

Advanced MDX – Performance, Optimization, and the MDX Script

Course Summary

Description

This is an advanced MDX course. Half of the first day is devoted to advanced MDX functions and sets. The second half of the day is focussed on effectively using the MDX Script. The second day is devoted completely to examining strategies for monitoring and optimizing MDX. Analysis Services in SQL Server 2012 is used in the course, but almost all of the content is also applicable to Analysis Services in SQL Server 2008.

Objectives

At the end of this course, students will be able to:

- Analyze and optimize the performance of MDX calculations
- Understand the situations where MDX can be modified to improve performance and those situations where other strategies must be used
- Re-write MDX calculations so that they will work with specific client applications
- Use all the MDX functions
- Use the MDX Script to create complex business logic
- Use the MDX Script to create complex Date Calculation Hierarchies

Topics

- Advanced Functions
- Advanced Set Topics
- The MDX Scripting
- Building a Time Calculation Hierarchy
- MDX Performance Analysis
- Enabling Block Computation
- Other Calculation Performance Strategies
- Using Subselects and Subcubes

Audience

This course is intended for individuals that are developing or supporting Analysis Services multidimensional cubes.

Prerequisites

Before attending this course, it is recommended that students have the following skills:

- Ability to create and understand MDX calculations
- Thorough understanding of sets and tuples
- Ability to use all the basic MDX functions
- It is recommended that the students have taken the SolidQ course MDX for Business Logic

NOTE: This is an Advanced MDX course. If a student is unfamiliar with MDX before they take this course, they are unlikely to receive much benefit from it.

Duration

Two days

Advanced MDX – Performance, Optimization, and the MDX Script

Course Outline

- I. Advanced Functions**
 - A. Using VBA and Excel Functions in MDX
 - B. Using the String Functions
 - C. StrToMember, StrToSet, StrToTuple, MemberToStr, SetToStr, TupleToStr
 - D. Using the Statistical Functions
 - E. StdDev, StdDevP, Var, VarP, Covariance, CovarianceN
 - F. LinRegIntercept, LinRegPoint, LinRegSlope, LinRegR2, LinRegVariance
 - G. Using the Drilldown (UI) Functions
 - H. DrilldownLevel, DrilldownMember, DrilldownLevelBottom, etc.
 - I. Using the KPI Functions
 - J. KPIGoal, KPIStatus, KPITrend, KPIValue, etc.
 - K. Using the LinkMember Function
 - L. Using the Generate Function to Create a Set
 - M. Using the Generate Function to Create a String
 - N. Using Recursion in MDX
 - O. *LAB: Using LinkMember and recursion*
- II. Advanced Set Topics**
 - A. The Autoexists Functionality
 - B. The Existing Keyword
 - C. The Exists Function
 - D. Dynamic and Static Sets
 - E. Referencing the Sets on the Axes
 - F. *LAB: Referencing the Set on the Rows*
- III. The MDX Scripting**
 - A. The MDX Script
 - B. The Calculate Command
 - C. Setting the Default Member
 - D. Creating a Calculation Hierarchy
 - E. Scoping and Assignments
 - F. Debugging the MDX Script
 - G. The Freeze Command
 - H. Assigning Formatting in the MDX Script
 - I. Using IF in MDX Scripting
 - J. Adding Values Not Directly in Source Data
 - K. Removing Inaccurate Values
 - L. Replacing IIF with MDX Scripting
 - M. Replacing Unary Operators with MDX Script Assignments
- IV. Building a Time Calculation Hierarchy**
 - A. Understanding Time Calculation Hierarchies
 - B. The Business Intelligence (BI) Wizard
 - C. Customizing Time Calculation Hierarchies
 - D. Current Period Calculations
 - E. Relative Date Period Calculations
 - F. Period To Date and Rolling Average Calculations
 - G. Comparison, Ratio, and Forecasting Calculations
 - H. Assigning Formatting
 - I. Creating a Second Hierarchy to Choose the Type of Date
 - J. *LAB: Adding New Time Calculations*
- V. MDX Performance Analysis**
 - A. Optimization Strategy
 - B. SSAS Cache, Windows Cache, and MDX Script Cache
 - C. Using SQL Server Profiler with SSAS
 - D. Determining Storage Engine Time and Formula Engine Time
 - E. Cache Warming Strategies
 - F. *LAB: Examining Cache Behavior*
- VI. Enabling Block Computation**
 - A. Understanding Block Computation
 - B. Situations that Prevent Block Computation
 - C. Monitoring for Block Computation
 - D. Rewriting Calculations to Use Block Computation
 - E. Removing Named Sets from Aggregation Functions
 - F. Replacing Filter with Crossjoin or Exists
 - G. Replacing Count(Filter) with Sum(IIF)
 - H. Removing Late Binding
 - I. Replacing User-Defined Functions
 - J. Replacing LinkMember
 - K. When Block Computation is Slower
 - L. *LAB: Fixing Calculations to Use Block Computation*
- N. LAB: Writing MDX Script Code**

Advanced MDX – Performance, Optimization, and the MDX Script

Course Outline (cont'd)

VII. Other Calculation Performance Strategies

- A. Simplifying Complex Calculations
- B. Making Sets as Small as Possible
- C. Ordering in Set Multiplication
- D. Using Query-Created Cache
- E. Fixing Non-Varying Expressions
- F. Other Best Practices
- G. IIF Function Hints
- H. Adding Attributes for Specialized Grouping
- I. Adding Attributes to Avoid Calculations
- J. Consolidating Attributes into Same Dimension
- K. Adding Measures to Avoid Querying the Leaf Level
- L. Creating Dummy Objects in the Cube
- M. Replacing Calculations with Cube Features
- N. Removing the Non_Empty_Behavior Property
- O. Aggregation and Partition Strategy
- P. LAB: Using MDX Calculation Performance Enhancement Strategies

VIII. Using Subselects and Subcubes

- A. Understanding Subselect and Subcubes
- B. Subselect and Slicing
- C. Cache and Performance
- D. Subselect and Subcube Syntax
- E. Referencing Members Excluded by Subselect
- F. Client Applications and Multiple Member Slicing
- G. Writing Calculations for Subselect
- H. LAB: Modifying a Calculation for Subselect