1. Record Nr. UNINA9910972860803321

Autore	Webb Chris
Titolo	Expert cube development with Microsoft SQL Server 2008 Analysis Services : design and implement fast, scalable, and maintainable cubes / / Chris Webb, Alberto Ferrari, Marco Russo
Pubbl/distr/stampa	Birmingham, UK, : Packt Publishing, 2009
ISBN	9786612238154 9781282238152 1282238159 9781847197238 184719723X 9781847197221 1847197221
Edizione	[1st ed.]
Descrizione fisica	1 online resource (360 p.)
Collana	From technologies to solutions
Altri autori (Persone)	FerrariAlberto RussoMarco
Disciplina	005.7585
Soggetti	Client/server computing OLAP technology Data mining
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Note generali	Includes index.
Nota di contenuto	Intro -- Expert Cube Development with Microsoft SQL Server 2008 Analysis Services -- Table of Contents -- Expert Cube Development with Microsoft SQL Server 2008 Analysis Services -- Credits -- About the Authors -- About the Reviewers -- Preface -- What this book covers -- What you need for this book -- Who this book is for -- Conventions -- Reader feedback -- Customer support -- Downloading the example code and database for the book -- Errata -- Piracy -- Questions -- 1. Designing the Data Warehouse for Analysis Services -- The source database -- The OLTP database -- The data warehouse -- The data mart -- Data modeling for Analysis Services -- Fact tables and dimension tables -- Star schemas and snowflake schemas -- Junk dimensions -- Degenerate dimensions -- Slowly Changing Dimensions

-- Bridge tables, or factless fact tables -- Snapshot and transaction fact tables -- Updating fact and dimension tables -- Natural and surrogate keys -- Unknown members, key errors, and NULLability -- Physical database design for Analysis Services -- Multiple data sources -- Data types and Analysis Services -- SQL queries generated during cube processing -- Dimension processing -- Dimensions with joined tables -- Reference dimensions -- Fact dimensions -- Distinct count measures -- Indexes in the data mart -- Usage of schemas -- Naming conventions -- Views versus the Data Source View -- Summary -- 2. Building Basic Dimensions and Cubes -- Choosing an edition of Analysis Services -- Setting up a new Analysis Services project -- Creating data sources -- Creating Data Source Views -- Designing simple dimensions -- Using the 'New Dimension' wizard -- Using the Dimension Editor -- Adding new attributes -- Configuring a Time dimension -- Creating user hierarchies -- Configuring attribute relationships -- Building a Simple Cube -- Using the 'New Cube' wizard -- Deployment.
Processing -- Summary -- 3. Designing More Complex Dimensions -- Grouping and Banding -- Grouping -- Banding -- Slowly Changing Dimensions -- Type I SCDs -- Type II SCDs -- Modeling attribute relationships on a Type II SCD -- Handling member status -- Type III SCDs -- Junk dimensions -- Ragged hierarchies -- Parent/child hierarchies -- Ragged hierarchies with HideMemberIf -- Summary -- 4. Measures and Measure Groups -- Measures and aggregation -- Useful properties of measures -- Format String -- Display folders -- Built-in measure aggregation types -- Basic aggregation types -- Distinct Count -- None -- Semi-additive aggregation types -- By Account -- Dimension calculations -- Unary operators and weights -- Custom Member Formulas -- Non-aggregatable values -- Measure groups -- Creating multiple measure groups -- Creating measure groups from dimension tables -- Handling different dimensionality -- Handling different granularities -- Non-aggregatable measures: a different approach -- Using linked dimensions and measure groups -- Role-playing dimensions -- Dimension/measure group relationships -- Fact relationships -- Referenced relationships -- Data mining relationships -- Summary -- 5. Adding Transactional Data such as Invoice Line and Sales Reason -- Details about transactional data -- Drillthrough -- Actions -- Drillthrough actions -- Drillthrough Columns order -- Drillthrough and calculated members -- Drillthrough modeling -- Drillthrough using a transaction details dimension -- Drillthrough with ROLAP dimensions -- Drillthrough on Alternate Fact Table -- Drillthrough recap -- Many-to-many dimension relationships -- Implementing a many-to-many dimension relationship -- Advanced modelling with many-to-many relationships -- Performance issues -- Summary -- 6. Adding Calculations to the Cube -- Different kinds of calculated members.
Common calculations -- Simple calculations -- Referencing cell values -- Aggregating members -- Year-to-dates -- Ratios over a hierarchy -- Previous period growths -- Same period previous year -- Moving averages -- Ranks -- Formatting calculated measures -- Calculation dimensions -- Implementing a simple calculation dimension -- Calculation dimension pitfalls and problems -- Attribute overwrite -- Limitations of calculated members -- Calculation dimension best practices -- Named sets -- Regular named sets -- Dynamic named sets -- Summary -- 7. Adding Currency Conversion -- Introduction to currency conversion -- Data collected in a single currency -- Data collected in a multiple currencies -- Where to perform currency conversion -- The Add Business Intelligence Wizard -- Concepts and

prerequisites -- How to use the Add Business Intelligence wizard -- Data collected in a single currency with reporting in multiple currencies -- Data collected in multiple currencies with reporting in a single currency -- Data stored in multiple currencies with reporting in multiple currencies -- Measure expressions -- DirectSlice property -- Writeback -- Summary -- 8. Query Performance Tuning -- How Analysis Services processes queries -- Performance tuning methodology -- Designing for performance -- Performance-specific design features -- Partitions -- Why partition? -- Building partitions -- Planning a partitioning strategy -- Unexpected partition scans -- Aggregations -- Creating an initial aggregation design -- Usage-based optimization -- Monitoring partition and aggregation usage -- Building aggregations manually -- Common aggregation design issues -- MDX calculation performance -- Diagnosing Formula Engine performance problems -- Calculation performance tuning -- Tuning algorithms used in MDX -- Using named sets to avoid recalculating set expressions.
Using calculated members to cache numeric values -- Tuning the implementation of MDX -- Caching -- Formula cache scopes -- Other scenarios that restrict caching -- Cache warming -- Create Cache statement -- Running batches of queries -- Scale-up and scale-out -- Summary -- 9. Securing the Cube -- Sample security requirements -- Analysis Services security features -- Roles and role membership -- Securable objects -- Creating roles -- Membership of multiple roles -- Testing roles -- Administrative security -- Data security -- Granting read access to cubes -- Cell security -- Dimension security -- Applying security to measures -- Dynamic security -- Dynamic dimension security -- Dynamic security with stored procedures -- Dimension security and parent/child hierarchies -- Dynamic cell security -- Accessing Analysis Services from outside a domain -- Managing security -- Security and query performance -- Cell security -- Dimension security -- Dynamic security -- Summary -- 10. Productionization -- Making changes to a cube in production -- Managing partitions -- Relational versus Analysis Services partitioning -- Building a template partition -- Generating partitions in Integration Services -- Managing processing -- Dimension processing -- Partition processing -- Lazy Aggregations -- Processing reference dimensions -- Handling processing errors -- Managing processing with Integration Services -- Push-mode processing -- Proactive caching -- Analysis Services data directory maintenance -- Backup -- Copying databases between servers -- Summary -- 11. Monitoring Cube Performance and Usage -- Analysis Services and the operating system -- Resources shared by the operating system -- CPU -- Memory -- I/O operations -- Tools to monitor resource consumption -- Windows Task Manager -- Performance counters -- Resource Monitor -- Analysis Services memory management.
Memory differences between 32 bit and 64 bit -- Controlling the Analysis Services Memory Manager -- Out of memory conditions in Analysis Services -- Sharing SQL Server and Analysis Services on the same machine -- Monitoring processing performance -- Monitoring processing with trace data -- SQL Server Profiler -- ASTrace -- XMLA -- Flight Recorder -- Monitoring Processing with Performance Monitor counters -- Monitoring Processing with Dynamic Management Views -- Monitoring query performance -- Monitoring queries with trace data -- Monitoring queries with Performance Monitor counters -- Monitoring queries with Dynamic Management Views -- MDX Studio -- Monitoring usage -- Monitoring Usage with Trace Data -- Monitoring usage with Performance Monitor counters -- Monitoring usage with Dynamic

Management Views -- Activity Viewer -- How to build a complete monitoring solution -- Summary -- Index.

Sommario/riassunto

Design and implement fast, scalable and maintainable cubesKey FeaturesA real-world guide to designing cubes with Analysis Services 2008Model dimensions and measure groups in BI Development StudioImplement security, drill-through, and MDX calculationsLearn how to deploy, monitor, and performance-tune your cubeFilled with best practices and useful hints and tipsBook DescriptionMicrosoft's SQL Server Analysis Services 2008 is an OLAP server that allows users to analyze business data quickly and easily. However, designing cubes in Analysis Services can be a complex task: it's all too easy to make mistakes early on in development that lead to serious problems when the cube is in production. Learning the best practices for cube design before you start your project will help you avoid these problems and ensure that your project is a success. This book offers practical advice on how to go about designing and building fast, scalable, and maintainable cubes that will meet your users' requirements and help make your Business Intelligence project a success. This book gives readers insight into the best practices for designing and building Microsoft Analysis Services 2008 cubes. It also provides details about server architecture, performance tuning, security, and administration of an Analysis Services solution. In this book, you will learn how to design and implement Analysis Services cubes. Starting from designing a data mart for Analysis Services, through the creation of dimensions and measure groups, to putting the cube into production, we'll explore the whole of the development lifecycle. This book is an invaluable guide for anyone who is planning to use Microsoft Analysis Services 2008 in a Business Intelligence project.What you will learnBuild a data mart suitable for use with Analysis ServicesCreate and configure an Analysis Services project in Business Intelligence Development StudioUse the Dimension Wizard and the Dimension Editor to build dimensionsCreate measure groups and associate them with dimensionsAdd calculations to the cube, including implementing currency conversion and a date tool dimensionExplore the security model, including dimension security and cell security, and implement dynamic securityTune queries to get the best possible performanceAutomate processing and partition creationMonitor your cube to see who s actually using itWho this book is forThis book is aimed at Analysis Services developers who already have some experience but who want to go into more detail on advanced topics, and who want to learn best practices for cube design.