



sqlbi.com

## Analysis Services Best Practices

presented by  
Marco Russo  
[marco@sqlbi.com](mailto:marco@sqlbi.com)

sqlbi.com



## Who am I

- o BI Expert and Consultant
- o Problem Solving
- o Complex Project Assistance
- o DataWarehouse Assesments and Development
- o Courses, Trainings and Workshops
- o Microsoft Business Intelligence Partner
- o Book Writer



## Latest conferences

- o PASS Europe 2009 – Neuss – Germany
- o PASS 2009 – Seattle – USA
- o SQL Conference 2010 – Milan – Italy
- o Teched 2010 – New Orleans – USA
- o 24 Hours of PASS 2010 – Online
- o PASS 2010 – Seattle – USA

sqlbi.com



## Agenda

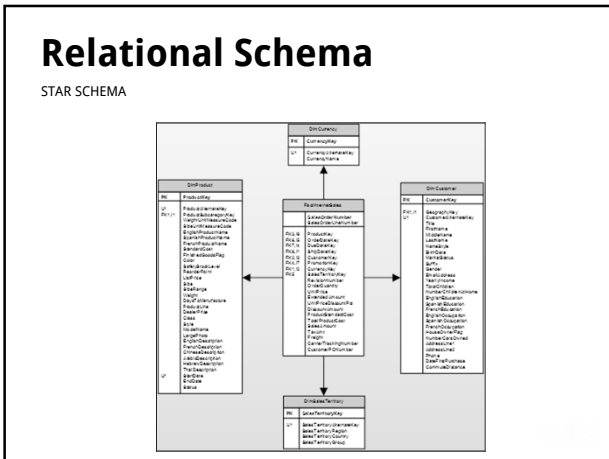
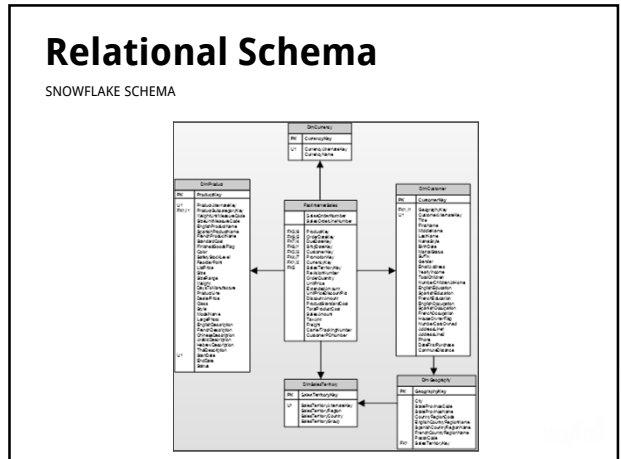
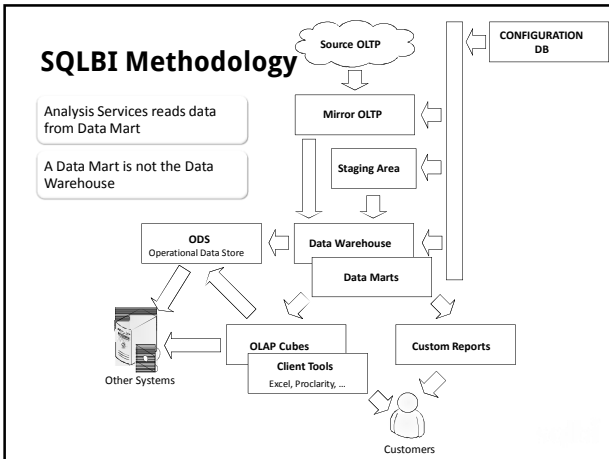
- o Relational Schema
- o Decoupling Layer
- o Dimensional Patterns
- o Slowly Changing Dimensions
- o Junk Dimensions
- o Parent-Child Hierarchies
- o Role Dimensions
- o Drill-through
- o Calculation Dimensions

DATA SOURCE (RELATIONAL MODELING)

## Relational Schema

sqlbi.com



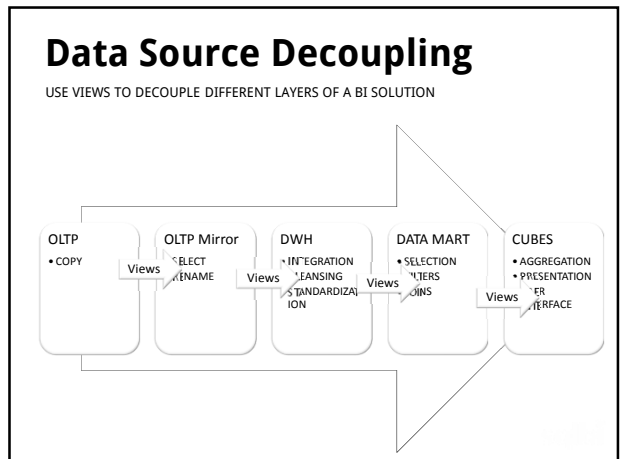


- ### Relational Schema
- STAR VS. SNOWFLAKE SCHEMA
- Options for dimensions from snowflake schema:
    - Transform into a star schema by using views
    - Transform into a star schema by using DWV queries
    - Join tables in SSAS dimensions
    - Referenced Dimension
  - Ideal solution
    - Use SQL views to generate a star schema
    - The star schema eliminates ambiguity

DATA SOURCE (RELATIONAL MODELING)

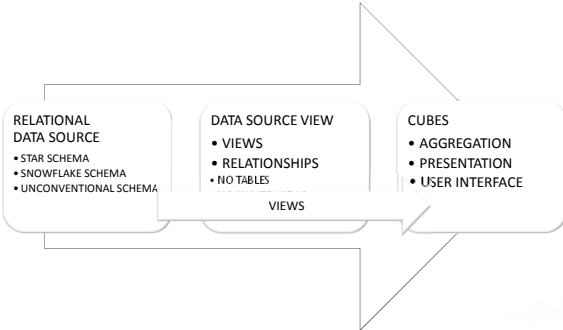
## Decoupling Layer

sqlbi.com



## Data Source Decoupling

DECOUPLE SSAS MODEL - RDBMS STRUCTURE



## Data Source Decoupling

DEFINE VIEWS FOR SSAS CUBES

### Model

- One view for each dimension
- Namespace = cube name (Common for shared dimensions, if necessary)
- Convert a snowflake in a star schema for SSAS

### Columns

- Number of view's columns = number of attributes of dimension
- Column name = attribute name
  - Use a naming convention for attribute keys
- Simple calculations or string operations in views
  - Avoid calculated column in DSV
- Default values when necessary
  - Eliminates NULL (i.e. in case of a LEFT JOIN)

## Data Source Decoupling

WHY TO USE VIEWS?

Decoupling interface	Stored in DB	Textual
Easy to change	Database objects	Can be optimized

SURROGATE KEY, ATTRIBUTE KEY  
GROUPING, BANDING

## Dimensional Patterns

sqlbi.com



## Surrogate Keys

WHY YOU SHOULDNT EXPOSE SURROGATE KEYS

### Benefits

- Independency from application keys
- SCD modeling
- Star schema optimization

### Drawbacks with Analysis Services

- MDX query with surrogate keys
  - Don't have a semantic value
  - Are invalidated if the cube is reprocessed (i.e. nightly one-shot processing)
- Impact on custom security and reporting

## Surrogate Keys

PATTERN FOR NOT EXPOSING SURROGATE KEYS IN SSAS


### Dimensione Customer

- `_CustomerKey`
- KeyColumns: ID\_Customer (surrogate key)
- NameColumn: (none)
- AttributeHierarchyVisible: False
- AttributeRelationship
  - Attribute: Customer
  - Cardinality: One
- Customer
  - KeyColumns: COD\_Customer (business key)
  - NameColumn: CustomerName
  - AttributeRelationship
    - SCD Type I – all other attributes are related to this one
    - SCD Type II – evaluate on a case-by-case basis (impact on KeyColumns)

## Surrogate Keys

PATTERN ATTRIBUTE HIERARCHY FOR SCD TYPE I

**Legend**  
 AttributeHierarchyVisible = True  
 AttributeHierarchyVisible = False




Hierarchies

- Country / City / Customer
- Category / Customer
- Area / Customer

## Surrogate Keys

PATTERN ATTRIBUTE HIERARCHY FOR SCD TYPE II

**Legend**  
 AttributeHierarchyVisible = True  
 AttributeHierarchyVisible = False



Hierarchies

- Country / City / Customer (*Customer-City*)
- Category / Customer (*Customer-Category*)
- Area / Customer (*Customer-Area*)

## Attribute Keys

PATTERN FOR CREATING NATURAL HIERARCHIES FROM CUSTOM HIERARCHIES

Use unique key for attribute

- i.e. Subcategory code if it is unique
- Don't create one in ETL if it doesn't exist (use composite key instead)

If it doesn't exist, use attribute with composite key

- Set of keys corresponding to attribute hierarchy
- Set AttributeHierarchyVisible = False


Use decoupling views to generate columns for appropriate descriptions

- Natural hierarchy: you may reference the parent's name
- Browsible attribute: use a simple description
- i.e. «February 2010» hierarchy - «February» attribute

## Attribute Keys

PATTERN FOR PRODUCT DIMENSION

**Legend**  
 AttributeHierarchyVisible = True  
 AttributeHierarchyVisible = False



Hierarchies

- Color / Category (*CategoryColor*) / Product
- Category / Color (*ColorCategory*) / Product

CategoryColor

- KeyColumns: COD\_Category, COD\_Color
- NameColumn: Category


ColorCategory

- KeyColumns: COD\_Color, COD\_Category
- NameColumn: Color

## Attribute Keys

PATTERN FOR DATE DIMENSION

**Legend**  
 AttributeHierarchyVisible = True  
 AttributeHierarchyVisible = False



Hierarchies

- Year / Quarter (*QuarterYear*) / Month (*MonthYear*) / Day (*DayYearMonthly*)
- Year / Month (*MonthYear*) / Day (*DayYearMonthly*)
- Year (*YearWeekly*) / Week (*WeekYear*) / Day (*DayYearWeekly*)

## Grouping

ATTRIBUTE GROUPING - AUTOMATIC VS. MANUAL

Automatic grouping

- DiscretizationMethod
  - EqualAreas (same number of elements)
  - Clusters (Data Mining algorithm)
  - Automatic
- DiscretizationBucketCount
  - Number of groups

Manual grouping

- Business logic in view for dimension on Data Mart

```

CASE WHEN weight IS NULL
      OR weight < 0 THEN 'N/A'
      WHEN weight < 10 THEN '0-10kg'
      WHEN weight < 20 THEN '10-20kg'
      ELSE '20kg or more'
END
  
```

## Banding

CATEGORIZE A MEASURE ACCORDING TO RANGES OF VALUES

<p><b>Direct relationship</b></p> <ul style="list-style-type: none"> <li>• Range granularity is the same as dimension granularity</li> <li>• Impact of range change             <ul style="list-style-type: none"> <li>• Fact table ETL</li> <li>• Dimension ETL</li> <li>• Reprocess Dimension</li> <li>• Reprocess Cube</li> </ul> </li> </ul>	<p><b>Indirect relationship</b></p> <ul style="list-style-type: none"> <li>• Discretization and constant granularity</li> <li>• Impact of range change             <ul style="list-style-type: none"> <li>• Dimension ETL</li> <li>• Reprocess Dimension</li> </ul> </li> </ul>

SSAS DESIGN FOR SCD TYPE I AND TYPE II

## Slowly Changing Dimensions in SSAS

sqlbi.com



## Type I SCD

ISSUES WITH PROCESS UPDATE

### Process Update

- Doesn't detect duplicate key errors
- Process Full would fail according with error handling settings
- Flexible aggregations are regenerated

### Existing MDX Queries

- Unique Name may change
- MDXMissingMemberMode – Ignore (Default)
- Missing elements are ignored
- Rows/columns disappear without a warning
- MDXMissingMemberMode – Error
- Need to change the query
- Some clients require to rebuild the query from scratch

## Type II SCD

SSAS MODELING

### Process Dimension

- Process Add instead of Process Update
- Attribute relationship: always Rigid

### SSAS Dimension Modeling

- There is no specific dimension type for SCD
- Surrogate key → CustomerKey
  - CustomerSCD attribute – (Primary key attribute)
- Business key → CustomerAlternateKey
  - Customer attribute

## Type II SCD

ATTRIBUTE RELATIONSHIP AND HIERARCHIES

<p><b>“Standard” Modeling</b></p> <ul style="list-style-type: none"> <li>• Only «visible» attributes</li> <li>• CustomerSCD should be invisible             <ul style="list-style-type: none"> <li>• Otherwise may be confusing for end user</li> </ul> </li> <li>• Country-City-Customer hierarchy is not natural</li> </ul>	<p><b>“Ideal” Modeling</b></p> <ul style="list-style-type: none"> <li>• Invisible attribute “Customer in City”</li> <li>• Composite key City + CustomerAlternateKey</li> <li>• Used only in hierarchy Country-City-Customer at Customer level</li> </ul>

CONSOLIDATE INDEPENDENT ATTRIBUTES

## Junk dimension

sqlbi.com



## Junk Dimension

MODELING

Transaction Flag Key	
→	Revision Number
→	Transaction Type
→	Valid Flag

### Junk dimension in relational modeling

- Ideal solution for performance and consumed space
- New attribute requires reprocess of fact table
- Unless a valid default exists for existing data

### Junk dimension only on cube

- Create dimension based on a view
- Use CROSS JOIN
- Creates also combinations of values that are never used
- A new attribute always require cube reprocess

## Junk Dimension

PERFORMANCE AND SEMANTIC

Transaction Flag Key	
→	Revision Number
→	Transaction Type
→	Valid Flag

### Reduce aggregation complexity

- Fewer dimensions
- Same number of attributes, but optimization on AttributeKey more efficient
- Better performance in both query and process

### Usability

- Improves if dimension name is intuitive
- Use more than one junk dimension if a single name is not good
- Set Attribute Key hierarchy to invisible
- Carefully choose names
  - Otherwise you might invalidate existing queries by changing them later

LIMITATIONS OF STANDARD PARENT-CHILD HIERARCHIES  
MODELING MULTIPLE HIERARCHIES

## Parent-child Hierarchies

sqlbi.com



## Parent-Child Hierarchies

LIMITATIONS

### Functional limits

- Only one parent-child hierarchy per dimension
- Must use the dimension's Attribute Key
- Unique name for a member includes surrogate keys

### Performance issues

- Increase processing time
- MDX calculation is more complex (especially SCOPE, autoexists e CurrentMember)
- Slower query response (Good up to some thousand members)

### Alternative

- Parent-child naturalizer in BIDS Helper

## Parent-Child Hierarchies

STANDARD PARENT-CHILD HIERARCHIES IN ANALYSIS SERVICES

### Pros

- Internal optimization is good enough (if not too many members)
- Easy to create and maintain for SSAS developer


### Cons

- Rigid structure
- Only one parent-child hierarchy per dimension
- User cannot add custom parent-child hierarchies
- Defining a separate dimension for each parent-child hierarchy (made with the same leaf nodes) increases the number of aggregations

## Parent-Child Hierarchies

STANDARD PARENT-CHILD VS MULTIPLE HIERARCHIES

- A multiple hierarchy IS NOT a parent-child with just some more nodes
  - Parent-child data is partitioned
  - Multiple hierarchy data are always the same but with different views and might appear as duplicated



## Parent-Child Hierarchies

MULTIPLE HIERARCHIES

### Pros

- Flexible structure (data-driven approach)
- A single dimension defines many hierarchies by using the same data
- User can define new hierarchies without changing the cube structure
- A single dimension to aggregate is better
- Easy to use? It depends...


### Cons

- Optimization is in charge of the SSAS developer
- They are complex to create and maintain


## Parent-Child Hierarchies

RELATIONAL MODEL FOR MULTIPLE HIERARCHIES

Standard Parent-Child  
The relationship with Fact\_Balance is direct, each transaction has a related account



Multiple Hierarchies  
The relationship with Fact\_Balance is made through a many-to-many relationship by using a bridge table



The Branch dimension loses the hierarchy and is flat

PROS & CONS, POSSIBLE ALTERNATIVES

## Role Dimensions

sqlbi.com



## Role Dimensions

PROS AND CONS

### Pros

- Only one physical dimension in Analysis Services
- A change in a dimension applies to all the "roles"
- A single dimension process is required

### Cons

- A single "role" cannot be modified
- Attribute and hierarchy names cannot be changed
- User interface in some clients makes it difficult to understand to which role each hierarchy belongs to

## Role Dimensions

ALTERNATIVES

### Duplicate dimension in SSAS

- Create a view for each dimension
- The view can filter rows/columns
- Rename fields in a meaningful name (for the role)
- Name of hierarchies can be role-based
- Keep an alignment between view and role

COLUMN ORDER  
CALCULATED MEMBERS  
MEASURES OR DIMENSIONS  
MOLAP VS ROLAP

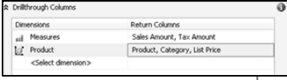
## Drill-through

sqlbi.com



## Drillthrough

COLUMNS ORDER



Change XML


- BIDS editor doesn't support column order mixing several dimensions
- After a change in MDX the graphical editor cannot be used anymore (otherwise you would lose custom order)

```

<Actions>
<Action xsi:type="DrillthroughAction">
<ID>Drillthrough Action</ID>
...
<Columns>
<Column xsi:type="MeasureBinding">
<MeasureID>Sales Amount</MeasureID>
</Column>
<Column xsi:type="MeasureBinding">
<MeasureID>Tax Amount</MeasureID>
</Column>
<Column xsi:type="CubeAttributeBinding">
<CubeID>Adv DM 1</CubeID>
<CubeDimensionID>Product</CubeDimensionID>
<AttributeID>Products</AttributeID>
<Type>All</Type>
</Column>
<Column xsi:type="CubeAttributeBinding">
<CubeID>Adv DM 1</CubeID>
<CubeDimensionID>Product</CubeDimensionID>
<AttributeID>Category</AttributeID>
<Type>All</Type>
</Column>
<Column xsi:type="CubeAttributeBinding">
<CubeID>Adv DM 1</CubeID>
<CubeDimensionID>Product</CubeDimensionID>
<AttributeID>List Price</AttributeID>
<Type>All</Type>
</Column>
</Columns>
</Actions>
    
```

## Drillthrough

ORDINAMENTO COLONNE



Change XML

- Graphical editor may show an incorrect order
- Always group by measures or dimensions

```

<Actions>
<Action xsi:type="DrillthroughAction">
<ID>Drillthrough Action</ID>
...
<Columns>
<Column xsi:type="MeasureBinding">
<MeasureID>Sales Amount</MeasureID>
</Column>
<Column xsi:type="CubeAttributeBinding">
<CubeID>Adv DM 1</CubeID>
<CubeDimensionID>Product</CubeDimensionID>
<AttributeID>Products</AttributeID>
<Type>All</Type>
</Column>
<Column xsi:type="CubeAttributeBinding">
<CubeID>Adv DM 1</CubeID>
<CubeDimensionID>Product</CubeDimensionID>
<AttributeID>List Price</AttributeID>
<Type>All</Type>
</Column>
<Column xsi:type="MeasureBinding">
<MeasureID>Tax Amount</MeasureID>
</Column>
<Column xsi:type="CubeAttributeBinding">
<CubeID>Adv DM 1</CubeID>
<CubeDimensionID>Product</CubeDimensionID>
<AttributeID>Category</AttributeID>
<Type>All</Type>
</Column>
</Columns>
</Actions>
    
```

## Drillthrough

NOT SUPPORTED ON CALCULATED MEMBERS

- Calculated measures
  - Create regular measures and use SCOPE assignments
  - Cost: MOLAP space for measure storage
- Calculated members on dimensions
  - Works with utility dimensions (i.e. DateTool)
  - Consider semantic meaning (i.e. YTD)
- Use Action – Rowset
  - You can specify a DRILLTHROUGH operation
  - Query generated server-side
    - Cannot be controlled by the client – i.e. MAXROWS

## Drillthrough

OPTIONS FOR NUMERICAL ATTRIBUTES (I.E. DOCUMENT NUMBER)

Transaction Detail Dimension

- Supports supplementary information (also textual)
- Long process time (requires SELECT DISTINCT)
- Large dimension: consider to hide it

ROLAP Dimension

- Reduced processing time
- Heavy load at query time (full scan of fact table with a SQL query)

Measure on a separate fact table

- Only for numeric values
- Aggregation Function = None (only Enterprise)
- Same fact table if only one integer value is required
- Processing time affordable (doesn't require DISTINCT)

## Drillthrough

MOLAP VS. ROLAP

MOLAP

- Requires SELECT DISTINCT for each dimension
- Longer processing time if a dimension is added just for drillthrough attributes
- Fast query time, both on measures and on large dimensions


ROLAP

- Process time is immediate
- Slow queries (pressure on relational engine, SQL queries)
- Almost never a good choice

## Calculation Dimensions

TIME INTELLIGENCE WIZARD LIMITS  
TECHNIQUES AND APPLICATION IN DATETOOL

sqlbi.com





## Calculation Dimension

BASE CONCEPTS

Goal: apply the same calculation to different measures

- Year-to-Date
- Year-over-Year
- Difference over Year

Implementation

- “Virtual” space in the multidimensional model
- Each member in each attribute defines a calculation
- MDX Script contains the calculation formulas

## Time Intelligence Wizard

WIZARD LIMITS IN BI DEVELOPER STUDIO

Attribute Overwrite

- Create attributes on existing hierarchy
  - <http://tinyurl.com/attoverwrite>
  - <http://tinyurl.com/chrisattoverwrite>

Wizard generated formulas have limits on multiple hierarchies

- i.e. Year-Month and Year-Week

Use calculated members

- Drillthrough cannot be used
- Security cannot be applied
- Issues with clients like Excel 2007

## Calculation Dimension

DESIGN PATTERN

SQL Views

- View with member for each attribute
- CROSSJOIN between attributes for a dimension

Attributes properties

- Remove All Member (IsAggregatable = False)

Dimensions relationships

- Add dimension to the cube
- No relationship with any measure group

MDX Script

- Use SCOPE to assign the MDX formula to the corresponding member in the attribute

## Conclusions

- Leverage on data modeling
  - Relational
  - Multidimensional
- Use decoupling views
- Avoid MDX if possible
  - Calculate values inside ETL if you can
  - Simplify maintenance by splitting complex calculations in a higher number of simple calculations
- Hide “dangerous” attributes to the end-user
  - i.e. dimensions that supports drillthrough

# Questions and Answers

sqlbi.com



## Links

- SQLBI Website  
[www.sqlbi.com](http://www.sqlbi.com)
- PowerPivot Workshop  
[www.powerpivotworkshop.com](http://www.powerpivotworkshop.com)
- Marco Russo blog  
[www.sqlblog.com/blogs/marco\\_russo](http://www.sqlblog.com/blogs/marco_russo)
- Alberto Ferrari blog  
[www.sqlblog.com/blogs/alberto\\_ferrari](http://www.sqlblog.com/blogs/alberto_ferrari)

For any question contact us at  
[info@sqlbi.com](mailto:info@sqlbi.com)

**Thank you!**