

www.sqlbi.com



Microsoft Partner
Gold Business Intelligence
Gold Data Platform

SSAS
MAESTRO
by Microsoft



Toolkit for DAX Optimization

Marco Russo, SQLBI





Mail: marco@sqlbi.com

Twitter: @marcorus

Marco Russo



- Founder of www.sqlbi.com
- Consultant, Trainer, BI Pro, Book Writer
- SSAS Maestro, MVP, MCP



Agenda

○ Tools

- VertiPaq Analyzer
- SQL Server Profiler
- DAX Studio
- OLAP PivotTable Extensions

○ Techniques

- Know your data model
- Isolate query and measure
- Create repro query

Know your data model first

- Do you know..
 - The size in memory of each column
 - The cardinality of each table
 - The cardinality of each column
 - Hint: they are not the same!
- VertiPaq Analyzer
 - Simplify extraction of this data from DMVs
 - <http://www.sqlbi.com/tools/vertipaq-analyzer>

Demo

VertiPaq Analyzer



Understanding Query Plans

- Logical Query Plan
 - It is the logical flow of the query
 - Fired as standard text
 - Pretty hard to decode
- Physical Query Plan
 - The logical query plan executed by the engine
 - Can be very different from the logical one
 - Uses different operators
- VertiPaq Queries
 - Queries executed by the xVelocity engine

Tabular Two Engines

- Formula Engine (FE)
 - Handles complex expressions
 - Single threaded
- Storage Engine (VertiPaq / xVelocity) (SE)
 - Handles simple expressions
 - Executes queries against the database
 - Multithreaded

SQL Server Profiler



Trace Properties

General Events Selection

Review selected events and event columns to trace. To see a complete list, select the "Show all events" and "Show all columns" options.

Events	EventSubclass	Duration	CPUTime	IntegerData	SPID	TextData
Queries Events						
<input checked="" type="checkbox"/> Query End	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Query Processing						
<input checked="" type="checkbox"/> DAX Query Plan	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
<input checked="" type="checkbox"/> VertiPaq SE Query Cache Match	<input checked="" type="checkbox"/>				<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
<input checked="" type="checkbox"/> VertiPaq SE Query End	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Query End
Query end. Show all events
 Show all columns

EventSubclass (no filters applied)
Event Subclass provides additional information about each event class.

Column Filters...
Organize Columns...

Run Cancel Help

- Catches events from SSAS
 - Queries Events
 - Query Processing

Demo Setup & Flow Analysis

SQL Profiler



Detect slowest queries



- Isolate slowest query
 - A report might generate several queries
 - Use only Query End events in SQL Profiler to locate more expensive ones
 - Create a single repro query
- Isolate slowest report
 - If SQL Profiler is not available, identify slowest query
 - Excel: capture MDX with OLAP PivotTable Extension
 - <http://olappivottableextend.codeplex.com/>

Demo

OLAP PivotTable Extension



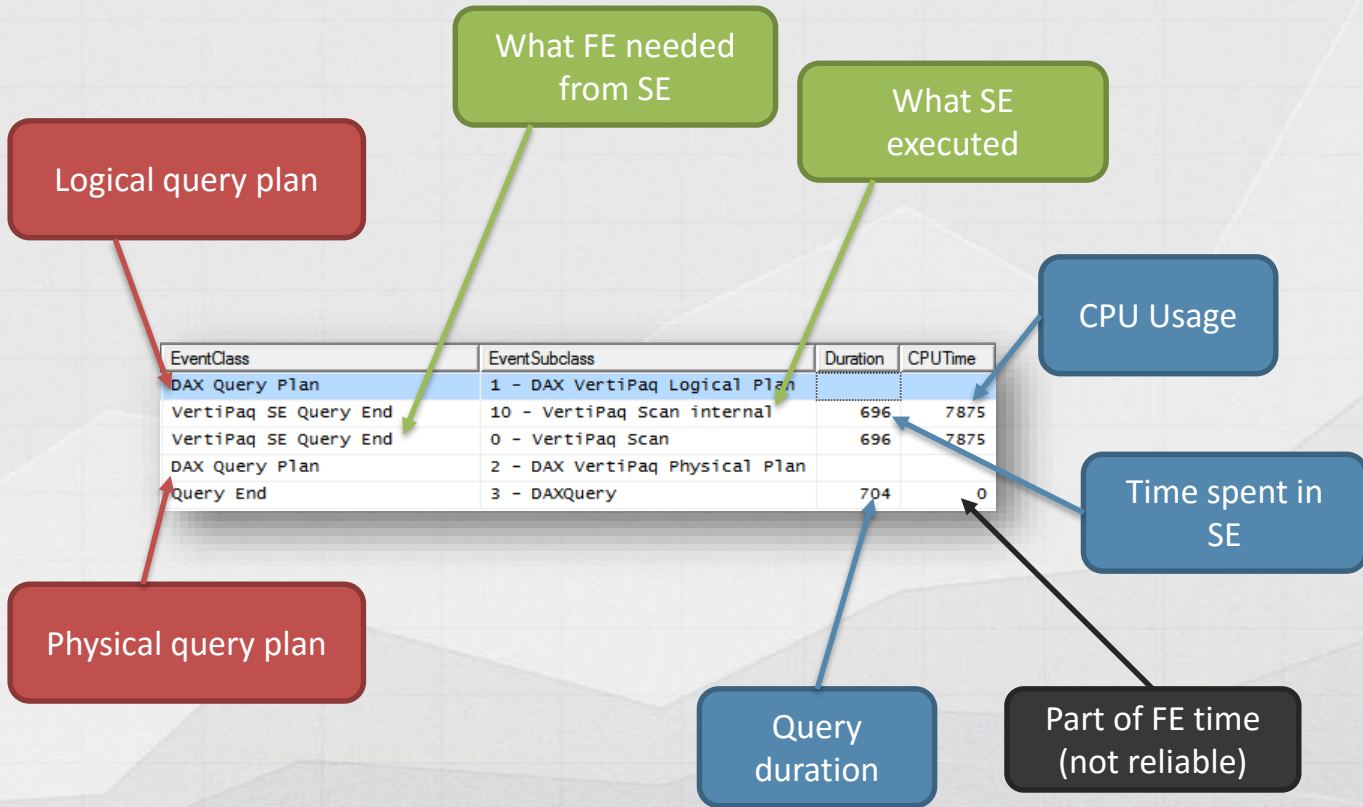
The simplest query



Let's profile a simple query to get acquainted with the DAX query engine

```
EVALUATE  
  
ROW (  
    "Result",  
    SUM ( Sales[SalesQuantity] )  
)
```

Profiler Information



Profiler Information

Formula Engine time computed as
Query Duration – Time spent in SE

704 - 696 = 8 milliseconds

EventClass	EventSubclass	Duration	CPUTime
DAX Query Plan	1 - DAX VertiPaq Logical Plan		
VertiPaq SE Query End	10 - VertiPaq Scan internal	696	7875
VertiPaq SE Query End	0 - VertiPaq Scan	696	7875
DAX Query Plan	2 - DAX VertiPaq Physical Plan		
Query End	3 - DAXQuery	704	0

CPU Usage

Time spent in SE

Query duration

Part of FE time
(not reliable)

Demo Query Plan & SE Events

SQL Profiler



Gathering meaningful numbers

- Values reported by the profiler might not be accurate on a busy server
- Always test on a dedicated server
- If not available, your workstation is better

What is a Data Cache?

Column Storage

Prod	Amt	Qty

Storage Engine (VertiPaq)

VertiPaq query (xMSQL)

Data Cache

Prod	SUM(Amt)	SUM(Qty)
...
...
...
...

Formula Engine

DAX Query Plan

Result

Prod	SUM(Amt)	SUM(Qty)
...
...
...
...

SE fills data cache

FE works on data caches

xmSQL

- Storage engine uses xmSQL
 - Syntax similar to SQL
 - Few options (compared to SQL)
 - Not intended for the user
- Simple mathematical operations
- Some (not all) aggregations
- Usage of model relationships

How FE uses Data Caches

- There is no indication of which data cache is used by which FE operator
- Only common sense and some geeky attitude helps
- For simple plans, it is easy
- For complex plans, it is nearly impossible
- Thus: optimize simple plans!

Cache



- DAX caches results of VertiPaq queries
 - MDX caches measures too
- FE bound queries will not benefit from cache

EventClass	EventSubclass	Duration	CPUTime
DAX Query Plan	1 - DAX VertiPaq Logical Plan		
VertiPaq SE Query End	10 - VertiPaq Scan internal	2015	23047
VertiPaq SE Query End	0 - VertiPaq Scan	2016	23047
VertiPaq SE Query End	10 - VertiPaq Scan internal	1	0
VertiPaq SE Query End	0 - VertiPaq Scan	1	0
DAX Query Plan	2 - DAX VertiPaq Physical Plan		
Query End	3 - DAXQuery	2032	16
DAX Query Plan	1 - DAX VertiPaq Logical Plan		
VertiPaq SE Query Cache M...	0 - VertiPaq Cache exact match		
VertiPaq SE Query End	0 - VertiPaq Scan	0	0
VertiPaq SE Query Cache M...	0 - VertiPaq Cache exact match		
VertiPaq SE Query End	0 - VertiPaq Scan	1	0
DAX Query Plan	2 - DAX VertiPaq Physical Plan		
Query End	3 - DAXQuery	3	0

Clear the Cache



Always remember to clear the cache prior to execute any performance test, otherwise numbers will be contaminated.

With DAX Studio it is a simple button click. With SSMS, you have to execute the following code snippet:

```
<Batch xmlns="http://schemas.microsoft.com/analysiservices/2003/engine">
  <ClearCache>
    <Object>
      <DatabaseID>Adventure Works DW Tabular</DatabaseID>
    </Object>
  </ClearCache>
</Batch>
```

The best all-in-one tool for developing DAX code

DAX Studio



DAX Studio

- <http://daxstudio.codeplex.com>
 - Free add-in for Excel
 - Standalone executable for Tabular and Power BI
- It can query
 - Analysis Services databases
 - Power Pivot data models
 - Power BI Desktop data models
 - Both DAX and MDX queries



Main Features

- DAX Query Editor
 - Full metadata available
 - IntelliSense enabled
 - Format DAX code
 - Through www.daxformatter.com website
- Catch Query Plan and Server Timings
- Save query plan and timings for later analysis

Complete Demo

DAX Studio



Conclusions

Create a repro query

Isolate MDX or DAX query

Identify DAX expression (usually a measure)


Modify DAX expression and test different timings

DAX Studio improves productivity

Use local DAX measure overriding the one defined in the data model

References


- The Definitive Guide to DAX
By Marco Russo, Alberto Ferrari – Microsoft Press
<http://www.sqlbi.com/books/the-definitive-guide-to-dax/>
- Understanding DAX Query Plans
<http://www.sqlbi.com/articles/understanding-dax-query-plans/>
- Understanding Distinct Count in DAX Query Plans
<http://www.sqlbi.com/articles/understanding-distinct-count-in-dax-query-plans/>



Tools

- VertiPaq Analyzer
 - <http://www.sqlbi.com/tools/vertipaq-analyzer>
- OLAP PivotTable Extensions
 - <http://olappivottableextend.codeplex.com/>
- DAX Studio
 - <http://daxstudio.codeplex.com>

Thank you!



Check our articles, whitepapers and courses on

www.sqlbi.com