

SSAS Tabular in the Real World

Lessons Learned



by Gerhard Brueckl



#419 | BRATISLAVA 2015

Gold sponsors



Platinum sponsor



About me

Gerhard Brueckl

From Austria 

Consultant, Trainer, Speaker

Working with Microsoft BI since 2006

Mainly focused on Analytic Databases

Analysis Services / SAP HANA

Power BI / O365 / Azure



SSAS
MAESTRO
by Microsoft

pmOne

<http://www.pmone.com>



@gbrueckl



blog.gbrueckl.at



gerhard@gbrueckl.at



#419 | BRATISLAVA 2015

Agenda

- Defining Requirements
- Choose Technology
- Design your Model
- Build your Model
- Deploy your Model
- Maintain your Model

Define Requirements

- Usually done by Business
- Not much we can do here
- Should get involved as soon as possible
 - Understand Business Requirements
 - Pre-Validation
 - Share your Experience

Define Requirements

Prototyping

- Easy and Fast Prototyping
 - In Excel and Power BI
 - (No additional Costs)
- “Quick Wins”
- Show the Business Value

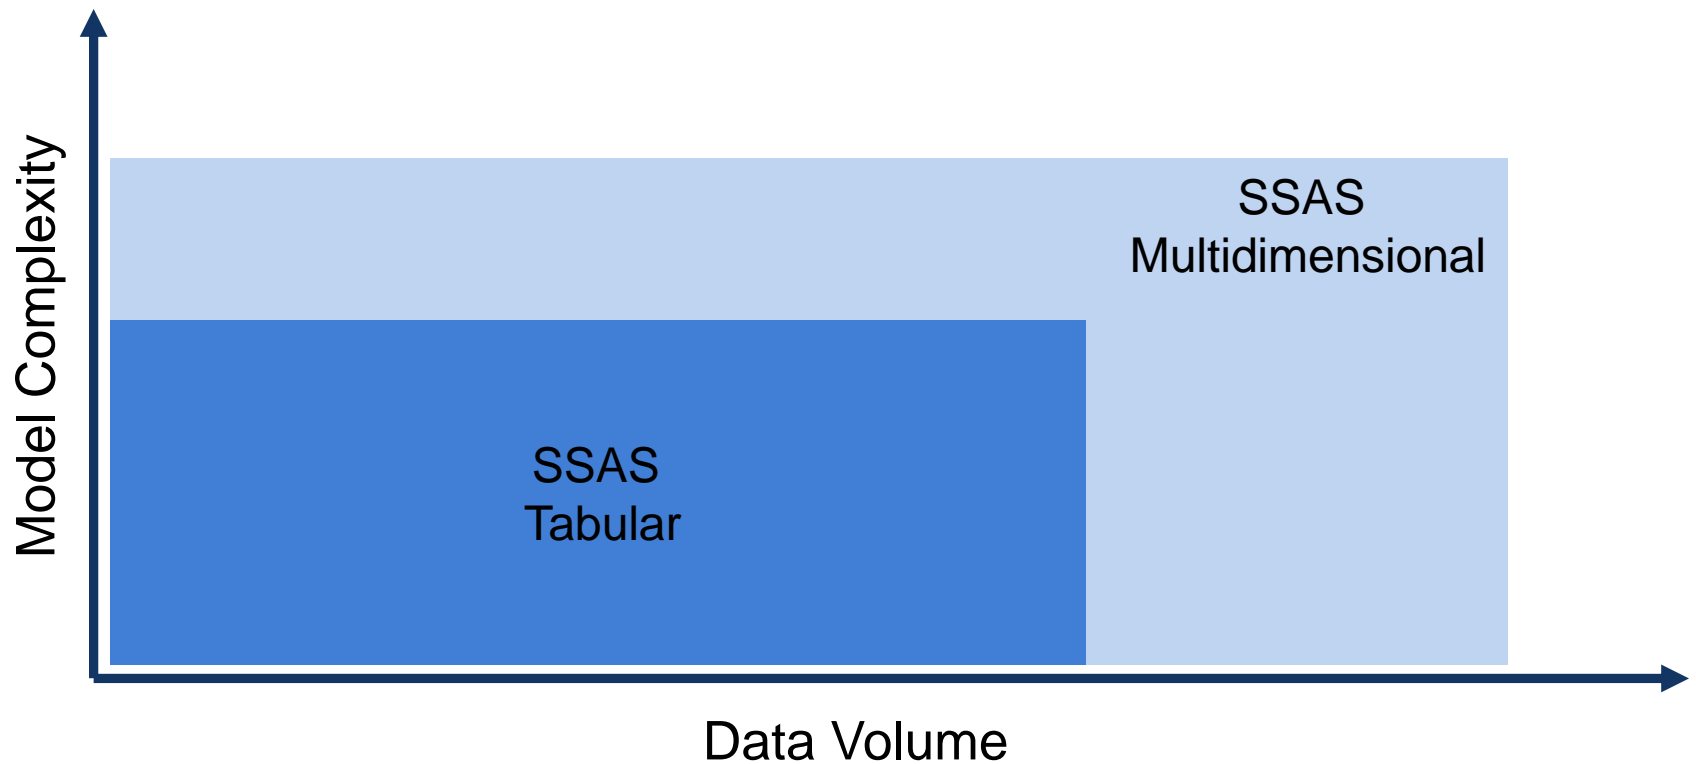
Choose Technology



Choose Technology

- When to choose Tabular
 - When NOT to choose Tabular (!)
- Data Volume
- Model Complexity
- Update Frequency
- Skill-Set

Choose Technology



Choose Technology

General

	Tabular	Multidimensional
Since	SQL 2012	1998 – OLAP Services 7.0 SQL 2005
Storage	Column Store In-Memory	MOLAP Disk
Compressrate	~ 10x	~ 3x
Design Pattern	Flexible Relationships Tables	Star Schema Facts vs. Dimensions
Modelling Features	Simple	Complex
Definition Language	DAX	MDX
Query Language	DAX and MDX	DAX and MDX

Choose Technology

Modelling Features

	Tabular	Multidimensional
Time Intelligence	Using DAX	Using MDX
User Hierarchies	Native	Native
- Leveled	Native	Native
- Parent-Child	Using DAX	Native
Many-to-Many	Using DAX	Native
Unary Operators	Using DAX	Native
Semi-Additive Measures	Using DAX	Native
Writeback	No	Yes
Role-Playing Dimensions	No	Yes

Choose Technology

Design Features

	Tabular	Multidimensional
Perspectives	Yes	Yes
Translations	Yes * (Metadata only)	Yes
Display Folders	Yes *	Yes
Actions	Yes *	Yes
Calculated Columns	Yes	Yes (in DSV)
Conditional Formatting	No	Yes
Generic Calculations	No	MDX Script

* Using BIDS Helper

Choose Technology

Why should I choose SSAS Tabular at all?!

- Short Development Lifecycle
 - Interactive Development
- Fast Performance
 - No Tuning required
- Fast-Changing Requirements

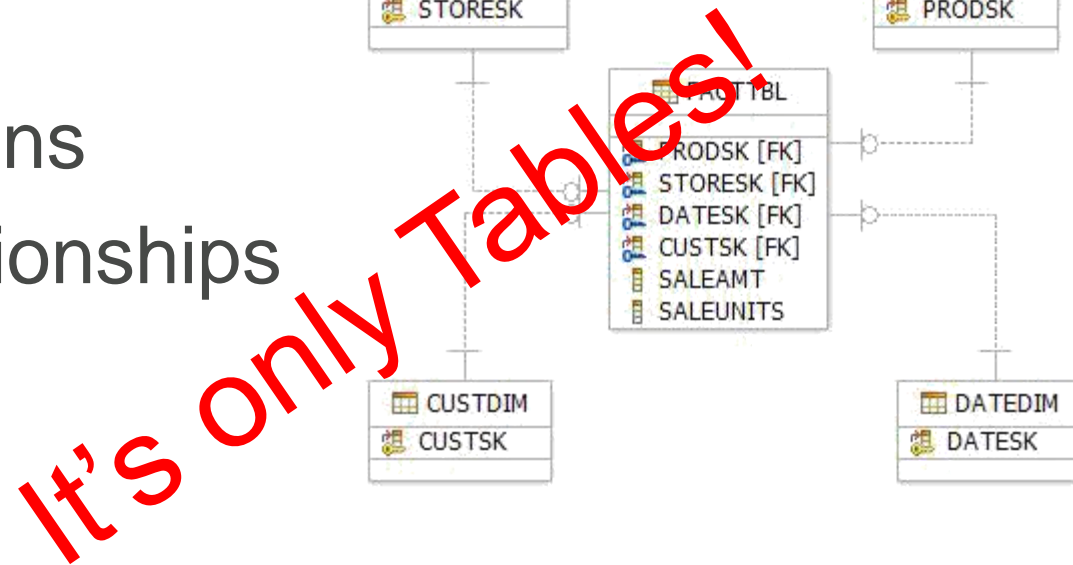
Design your Model



Design your Model

Classic Approach

- Facts
- Dimensions
- 1:n Relationships



Design your Model

- No need for separate Fact Tables
 - Fact ↔ Dimension
 - Performance
 - Query Time
 - Processing Time
- Line-Item Reporting

Cosmetics

- Dimension Name = MeasureGroup Name

Design your Model

Role-Playing Dimensions

Not natively supported in SSAS Tabular

- Active / Inactive Relationships
 - one Measure for each Relationship
- Import same Table multiple times
 - + flexible Analysis
 - Memory Consumption

Design your Model

Many-to-Many

Not natively supported in SSAS Tabular

- Workarounds using DAX
- Based on multiple One-to-Many Relationships

How about SQL 2016?

Design your Model

Many-to-Many

SQL 2016

- Only supports One-to-Many Relationships

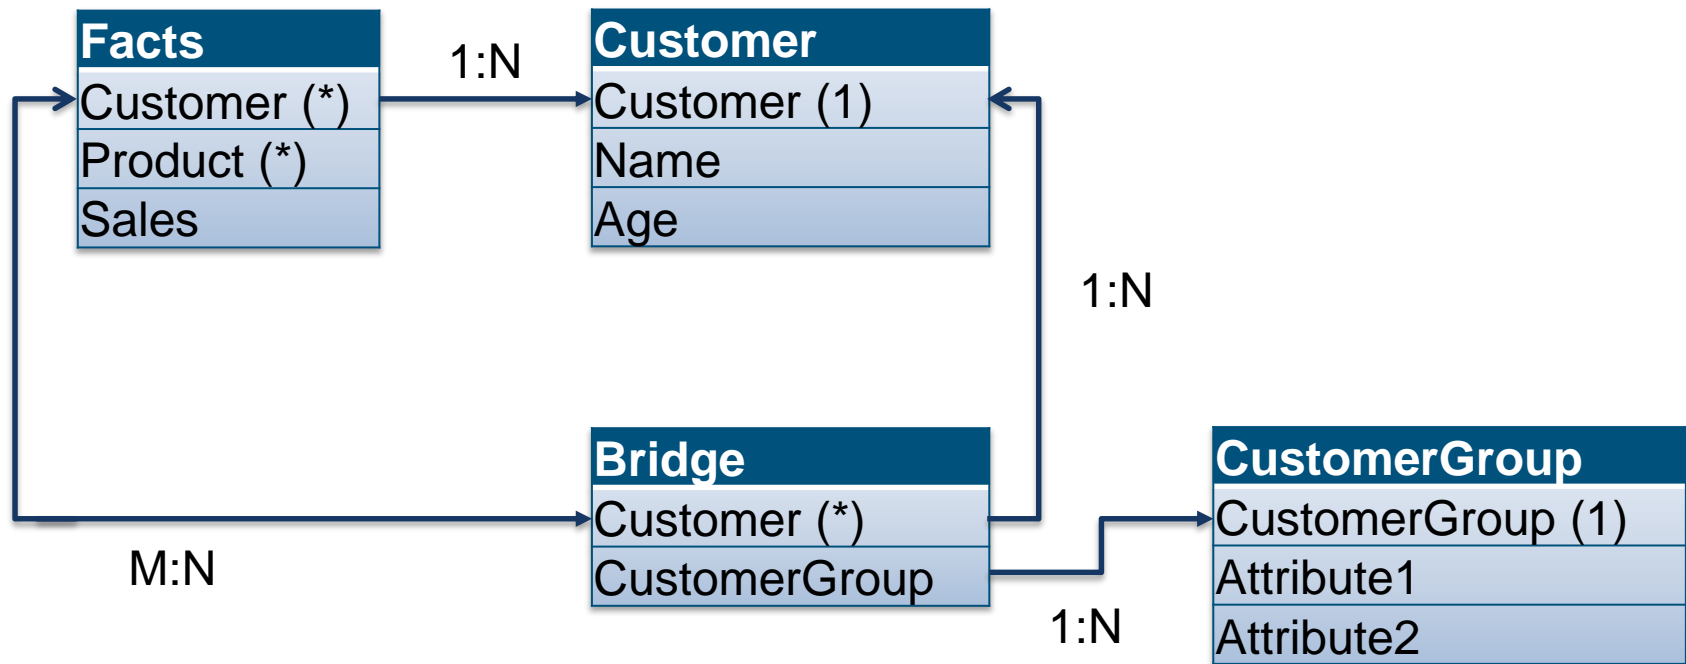
BUT

- Relationships can be bi-directional
- Filters flow in both directions

Design your Model

Many-to-Many

Classic Many-to-Many



Real Many-to-Many

Design your Model


The Problem of Distinct Count

- Special kind of Measure
- Not Additive
- Best Practice in SSAS MD
 - Use separate MeasureGroup

Design your Model

The Problem of Distinct Count

Customer	Product	Date	Sales
Peter	Beer	2015-01-07	20
John	Gin	2015-02-03	100
John	Beer	2015-02-03	30
Peter	Beer	2015-02-05	40



Product	Sales
Beer	90
Gin	100
Total	190

SUM ↓



Product	DC Customers
Beer	2
Gin	1
Total	??

Design your Model

The Problem of Distinct Count

Customer	Product	Date	Sales
Peter	Beer	2015-01-07	20
John	Gin	2015-02-03	100
John	Beer	2015-02-03	30
Peter	Beer	2015-02-05	40

Product	Customer	DC Customers
Beer	Peter	1
Beer	John	1
Gin	John	1
Total	{Peter, John}	2

Sales
60
30
100
190

Distinct

Count

Design your Model

The Problem of Distinct Count

DC in SSAS Tabular

- No Aggregations
 - No separate Table/MeasureGroup
 - DC on “Dimensions”
- Columnar Storage
 - Internally works with Distinct Values
 - Independent of DataType

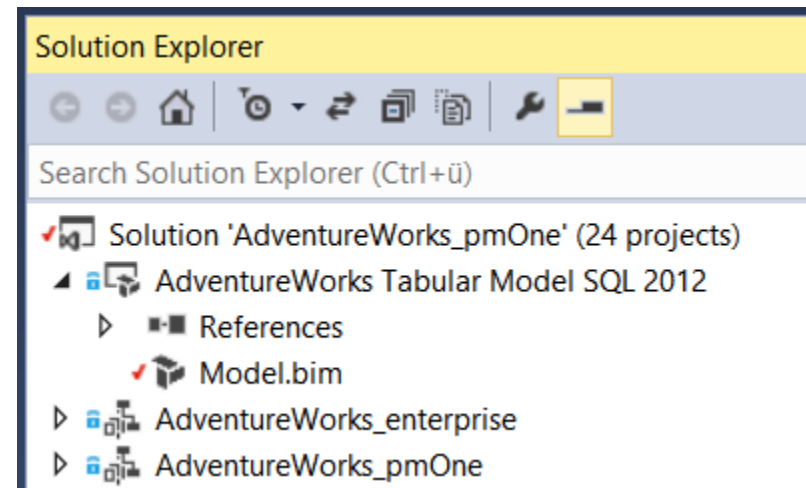
Build your Model



Build your Model

General Development

- SQL Server Data Tools (SSDT)
= Visual Studio
- Native TFS Support
 - Source Control
 - Bug Tracking
 - ...
- **BUT** its only one single File!



Build your Model

General Development

Multiple Developers

- Exclusive Locks
- BISM Normalizer
<http://bism-normalizer.com>
(20 day free trial)

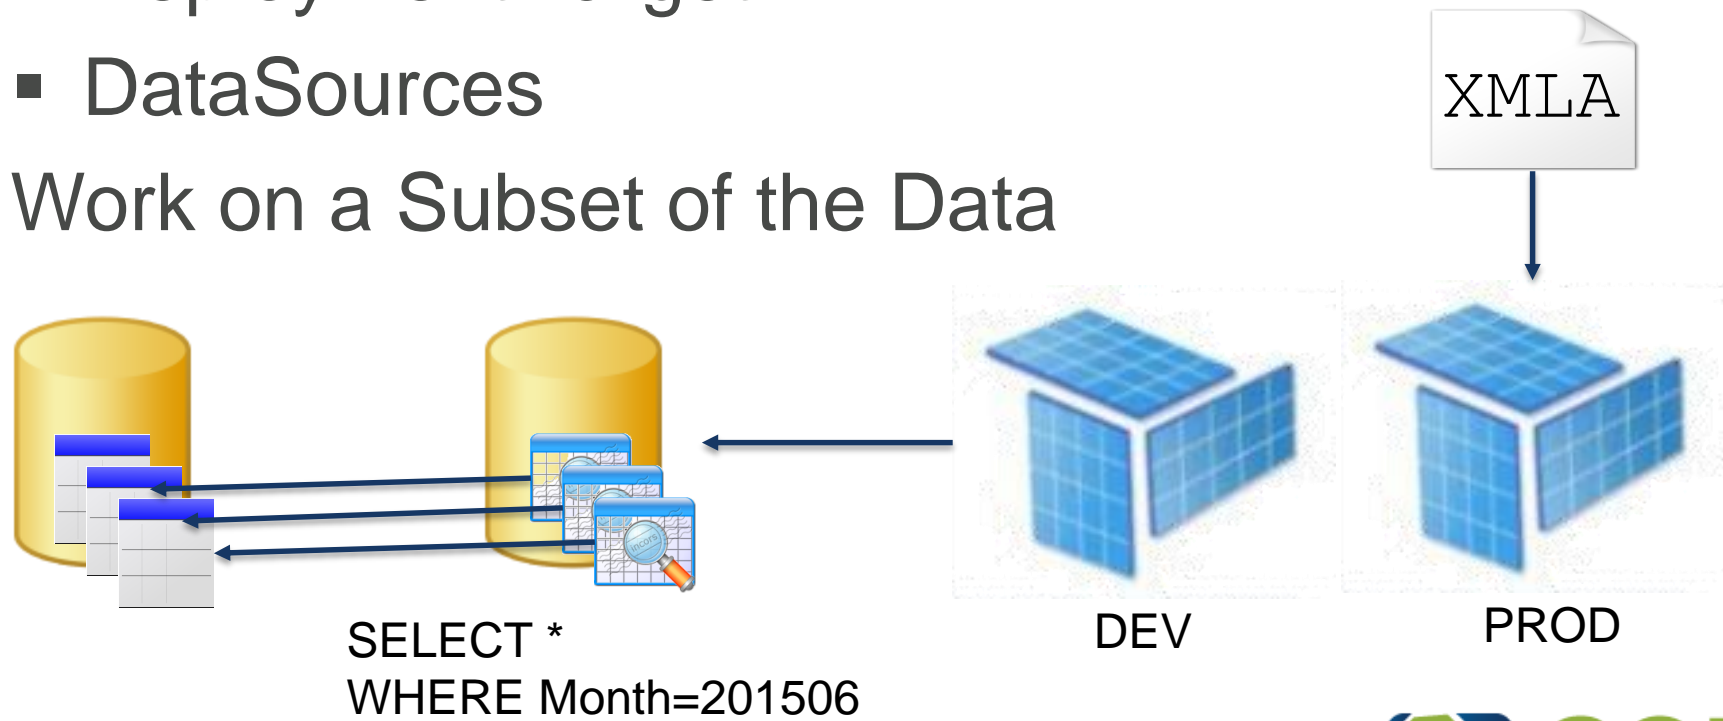
Build your Model

General Development

VS Configurations not working

- Deployment Target
- DataSources

Work on a Subset of the Data



Build your Model Business Logic

- Need to know DAX!
 - Nothing similar to MDX Script
 - No Generic Calculations
 - Only Measures
- One Measure for each Requirement

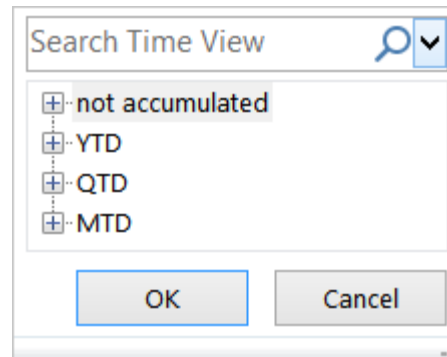
Build your Model Business Logic

8 Base Measures

- Internet Sales
 - Growth in Customer Base
 - Internet Average Sales Amount
 - Internet Extended Amount
 - Internet Freight Cost
 - Internet Gross Profit
 - Internet Sales Amount
 - Internet Standard Product Cost
 - Internet Tax Amount

X

4 Time Views



X

Time Comparison
(CY, PY, Growth)

X

Scenario
(Actual, Plan, FC)

X

Aggregation
(Sum, LNE, ...)

DAX Measures: 864

Build your Model Business Logic

So we have 200+ Measures ...

- Development
 - Slows down VS
- Maintenance
 - Change Requests
 - Nest your Measures!
- (Usability)

Build your Model

Perspectives

- For Usability of the End User
- Can also be used to ease Development
- New Measures/Columns need to be added manually each time!

Build your Model Power BI

Story by Microsoft:

- Build in Power BI (Personal/Team BI) –
Deploy to SSAS (Enterprise BI)

- Not that easy
 - Data Loading Patterns
 - Validate Business Logic
 - Performance

Build your Model

Other Issues

- IntelliSense / AutoComplete
- Formattings
- Timeout for Table-Changes
- ...

Deploy your Solution



Deploy Solution

Same as for Multidimensional

- Deploy via VS
- Deployment Wizard
- XMLA

- BISM Normalizer

Maintain your Solution



Maintain Solution

Same as for Multidimensional

- SSMS
 - SSIS
 - (XMLA)
-
- No Live-View in VS

Maintain Solution Processing

Parallel Processing

- Supported, but for Tables only
 - Not for Partitions (!)
- Supported in SQL 2016 !!!

Maintain Solution Processing

“Tables” vs. Dimensions/MeasureGroups

- Same XMLA
- Dimension=MeasureGroup
- Partitions via MeasureGroup

Maintain Solution Processing Options

ProcessRecalc

- On Database-Level only
- Recalculates
 - Hierarchies
 - Relationships
 - Calculated Columns

- Is done after each ProcessFull/Add/Default (!)

Maintain Solution Processing Options

ProcessData

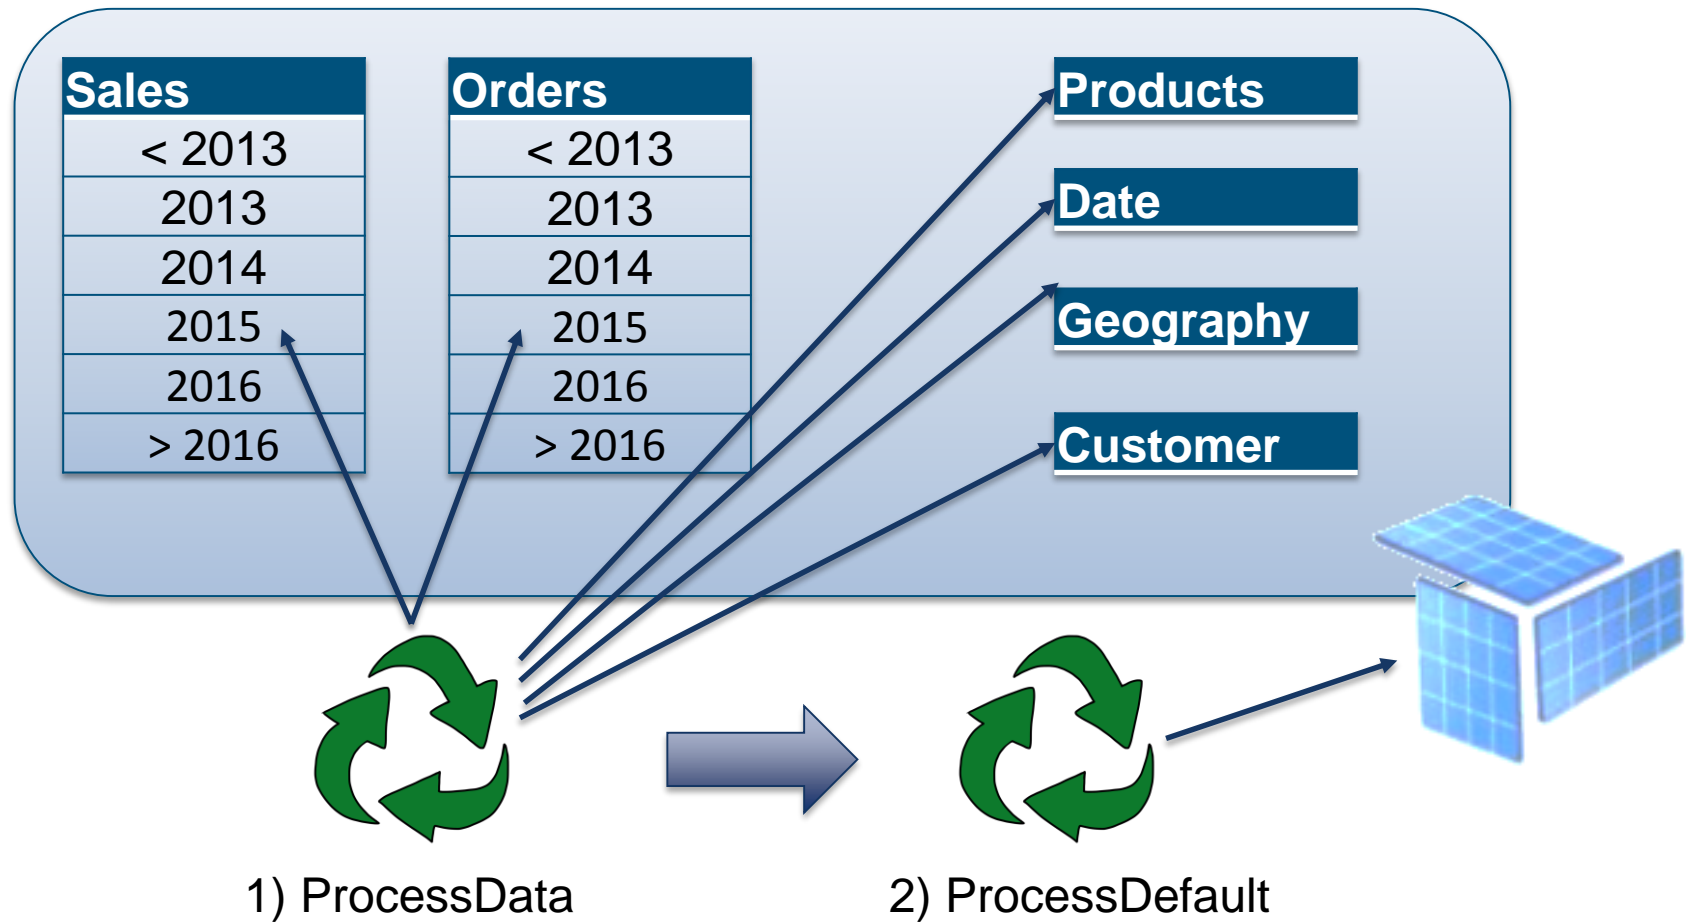
- On Table or Partition
- No automatic ProcessRecalc (!)
 - Invalidates
 - Calculated Columns
 - Relationships
 - Hierarchies

Maintain Solution Processing Options

ProcessData

- On Table or Partition
- No automatic ProcessRecalc (!)
 - Invalidates
 - Calculated Columns
 - Relationships
 - Hierarchies

Maintain Solution Processing Options



Maintain Solution Monitoring

Memory, Memory, Memory

- Consider at least 2-3 times the amount of Memory for Processing
- Avoid Paging at all Costs!

Maintain Solution Monitoring

Same as for Multidimensional

- PerfMon Counters
- SQL Server Profiler (deprecated!)
- Extended Events
- DMVs

- No Query Log

Lessons Learned

- Strengths and Weaknesses of SSAS Tabular
- Impacts on your Project
- Deal with certain issues
- Have fun!