



# Empowering Business through **Data**

1205	\$9,079	\$2,387	\$3,642	\$1,989	\$5,183	\$4,338	\$7,052	\$11,975	\$13,116
1282	\$5,442	\$8,753	\$9,259	\$3,555	\$1,251	\$6,735	\$8,987	\$13,990	\$15,244
1505	\$9,848	\$8,638	\$9,322	\$3,021	\$1,251	\$6,735	\$8,987	\$13,990	\$15,244
1071	\$5,459	\$8,638	\$4,472	\$3,068	\$1,251	\$6,735	\$8,987	\$13,990	\$15,244
1366	\$0	\$3,310	\$7,267	\$3,003	\$1,251	\$6,735	\$8,987	\$13,990	\$15,244
1053	\$5,573	\$8,617	\$2,162	\$1,989	\$3,642	\$2,467	\$7,951	\$10,143	\$11,119
1024	\$5,543	\$9,424	\$3,905	\$1,251	\$6,735	\$1,237	\$8,123	\$9,513	\$9,463
1556	\$3,589	\$9,424	\$2,110	\$1,251	\$6,735	\$1,237	\$8,123	\$9,513	\$9,463
1710	\$7,317	\$6,329	\$8,143	\$7,338	\$3,476	\$8,118	\$5,790	\$8,215	\$10,799
1389	\$8,413	\$1,571	\$4,472	\$4,338	\$5,4	\$3,072	\$8,959	\$8,190	\$9,477
1389	\$8,413	\$3,398	\$4,472	\$4,338	\$5,4	\$3,072	\$8,959	\$8,190	\$9,477
1389	\$8,451	\$3,398	\$3,000	\$1,989	\$3,642	\$2,467	\$7,951	\$10,143	\$11,119

SQL SERVER 2016

# SSAS Monitoring and Performance Tuning with Extended Events



## **Products**

Making Business  
Intelligent



## **Services**

Consultation,  
Mentoring,  
Solutions



## **Training**

Online Training,  
Webinars, and  
More

**Delora Bradish**

*Senior Consultant*

[dbradish@pragmaticworks.com](mailto:dbradish@pragmaticworks.com)

[www.delorabradish.com](http://www.delorabradish.com)

# Who am I?

[www.delorabradish.com](http://www.delorabradish.com)



**Pragmatic Works Sr. Consultant**

Hands-on lead of MS BI consulting engagements nationwide leading teams in OLAP data modeling, SSIS, SSAS multidimensional, SSAS tabular and Power BI

**Microsoft Certified**

- MCSE Business Intelligence
- MCSA SQL Server 2012
- MCSE Data Platform
- MCITP Database Developer 2008
- MCITP Database Administrator 2008
- MCTS SQL Server 2008 Database Development
- MCTS SQL Server 2008 Implementation & Maintenance



Grand Tetons



Microsoft



Sedona, AZ

1. SSAS Performance Optimization 10K Foot View
2. Profiler vs Extended Events
3. SQL Server Instance QueryLog Properties
4. SQL Server 2016 Profiler & Extended Events Demos
5. Interpreting Extended Events Data
6. Extended Events Integration with SSIS Demo

### Everything not SSAS Query Logging

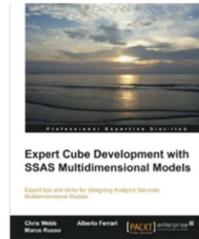
- Page Splits
- Uncommitted Transactions
- Wait Statistics
- Latches
- Error Sessions

<http://www.jasonstrate.com/extended-event-resources/>

<http://www.jasonstrate.com/extended-event-resources/>

## MULTIDIMENSIONAL Performance Tuning 10K Foot View

1. Hierarchies
2. Partitions
3. Aggregations
4. MDX



**Pragmatic  
Works**

Recommended Reading: Expert Cube Development ... by Chris Webb, Alberto Ferrari and Marco Russo (hyperlink in photo)

Partitions: Why do we care?

1. Impacts cube processing speed
2. Impacts end-user query response times
3. Allows you to archive old data

Aggregations: Why do we care?

1. Aggregations are a shortcut to acquiring a SUM() or a COUNT(). Think of aggregations as pre-calculated totals.

2. Correctly designed aggregations = better query performance

Notes:

1. **Hierarchies**

2. **Partitions**

- a. Create at least 4GB equal-size partitions
- b. Align partitions with source system partitions
- c. Match the partitions slice property with the WHERE clause
- d. Sort data for best compression
- e. Only reprocess the current partition

3. **Aggregations**

- a. Create usage-based aggregation designs
- b. Do not exceed 1/3 the size of the partition
- c. Place aggregations on attributes used in hierarchies

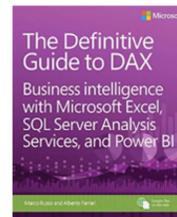
4. **MDX**

1. Exhaust all other options for business logic
2. Test and Optimize!

## TABULAR

# Performance Tuning 10K Foot View

1. Cardinality
2. Partitions
3. Memory Limits
4. Hardware



Memory \ HardMemoryLimit	0	0	0
Memory \ LowMemoryLimit	65	65	65
Memory \ TotalMemoryLimit	80	80	80
Memory \ VertiPaqMemoryLimit	60	60	60

**Pragmatic  
Works**

Recommended Reading: (hyperlinks in photos)

1. Definitive Guide to DAX by Alberto Ferrari and Marco Russo
2. DAX Patterns by Marco Russo, Alberto Ferrari and Daniele Perilli
3. [www.daxpatterns.com](http://www.daxpatterns.com)
4. Performance Tuning of Tabular Models → <https://msdn.microsoft.com/en-us/library/dn393915.aspx>

Performance Tuning Tabular 2012 Models → <https://msdn.microsoft.com/en-us/library/dn393915.aspx>

Please review Rachel Martin's PPTX on optimizing PowerPivot / tabular models. Her slide deck is not copied here.

Original Kasper de Jonge's Server Memory Report for Excel 2013 → <http://www.kasperonbi.com/what-is-using-all-that-memory-on-my-analysis-server-instance/>

See **Extended Events** slide for troubleshooting long running queries.

Please read **chapters 13, 14, 15 & 16** of Alberto Ferrari and Marco Russo's book, The Definitive Guide to DAX: Business Intelligence with Microsoft Excel, SQL Server Analysis Services, and Power BI →  
[https://books.google.com/books/about/The\\_Definitive\\_Guide\\_to\\_DAX.html?id=sJm9CgAAQBAJ&source=kp\\_cover&hl=en](https://books.google.com/books/about/The_Definitive_Guide_to_DAX.html?id=sJm9CgAAQBAJ&source=kp_cover&hl=en)

**Note:** It is useless to have two or more sockets on the same server as SSAS Tabular does not recognize the NUMA architecture which splits memory between different sockets.

From <https://msdn.microsoft.com/en-us/library/ms174514.aspx>

HardMemoryLimit → threshold at which Analysis Services begins rejecting requests outright due to memory pressure.

LowMemoryLimit → IS APPLICABLE TO TABULAR. The MSDN write up is misleading. You have to read the entire article. Set LowMemoryLimit to about 80% OR LESS of the TotalMemoryLimit.

TotalMemoryLimit → An upper threshold at which Analysis Services begins releasing memory more aggressively to make room for requests that are in execution as well as new high priority requests.

VertiPaqMemoryLimit → For tabular instances, a lower threshold at which the server first begins releasing memory allocated to infrequently used objects.

### **Delora's notes on tabular memory settings**

#### **VertiPaqMemoryLimit**

Controls maximum amount of memory that can be used by VertiPaq objects (dictionaries and segments)

If we turn VertiPaq paging on, the VertiPaq memory limit defaults to about 60%

If the VertiPaqMemoryLimit passes above the LowMemoryLimit, SSAS will start cleaning out the cache

If memory usage climbs, cache cleaning will become more aggressive.

If memory passes above TotalMemoryLimit (default 80%), all caches not in use are cleaned out

## Set LowMemoryLimit to about 80% OR LESS of the TotalMemoryLimit

HardMemoryLimit -- if memory exceeds this limit, active sessions will be cancelled in an attempt to get below the HardMemoryLimit;

defaults to 0 (is 1/2 way between the TotalMemoryLimit and server memory limit)

If there are 10GB on the server, and TotalMemoryLimit is set to 8GB, the HardMemoryLimit will set to 9GB

Do not set TotalMemoryLimit close to the physical memory on the machine b/c by the time you trigger the HardMemoryLimit, the gap is too narrow

Slide notes:

1. In tabular models, **cardinality** reigns supreme. Denormalization and parsing columns into multiple columns are two possible solutions to columns with high cardinality (columns containing multiple unique values, like string comments).
2. Set your LowMemoryLimit to 80% OR LESS of the TotalMemoryLimit (which defaults to 80% of total server memory)
3. Hardware priorities
  1. CPU clock speed of 3 GHz or more
  2. Memory speed of 1833, 2133 or 2400 MHz
  3. Number of cores
  4. Memory size used to store the entire DB, execute process operations, and execute queries.

1. Business Logic
2. Query Performance
3. Extended Events
4. Partitions
5. Snapshot Fact Tables
6. Warm the Cache of MultiD Cubes

Effects of clearing the cache for tabular models and MultiD cubes →  
<https://msdn.microsoft.com/en-us/library/hh230974.aspx>

Performance Tuning Tabular 2012 Models → <https://msdn.microsoft.com/en-us/library/dn393915.aspx>

Please review Rachel Martin's PPTX on optimizing PowerPivot / tabular models. Her slide deck is not copied here.

Original Kasper de Jonge's Server Memory Report for Excel 2013 →  
<http://www.kasperonbi.com/what-is-using-all-that-memory-on-my-analysis-server-instance/>

Notes:

1. Push as much **business logic** as possible to the data model. Just because you can calculate it in DAX does not mean you should. Both DAX and MDX should be considered a solution of last resort.
2. End user **query performance** trumps all. You often will incur additional ETL work and longer model processing times to speed query times.

3. Both DAX and MDX can often be written in more than one way and yet achieve the same output. Look at query durations via **Extended Events** and focus on optimizing the long-running queries.
4. Align database **partitions** with model partitions for faster processing.
5. Consider accumulating and periodic **snapshot fact tables** to speed up in-memory report writers
6. Warm the cache after cube processing so the first user of the day does not have to wait for their personal query to warm the cache for everyone.

### MultiD Cache

- Built in the formula engine when evaluating calculations
- Built in the storage engine for results of queries

### Tabular Cache

- Models are stored in memory (or direct query)
- Aggregations and calculations are performed at query time
- Exists as a result of MDX queries

1. Monitor Usage
2. Slow Running Queries
  - Storage Engine
  - Query Engine
3. Warm the Cache of MultiD Cubes

Features Not Supported in the Next Version of SQL Server →

<https://msdn.microsoft.com/en-us/library/ms143729.aspx>

Notes:

1. Monitor user logins, durations and execution count
2. Understand and correct Slow Running Queries
  1. Storage Engine
  2. Query Engine

3. Warm the Cache: this applies to MultiD.  
(You can clear the cache of a tabular model, but there is little point to warm the cache.)

PRAGMATIC WORKS SERVICES  
SQL Server 2016 Profiler

1. Being deprecated by Microsoft
2. The fastest way to log SSAS query activity for immediate analysis
3. Footprint is “heavier” than Extended Events, so not the best solution for scheduled logging
4. In **Event Selection** window, be sure to **Show All Properties** and select the events used by SSAS

---

*Pragmatic*  
**Works**

**Microsoft** | Business Intelligence

Features Not Supported in the Next Version of SQL Server →

<https://msdn.microsoft.com/en-us/library/ms143729.aspx>

1. Not a new SQL 2016 feature. You could use XMLA in earlier SQL versions.

<https://blog.crossjoin.co.uk/2012/05/05/using-xevents-in-ssas-2012/>

2. Graphical user interface in SSMS is a new 2016 feature.
3. Currently you can only add, not edit EE via SSMS

Chris Webb's BI Blog, **Using Xevents in SSAS 2012** →

<https://blog.crossjoin.co.uk/2012/05/05/using-xevents-in-ssas-2012/>

## MULTIDIMENSIONAL Query Logging

1. Designed for use by SSAS multidimensional cube  
usage based aggregations
2. SSAS tabular model activity is not logged here
3. Does not contain an actual SELECT statement

```
Log \ QueryLog \ CreateQueryLogTable
Log \ QueryLog \ QueryLogConnectionString
Log \ QueryLog \ QueryLogSampling
Log \ QueryLog \ QueryLogTableName
```

## TABULAR

# Query Logging

1. As of SQL Server 2016 CU1, it is possible to configure query logging in Properties of the tabular 2016 instance, but tabular query metadata is never written to the specified log table.

Log \ FlightRecorder \ Enabled	true
Log \ QueryLog \ CreateQueryLogTable	true
Log \ QueryLog \ QueryLogConnectionString	Provider=SQLNCLI11.1;...
Log \ QueryLog \ QueryLogSampling	1
Log \ QueryLog \ QueryLogTableName	SSASQueryLog

2. To capture tabular query detail there are at least five options:
  - a. Capture via **Extended Events** (preferred), Import to SQL Server via SSIS and report via SSRS
  - b. DMVs and SSIS
  - c. Activity Viewer from CodePlex
  - d. SQL Profile Trace (not a long-term solution) and SSIS
  - e. ATrace2012 from CodePlex, but not written for SQL 2016 yet

**Pragmatic  
Works**

Please see DTSX file named **Warm the SSAS Cube Cache.dtsx**. This is a SSIS package from another PW client, but it contains tasks for stopping, importing data, and restarting extended events.

Using DMVs to Monitor SSAS → <https://msdn.microsoft.com/en-us/library/hh230820.aspx>

Activity Viewer and ATrace2012 via CodePlex → <http://sqlsrvanalysisrvcs.codeplex.com/>

**Note:** The SSAS logging properties were originally intended for multidimensional cubes with the intent of creating usage-based aggregations from the log. Tabular models do not have pre-processed totals / aggregations, so the use of the query log would only be to see User Ids and query durations.



## 1. Key Multidimensional Extended Events

- Progress Reports \ Progress Report Begin
- Progress Reports \ Progress Report End
- Queries Events \ Query Begin
- **Queries Events \ Query End**
- Query Processing \ Execute MDX Script Begin
- Query Processing \ Execute MDX Script End
- Query Processing \ Get Data From Aggregation
- Query Processing \ Query Cube Begin
- Query Processing \ Query Cube End
- Query Processing \ **Query Subcube**
- **Query Processing \ Query Subcube Verbose**
- Query Processing \ Resource Usage

## 2. Use to trouble-shoot and retest long-running queries

## 3. Be sure to clear your cache!

**Query Duration** = SUM([duration] of Query End event)

**Storage Engine Duration** = SUM([duration] of Query Subcube event)

**Formula Engine Duration** = (Query Duration) – (Storage Engine Duration)



RMC = Right Mouse Click

```
<ClearCache xmlns="http://schemas.microsoft.com/analysiservices/2003/engine">
  <Object>
    <DatabaseID>MySSASdatabaseName</DatabaseID>
  </Object>
</ClearCache>
```

**Query Subcube Verbose**

[Dimension Name] (0,0,0) = there are 3 attributes in the [Dimension Name]

0 = default member was returned

\* (star) = all members were returned for the attribute. Scroll right and the \* (star) should appear by a particular attribute

A specific number = a specific member is being returned (the data ID, not the key value)

+ (plus) = greater than one member was returned

<http://blogs.msdn.com/b/sqlcat/archive/2007/03/05/ssas-partition-slicing.aspx>

<https://thomasivarssonmalmo.wordpress.com/2012/01/28/my-quick-guide-to-ssas-query-tuning-part-one/>



1. Use to **monitor** tabular models
  1. Query duration via **Query End** event
  2. DAX query plans via **Dax Query Plan** event
  3. VertiPaq cache information via **VertiPaq SE Query Cache Match** event
  4. VertiPaq storage engine information via **VertiPaq SE Query End** event
  5. End-user activity, number of requests etc.
2. Use to **understand and retest** long-running queries

**Query Duration** = SUM([duration] of Query End event)

**Storage Engine Duration** = SUM([duration] of VertiPaq Scan event subclass)

**Formula Engine Duration** = ([duration] of Query End event) less  
(Storage Engine Duration)

**Cores used in VertiPaq Operation** =  $\frac{([\text{CPU Time}] \text{ of VertiPaq Scan event subclass})}{([\text{Duration}] \text{ of VertiPaq Scan event subclass})}$

## Steps to Create Extended Event Session

In SSMS, connect to tabular instance

Expand Management

Right mouse click on Sessions --> New Session

On Events page, select events

Category "profiler" are those events that match up to SQL Server Profiler

Category "pureExeevent" are those events that are new to Extended Events

Click the <Configure> button and *for each event*, see the Event fields tab showing the data being returned

Click the back button in the <Select> button

Go to the Data Storage page -- you can have multiple targets

Event File -- XEL file that will be placed on hard disc at a default location of your SSAS log directory

Event string -- live stream of the data

Ring buffer

Retained in memory

Does not go to a file

Only available when the service is up and running. Data goes away when the SSAS service is restarted

Go to Advanced page

Retention mode - how do you want to handle events?

Capture all events?

Loose some events in favor of performance?

Xevents is pretty lightweight vs Profiler Trace

Dispatch latency -- how long do you want to wait until you serve the events from the buffer to the target for processing

Memory size -- the larger the buffer, the longer before we dispatch the data to the target

Partition mode

Memory per CPU

Memory per node

Script to new query editor window

Open XEL file in SSMS

File --> Open --> Merge Extended Events

From the XE Viewer, right mouse click to add columns to the event display window

### SSIS Control Flow

1. Stop Xevent Session
2. Import xEvent File
3. Delete xEvent file
4. Get MDX queries meeting criteria
5. Execute each MDX query
6. Delete old SQL DB xEvent data
7. Start xEvent Session

Monitoring SSAS 2016 Activity with Extended Events

<http://www.delorabradish.com/ssas/monitoring-performance-tuning-troubleshooting-sql-server-analysis-services-ssas-2016-activity-with-extended-events>

Profiler vs Extended Events for SSAS 2016 Database Administration

<http://www.delorabradish.com/dba/sql-server-profiler-2016-vs-extended-events-for-analysis-services-ssas-database-administration>

Performance Tuning of Tabular Models

<https://msdn.microsoft.com/en-us/library/dn393915.aspx>

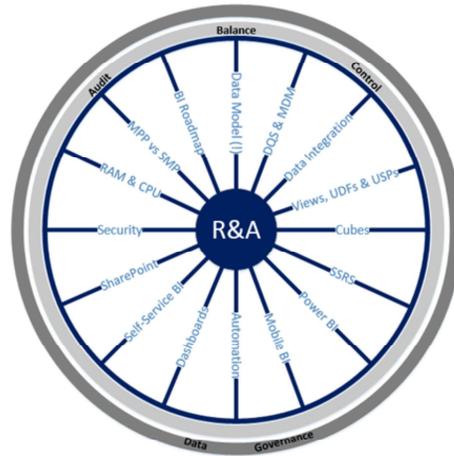
Reference of Profiler event numbers and names

<https://blogs.msdn.microsoft.com/danhardan/2008/09/09/sql-profiler-eventclass-eventsubclass-column-descriptions/>

**Pragmatic  
Works**

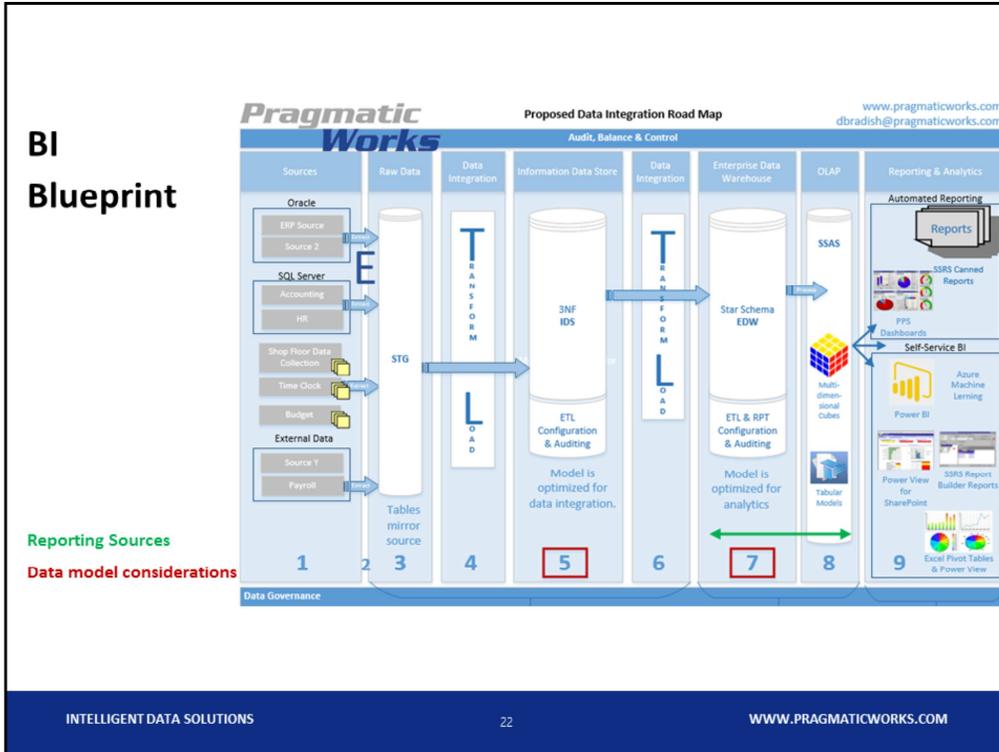
**Questions?**

## Purpose of BI = Reporting & Analytics

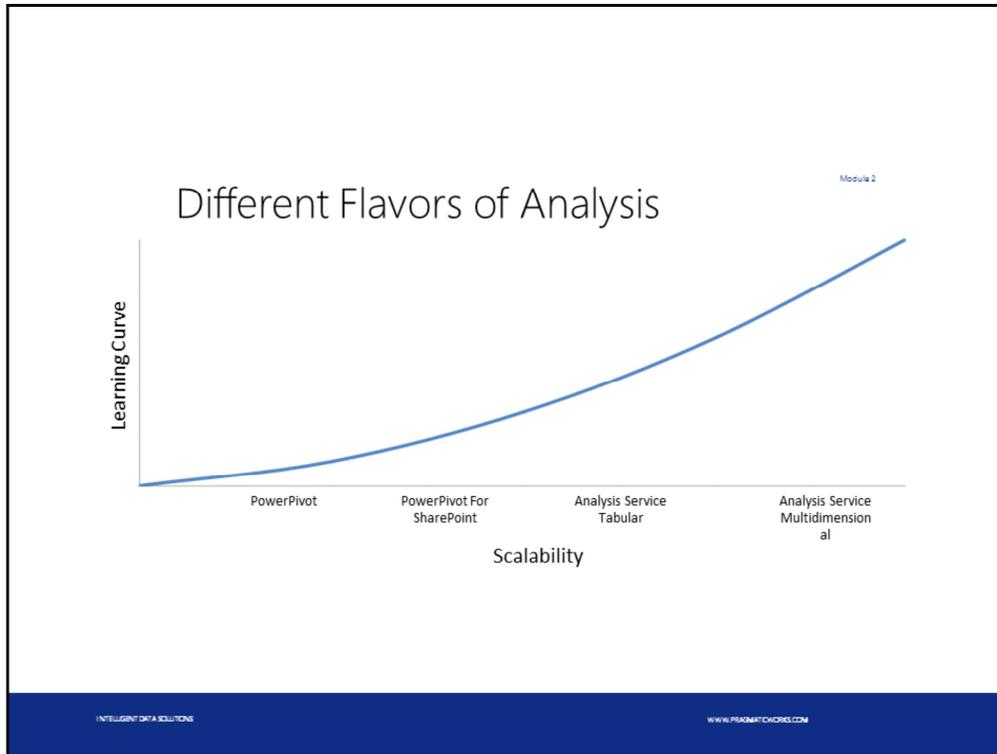


### Talking Points:

1. If you have not defined a single report, how will your end result be optimized for reporting?
2. How do we frequently sacrifice the permanent, R&A, on the alter of the immediate, a faster ETL?



Today's focus will be on modeling for pipes #5 and #7 (the red boxes) understanding that the plan is for all analytics to source from pipes #7 and #8 (the green arrow)



Notes from online course:

Different flavors of SSAS have different best use cases

Y axis = difficulty to learn, time to develop

X axis = scalability

Different flavors of SSA include

1. Powerpivot on your personal machine
2. Powerpivot for SharePoint which is a PP model deployed out to SharePoint environment that uses server resources for queries
3. SSAS tabular
4. SSAS multidimensional

PP for personal BI

- Least scalable
- Data size restrictions
- Limited to personal laptop
- Little learning curve, little complexity

### PP for SharePoint

- Same as PP on your personal machine, but is now PP deployed out to your SP environment
- Database is stored in Analysis Services running behind the scenes
- Using server resources
- More people can see it and use it
- A team BI solution

### SSAS tabular

- Available in SQL Server 2012
- Column compression
- Has features not in PP, role based security, partitioning
- Used in place of PP for SharePoint, but does not need SP
- Allows you to take your PP model right into SSAS via wizard

### SSAS Multi

- Most scalable, but also the most complex
- Take the most time to develop
- Requires a lot more expertise
- Requires knowledge of the DW, partitioning strategies, aggregations
- Steepest learning curve but the most scalable (TB data sources with 50 + GB cube sizes)