup 17
Setup Parameters and Silent Installation 19
Reinstalling Analysis Services 20
Stopping or Removing Analysis Services 21
Upgrading from an Earlier Version 21
Backward Compatibility 22
7.0 Analysis Services Client and Local Cube Support 23
Supported Migration Paths for Analysis Services Repositories 25
Archiving and Restoring Databases Between Versions of Analysis Services 26

Chapter 3 Analysis Services Architecture 27

Analysis Services Features 27
Ease of Use 28
Flexible Data Model 31
Scalability 34
Integration 36
Widely Supported APIs and Open Architecture 38
Server and Client Architecture 39
Server Architecture 39
Client Architecture 40
Object Architecture 42
Object Architecture Overview 43
Analysis Server 44
Databases 46
Data Sources 46

Dimensions 48
Levels and Members 76
Measures 87
Cells 101
Cubes 115
Partitions 136
Aggregations 148
Roles 150
Commands 154
Member Properties 159
Data Mining Models 164
Data Mining Columns 176
Security and Authentication 186
Administrator Security 187
End-User Security 188
Operational Considerations 229

Chapter 4 Data Warehousing and OLAP 233
About Data Warehouses 234
About OLAP 235
OLAP and Data Warehouses 235
Maintaining OLAP Data 237

Chapter 5 Administering Analysis Services 241
Before Administering Analysis Services 241
Administrative Tools 242
Analysis Manager 242
Using Active Directory with Analysis Services 248
Using Data Transformation Services with Analysis Services 249
Using MDX with Analysis Services 249
Msmdarch Command 250
Administrative Tasks 250
Configuring Analysis Servers 252
Creating Prerequisite Objects for Cubes 254
Building and Processing Cubes 261
Creating Security Roles 279
Managing Partitions 282
Enhancing Dimensions with Optional Features 289
Enhancing Cubes with Optional Features 296
Updating Cubes and Dimensions 315

Building and Using Data Mining Models 318
Archiving, Restoring, and Copying Data 328
Analyzing and Optimizing Performance 333
Automating and Scheduling Administrative Tasks 352

Chapter 6 Administrator's Reference 361

Specifications and Limits 361
SQL 364
OLE DB Compliance 368
OLE DB for Data Mining Compliance 372

Chapter 7 Multidimensional Expressions (MDX) 377

MDX Overview 378
Introduction to MDX 378
Key Concepts in MDX 379
Comparison of SQL and MDX 383
Basic MDX 384
The Basic MDX Query 385
Members, Tuples, and Sets 387
Axis and Slicer Dimensions 393
Establishing Cube Context 396
Advanced MDX 397
Creating and Using Property Values 398
Building Named Sets in MDX 415
Building Calculated Members in MDX 418
Building Caches in MDX 426
Building Calculated Cells in MDX 428
Creating and Using User-Defined Functions in MDX 431
Using Writebacks 434
Using DRILLTHROUGH to Retrieve Source Data 436
Understanding Pass Order and Solve Order 437
Effective MDX 448
Comments in MDX 448
Working with Empty Cells 449
Creating a Cell Within the Context of a Cube 454
Working with the RollupChildren Function 455
WHERE Clause Overrides 457
MDX Functions in Analysis Services 458
MDX Function Reference 458
Registered Function Libraries 551
User-Defined Functions with MDX Syntax 554

Chapter 8 How To Perform Common Analysis Services Tasks **557**
Configuring Analysis Servers 557
How to register an Analysis server 557
How to add a linked Analysis server using SQL Server Enterprise Manager 558
How to start the Migrate Repository Wizard 559
How to configure Analysis Services for the Web 559
Creating Cubes' Prerequisite Objects 560
How to create a database 560
How to specify a data source 561
How to create a shared dimension using the Dimension Wizard 561
How to create a shared dimension using Dimension Editor 562
How to browse a shared dimension 563
How to browse a private dimension 563
How to create a virtual dimension based on member properties 563
How to create a virtual dimension based on columns 564
Building and Processing Cubes 566
How to start the Cube Wizard 567
How to build a cube with Cube Editor 567
How to create a private dimension 568
How to add a multiple-column measure to a cube 569
How to start the Virtual Cube Wizard 569
How to start the Storage Design Wizard 569
How to process a cube 570
How to process a virtual cube 570
How to view an SQL statement 571
How to change the temporary file folder used by Analysis Services 571
How to browse sample data in Cube Editor 571
Creating Security Roles 572
How to create a database role 572
How to create a cube role, change its default values, and specify cell security 574
How to create a mining model role and change its default values 578
How to create a custom rule for dimension security in a database role 579
How to create a custom rule for dimension security in a cube role 581
How to create a custom rule for cell security 584
Managing Cube Storage 584
How to start the Partition Wizard 585
How to process a partition 585
How to merge two partitions 585
How to specify a data source for a linked cube 586

How to create a linked cube ... 586
How to process a linked cube ... 587
Enhancing Dimensions ... 587
How to create a member property in a shared dimension ... 588
How to create a custom rollup operator for a shared dimension ... 588
How to create a custom rollup operator for a private dimension ... 589
How to create or select a column to store custom member formulas ... 591
How to create a custom member formula in a write-enabled dimension ... 592
How to create member groups ... 592
How to create a dimension with multiple hierarchies in the Dimension Wizard ... 594
How to create a dimension with multiple hierarchies in Dimension Editor ... 595
How to add a hierarchy to an existing dimension ... 596
Enhancing Cubes with Optional Features ... 597
How to create calculated cells ... 598
How to import calculated cells into a virtual cube ... 599
How to edit a calculated cells definition ... 599
How to rename a calculated member ... 600
How to create a calculated member in a regular cube ... 600
How to import a calculated member into a virtual cube ... 602
How to edit a calculated member in a virtual cube ... 602
How to create a calculated member in a virtual cube ... 603
How to create a named set in a regular cube ... 604
How to import a named set into a virtual cube ... 605
How to edit a named set in a virtual cube ... 605
How to create a named set in a virtual cube ... 606
How to create an action in a regular cube ... 606
How to import an action into a virtual cube ... 607
How to edit an action in a virtual cube ... 607
How to create an action in a virtual cube ... 608
How to test an action ... 608
How to write-enable a cube ... 609
How to browse writeback data for a cube ... 610
How to delete writeback data for a cube and write-disable it ... 610
How to convert a cube's writeback data to a partition and write-disable the cube ... 610
How to specify drillthrough options for a regular cube ... 611
How to specify drillthrough options for a virtual cube ... 612
How to specify drillthrough options for a linked cube ... 613
How to specify drillthrough options for a partition ... 615
Updating Cubes and Dimensions ... 616

How to incrementally update a cube 616
How to refresh data in a cube 618
How to incrementally update a shared dimension 618
How to rebuild the structure of a shared dimension 618
Managing Data Mining Models 619
How to start the Mining Model Wizard 619
How to start Relational Mining Model Editor 620
How to start OLAP Mining Model Editor 620
How to process a data mining model 620
How to start Mining Model Role Manager 621
How to add a role to a data mining model 621
Archiving, Restoring, and Copying Data 622
How to archive an Analysis Services database using Analysis Manager 622
How to archive an Analysis Services database using the msmdarch command 623
How to restore an Analysis Services database using Analysis Manager 624
How to restore an Analysis Services database using the msmdarch command 625
How to copy and paste an object 626
Analyzing and Optimizing Performance 627
How to start the Usage Analysis Wizard 627
How to start the Usage-Based Optimization Wizard 627
Automating and Scheduling Administrative Tasks 627
How to create an Analysis Services Processing task 628
How to create a Relational Data Mining Prediction Query task 629

Index Topics in the SQL Server 2000 Reference Library 631

Introduction

Imagine holding your tuba as a carry-on item for your overbooked, middle-seat, coach-class flight from Seattle to Miami—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 in 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 found only 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 *topics found only 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.

CHAPTER 1

What's New In Analysis Services

Microsoft SQL Server 2000 extends and renames the former OLAP Services component, now called Analysis Services. Many new and improved features significantly enhance the analysis capabilities of the acclaimed OLAP Services introduced in SQL Server version 7.0. In this release, Analysis Services introduces data mining, which can be used to discover information in OLAP cubes and relational databases.

The following table details the sections found in this chapter, and provides a brief description of each.

Topic	Description
Cube Enhancements	New cube types and enhanced cube functionality substantially extend the scalability and functionality of Analysis Services.
Dimension Enhancements	New dimension and hierarchy types, features, and improvements extend the analysis capabilities of cubes.
Data Mining Enhancements	New in this release, data mining is integrated into online analysis and can be used to discover information in OLAP cubes and relational databases.
Security Enhancements	Security enhancements include using roles on cube cells and dimension members, additional authentication methods, and improved enforcement.
Client Connectivity Enhancements in PivotTable Service	Client applications can use many new features and enhancements such as data mining, HTTP or HTTPS connections, additional dimension types, and cell allocation for writeback.
Other Enhancements	Other enhancements provide a variety of new features including multiuser administration, MDX Builder, additional Multidimensional Expressions (MDX) functions, Virtual Cube Editor, support for Active Directory, and more.

Cube Enhancements

Microsoft SQL Server 2000 Analysis Services substantially extends the scalability and functionality of OLAP cubes. You can distribute cube data across multiple servers to provide more storage capacity, create linked cubes to distribute end-user access to information without duplicating cube data, create cubes that are updated in real time as data changes, and use a number of other new features to create cubes that address your specific business needs.

Distributed Partitioned Cubes

You can create distributed partitioned cubes by using remote partitions that distribute a cube's data among multiple Analysis servers. A distributed partitioned cube is administered on a central Analysis server. For more information, see Remote Partitions.

Real-Time OLAP

Real-time OLAP provides a multidimensional OLAP view of data that is continually updated as the underlying data changes. Real-time cubes implement real-time OLAP by using ROLAP storage for partitions and dimensions, new SQL Server 2000 indexed views for aggregations, and automatic notification by the SQL Server 2000 relational engine when data changes. Real-time cubes provide the capability to develop new categories of OLAP solutions such as call-center management, stock market analysis, or campaign management. For more information, see Real-Time Cubes.

Linked Cubes

A cube can be stored on a single Analysis server and then defined as a linked cube on other Analysis servers. End users connected to any of these Analysis servers can then access the cube. This arrangement avoids the more costly alternative of storing and maintaining copies of a cube on multiple Analysis servers. Linked cubes can be connected using TCP/IP or HTTP. To end users, a linked cube looks like a regular cube. For more information, see Linked Cubes.

Indexed Views for Aggregations

Indexed views for increased performance and flexibility are used instead of aggregation tables for ROLAP partitions if the partition's source data is stored in SQL Server 2000 and if certain criteria are met. For more information, see Indexed Views for ROLAP Partitions.

Cube Processing

You can use lazy aggregations to make cube data available to end users while aggregations are being calculated. When processing cubes for which the underlying data contains dimension key errors, you can elect to stop processing on key errors, stop processing after a specified number of errors, or ignore all key errors. You can have errors logged to a file for later review. For more information, see Processing Cubes.

Calculated Cells

You can specify formulas that apply to individual cells or to sets of cells in a cube. These formulas can contain conditional calculations that compute a new value for a cell or set of cells based on values in the cell or cells, or on values in other cells in the cube. Calculated cells use Multidimensional Expressions (MDX) expressions and you can specify calculations to be performed in multiple passes. Calculated cells can be used in complex financial modeling and budgeting applications; for example, you can specify a default value such as a percentage of a parent cell if the cell value is zero, or to use the actual value if it is not zero. For more information, see Calculated Cells.

Drillthrough

Client applications that support drillthrough can now allow end users to select a cube cell and retrieve a result set from the source data for that cell. You can use roles to control user access to the drillthrough functionality. For more information, see Specifying Drillthrough Options.

Actions

Actions enable end users to act upon the outcomes of their analyses. An action is a predefined operation that an end user can initiate upon a selected cube or portion of a cube. The operation can start an application with the selected item as a parameter or retrieve information about the selected item. A wizard is provided to help you create actions. For more information, see Actions.

Distinct Count

You can use the new **DistinctCount** aggregate function to analyze the number of unique occurrences of events or transactions in your data, such as unique users visiting a Web site. For more information, see DistinctCount.

Hidden Cube Elements

You can hide complete cubes, dimensions, levels, measures, or member properties from end users who browse cubes with client applications. The visibility of these objects is controlled by the **Visible** property. For more information, see Properties Pane (Cube Editor Data View) and Properties Pane (Dimension Editor Data View).