PASS SUMMIT 2013

Fast Performing SSAS Hierarchies

Tips & Tricks

October 15-18, 2013
Charlotte, NC

Jose Chinchilla, President & CEO
Agile Bay, Inc.

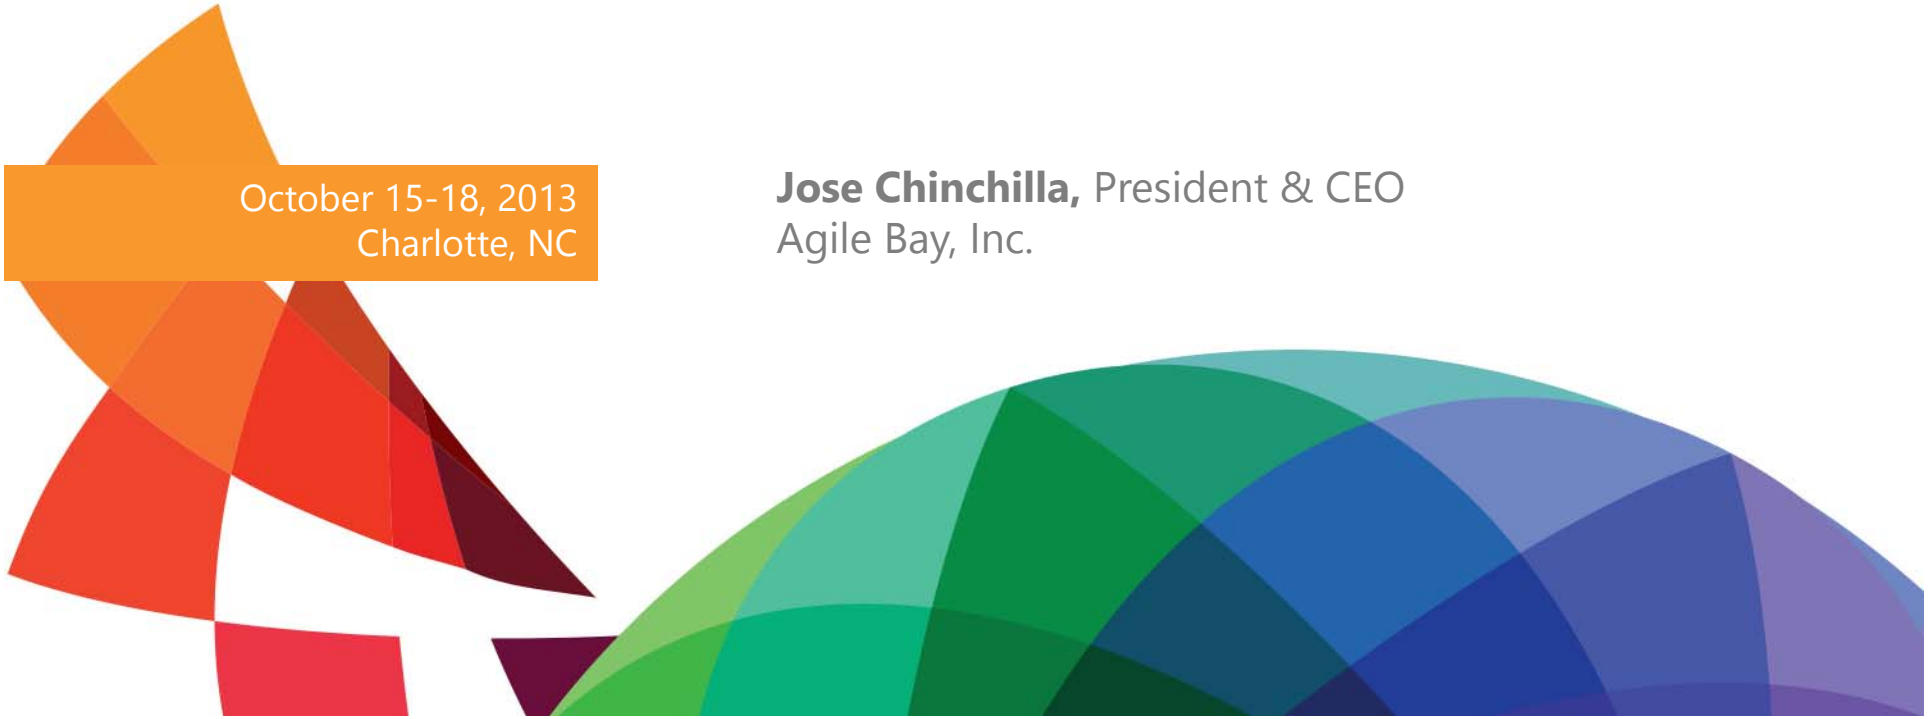

Jose Chinchilla

MCITP: Database Administrator, SQL Server 2008
MCITP: Business Intelligence SQL Server 2008

Current Positions:
President, Agile Bay, Inc.
President, Tampa Bay Business Intelligence User Group
Regional Mentor, PASS Greater Southeast

Blog: http://www.sqljoe.com
Twitter: @sqljoe
Linked-in: http://www.linkedin.com/in/josechinchilla
Email: jchinchilla@sqljoe.com

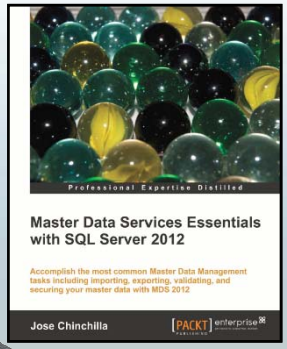

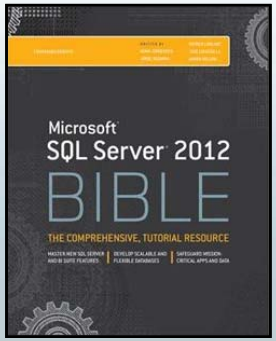

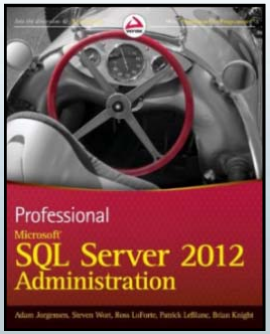

Customers & Partners

Please silence cell phones

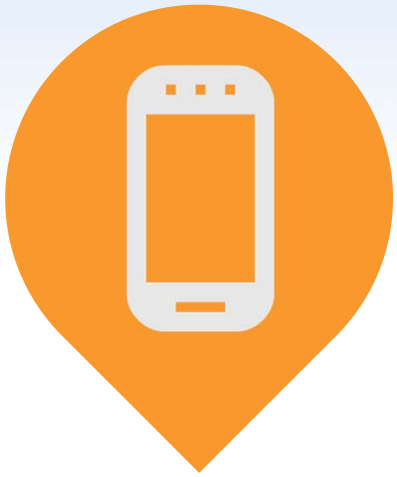

Explore Everything PASS Has to Offer

24 HOURS OF PASS

Free SQL Server and BI Web Events

PASS SQL saturday

Free 1-day Training Events

PASS SQLRally

Regional Event

PASS SUMMIT

This is Community

PASS BUSINESS ANALYTICS CONFERENCE

Business Analytics Training

PASS CHAPTERS

Local User Groups Around the World

Session Recordings

PASS CommunityCONNECTOR

PASS Newsletter

PASS VIRTUAL CHAPTERS

Free Online Technical Training

PASS

Session Evaluations

Your feedback is important and valuable.

3 ways to access

Submit by 5pm Friday Oct. 18 to **WIN** prizes

Go to **passsummit/evals**

Download the GuideBook App and search: **PASS Summit 2013**

Follow the QR code link displayed on session signage throughout the conference venue and in the program guide

PASS

Agenda: Part 1

- **Types & classification of hierarchies**
  - Natural & unnatural hierarchies
  - Attribute hierarchies vs. User hierarchies
  - Parent-child hierarchies
  - Balanced, unbalanced & ragged hierarchies
- **Demo**
  - Designing & optimizing hierarchies
    - Dim Products
    - Dim Reseller
    - Dim Geography
    - Dim Employee
    - Dim Date
    - Dim Account
- **Break**

Agenda: Part 2

- **Troubleshooting performance**
  - DMVs
  - Performance Monitor
- **Aggregations & Partitions**
  - Portioning strategy
  - Default aggregations
  - Usage based optimizations
- **Drill-down or Drill-through?**
- **Demo**
  - Adding Aggregations and Partitions
  - Adding Drill-through actions

Hierarchies

Types & Classification

Hierarchies everywhere

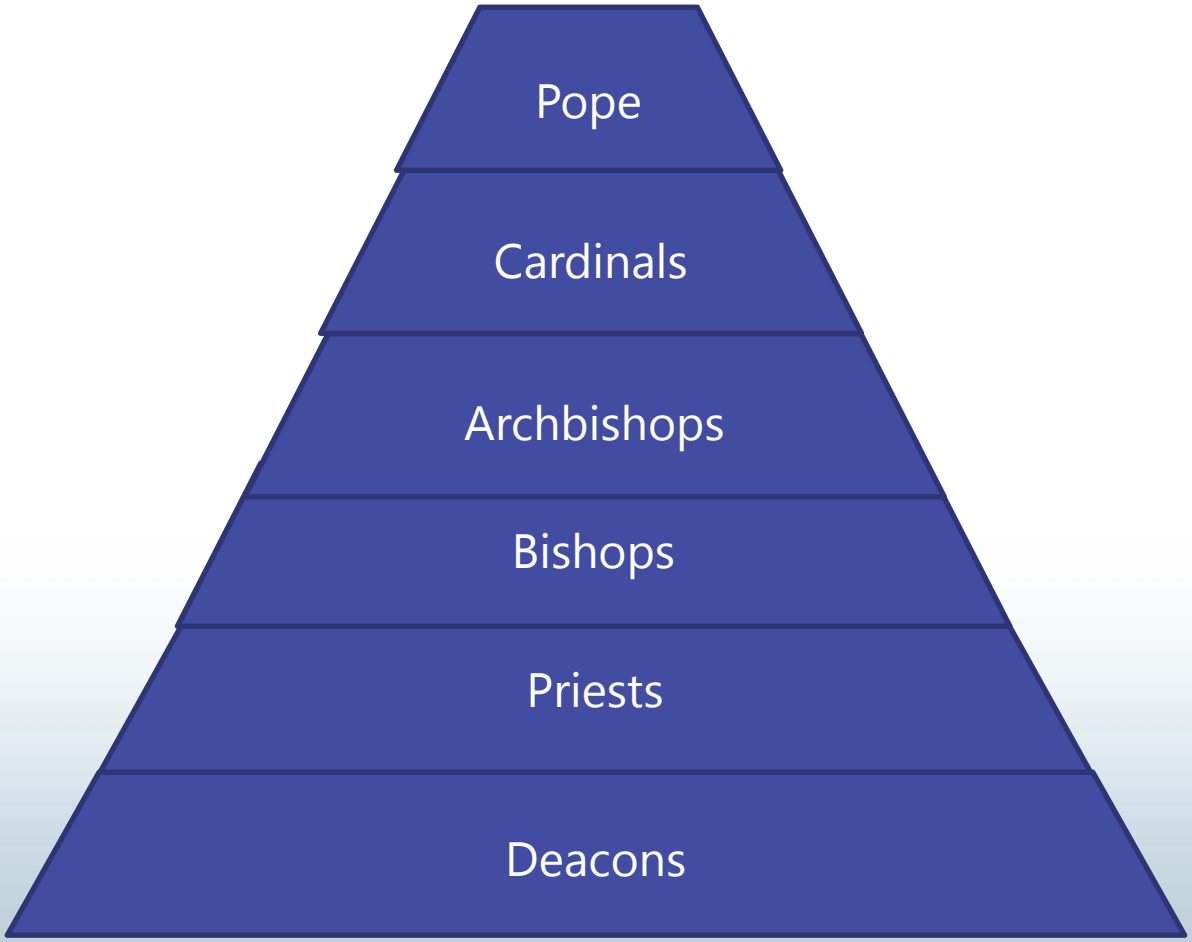

Hierarchies everywhere

Maslow's Hierarchy of needs

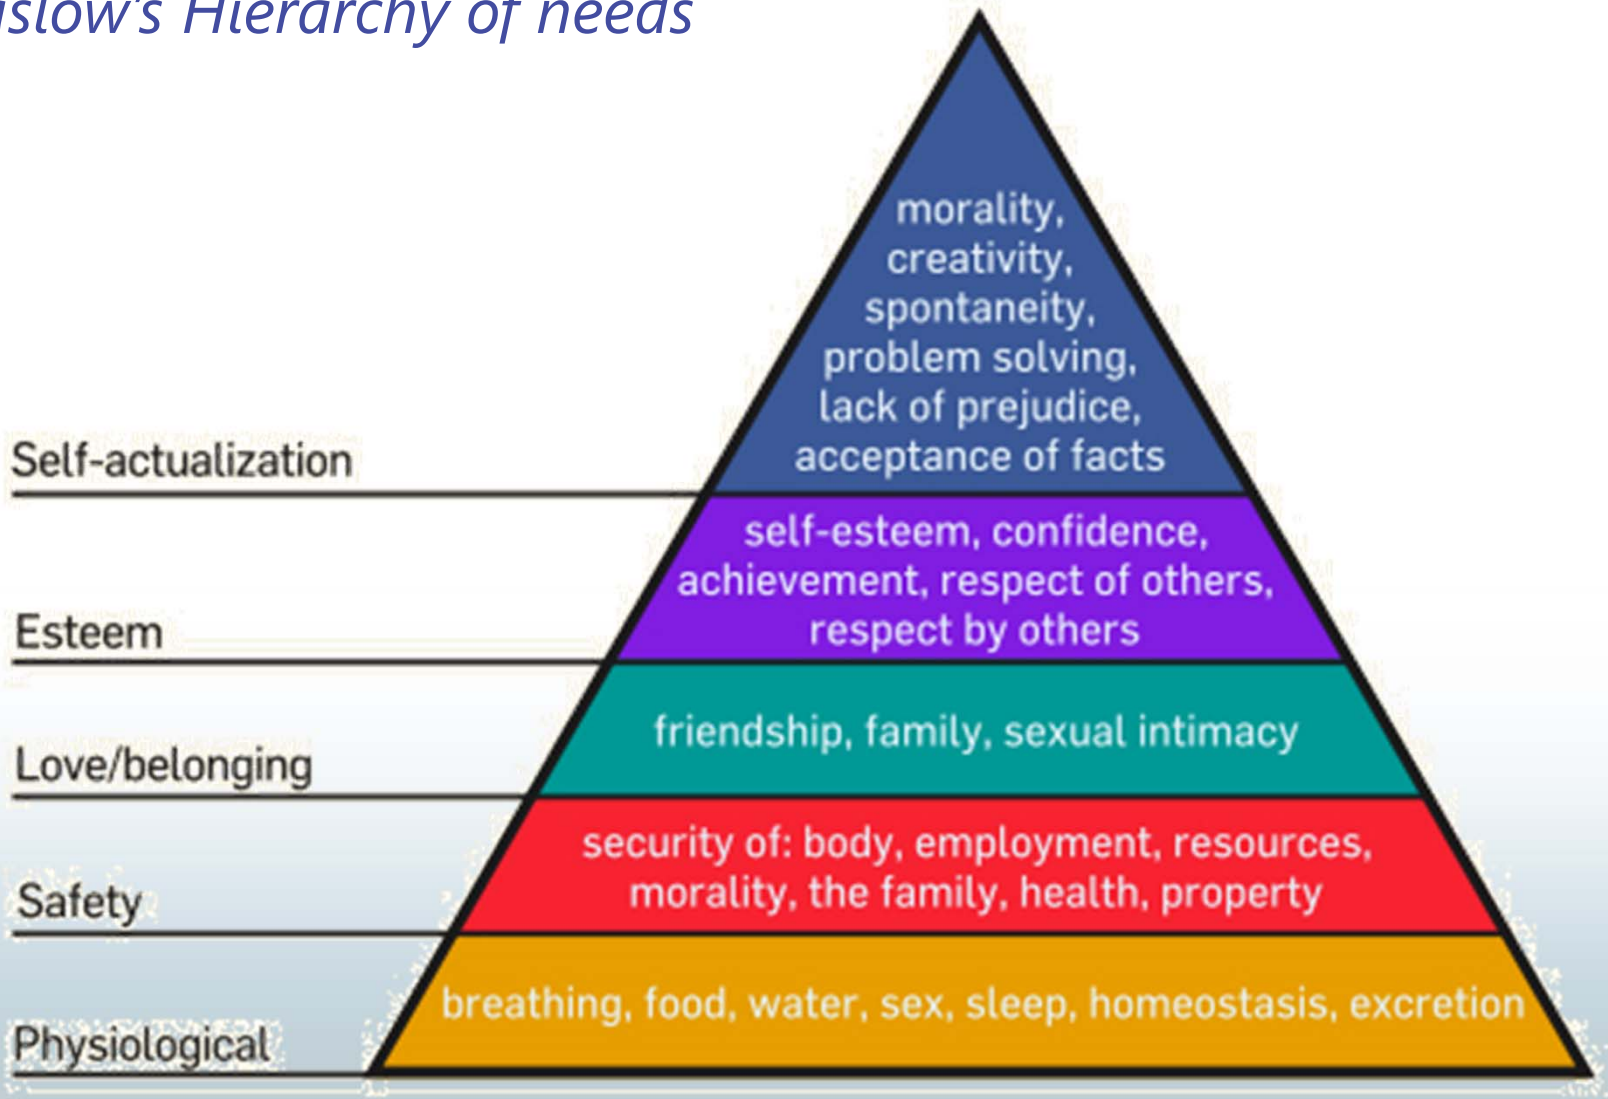

HIERARCHY OF BEARDS.

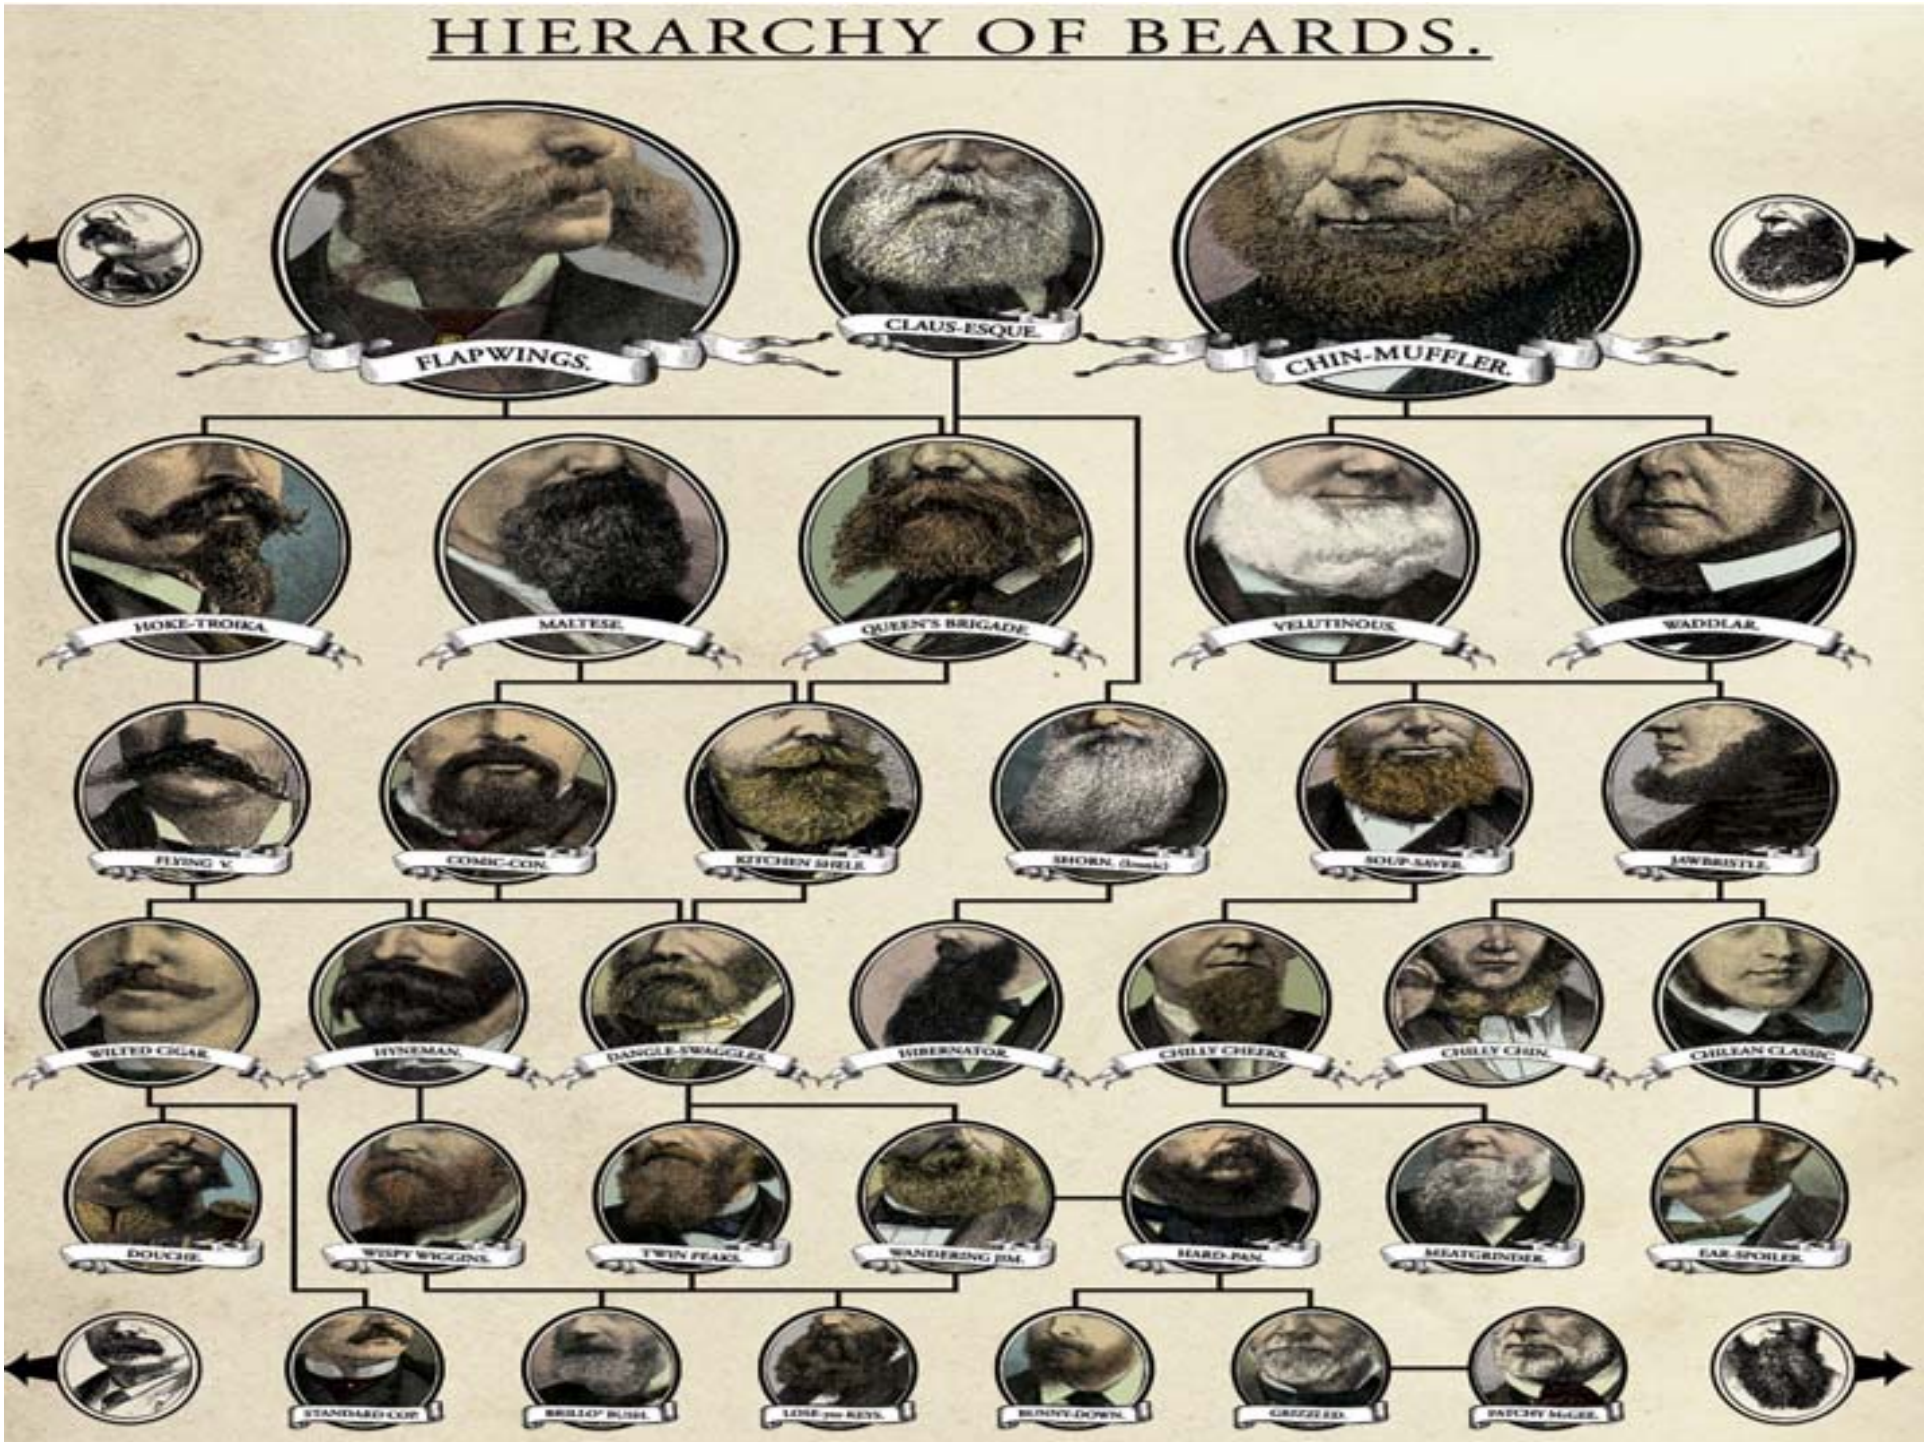

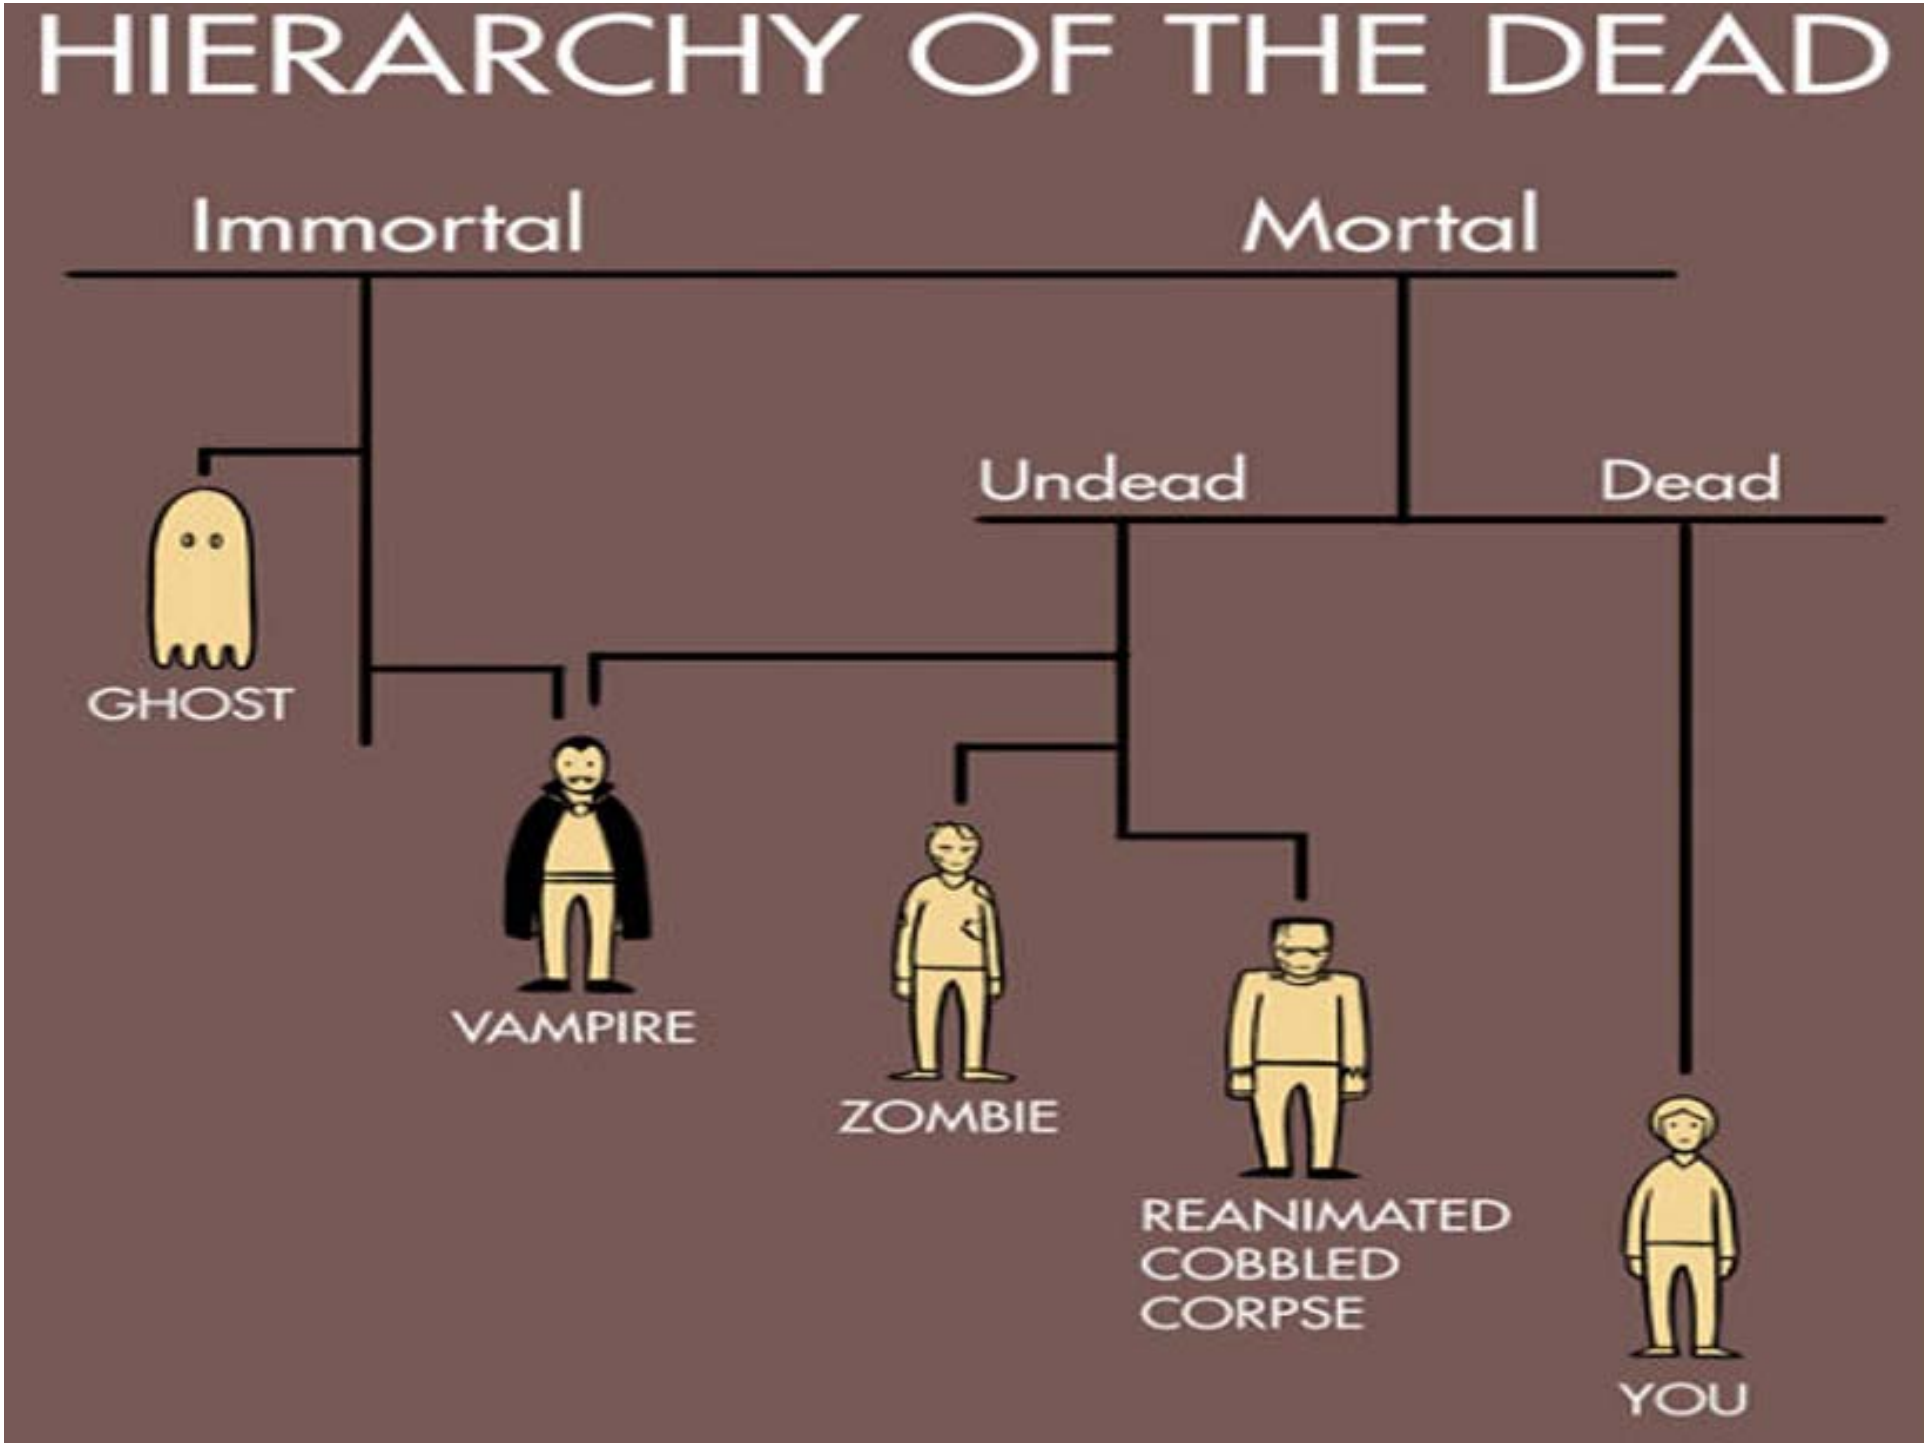
HIERARCHY OF THE DEAD
Immortal
Mortal
Undead
Dead
GHOST
VAMPIRE
ZOMBIE
REANIMATED
COBBLED
CORPSE
YOU

Executives

Analysts

Developers

DBAs

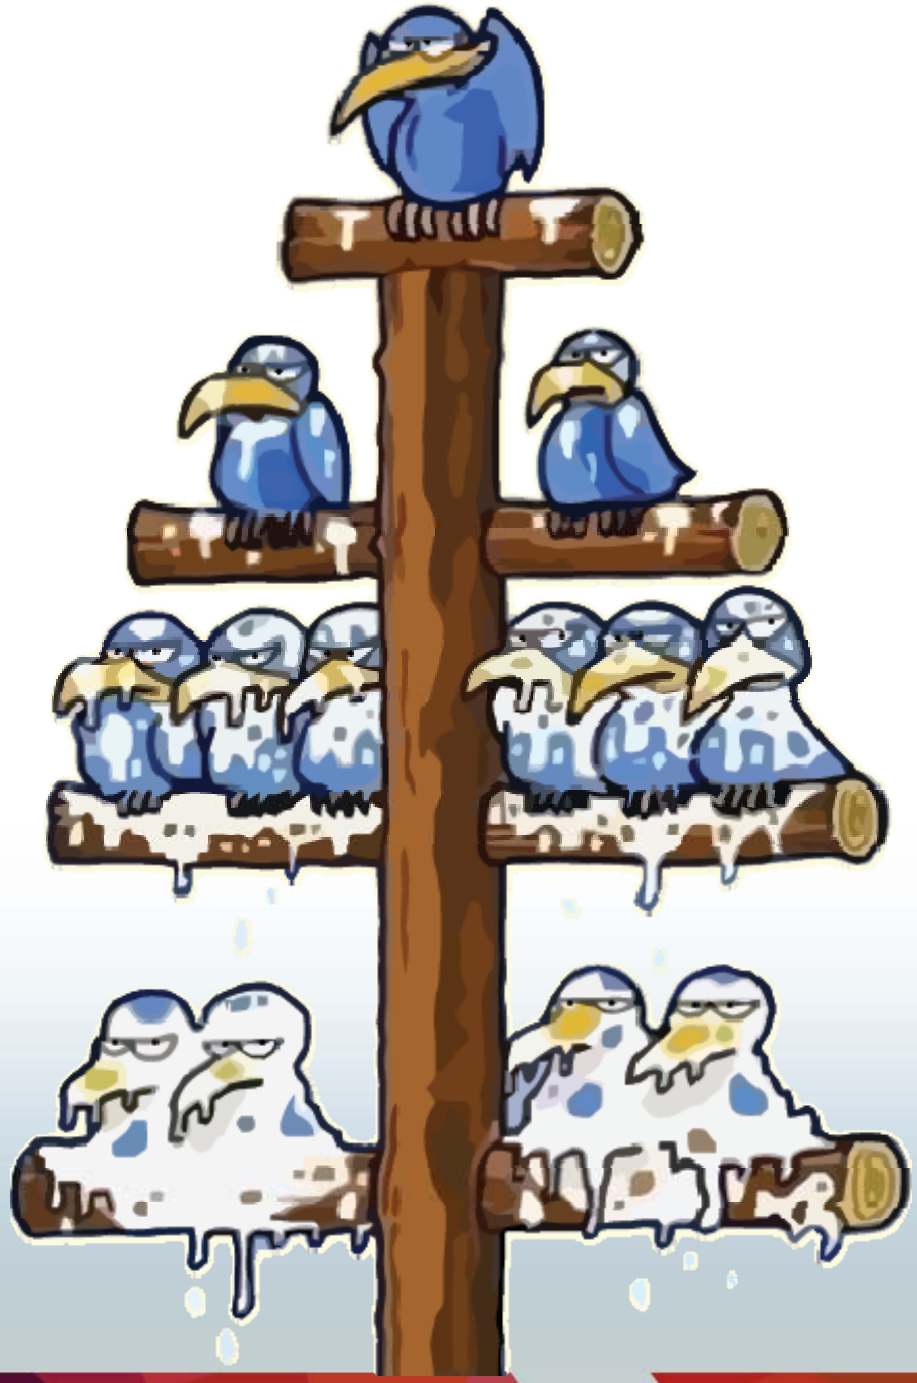

Blame Flow

Analysis Services Hierarchies

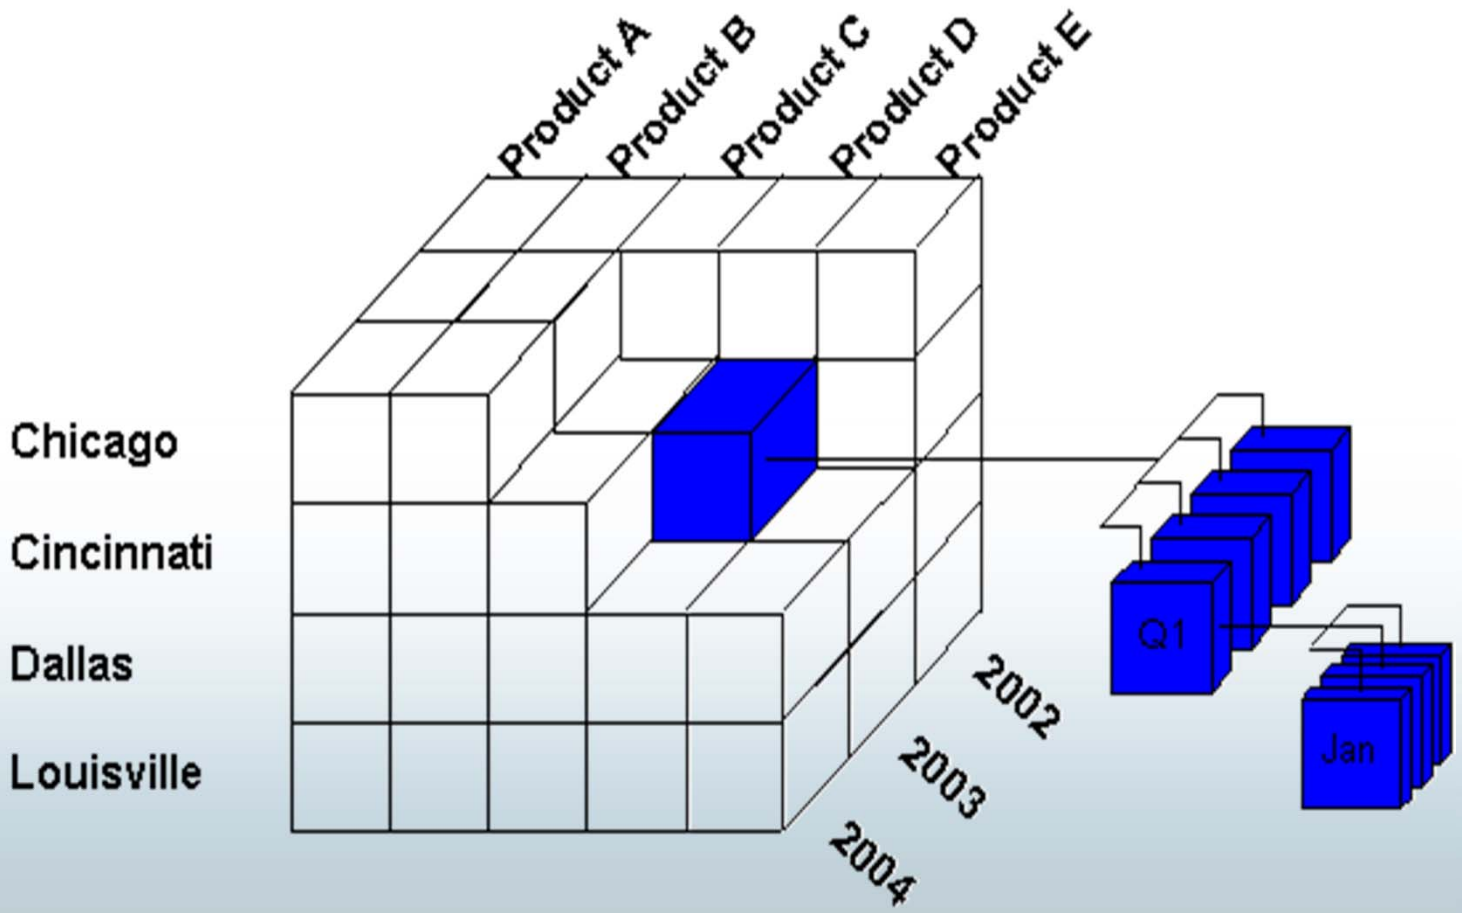

Natural & Unnatural Hierarchies

Natural Hierarchies

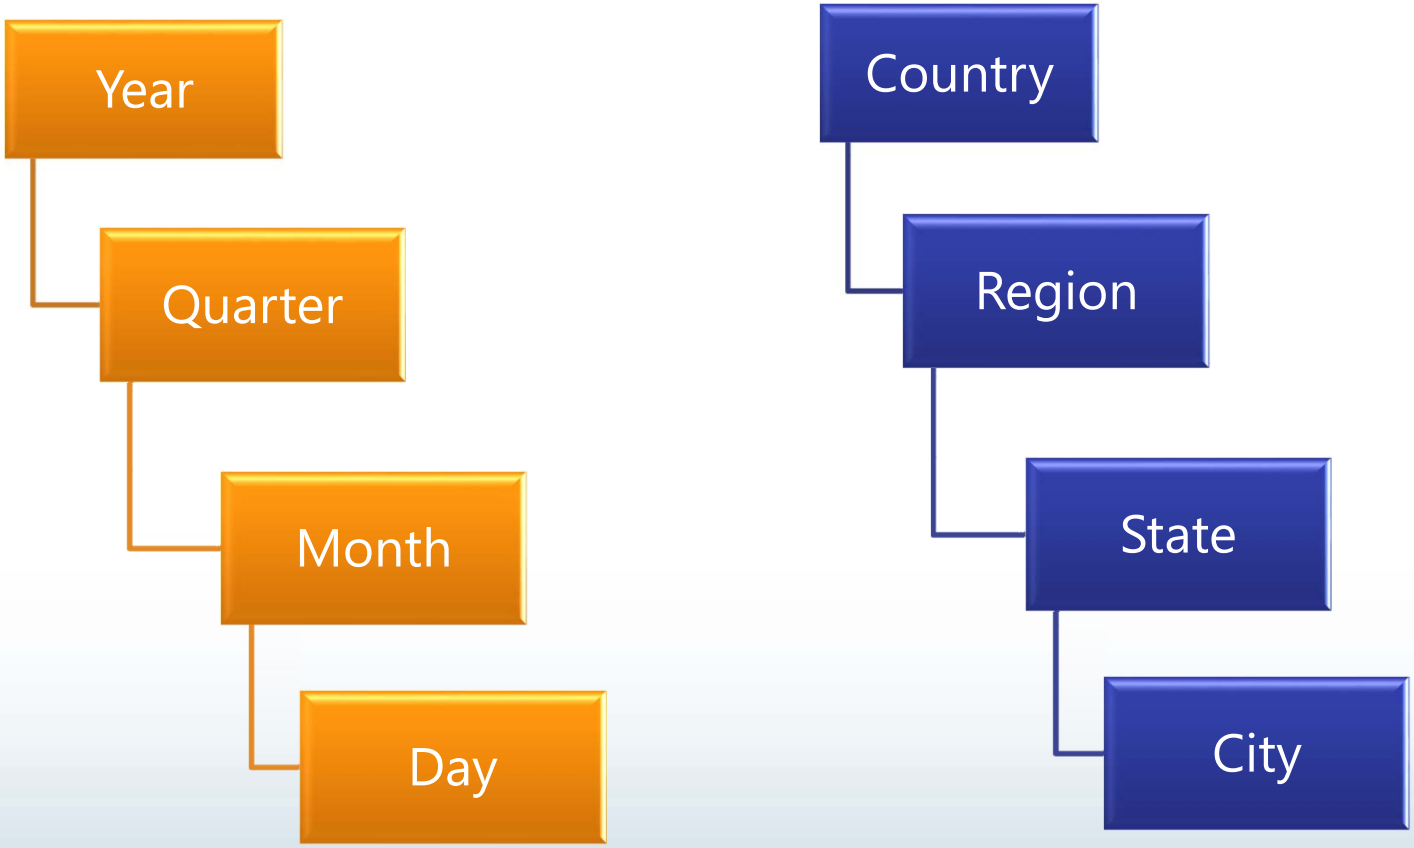

Unnatural Hierarchies

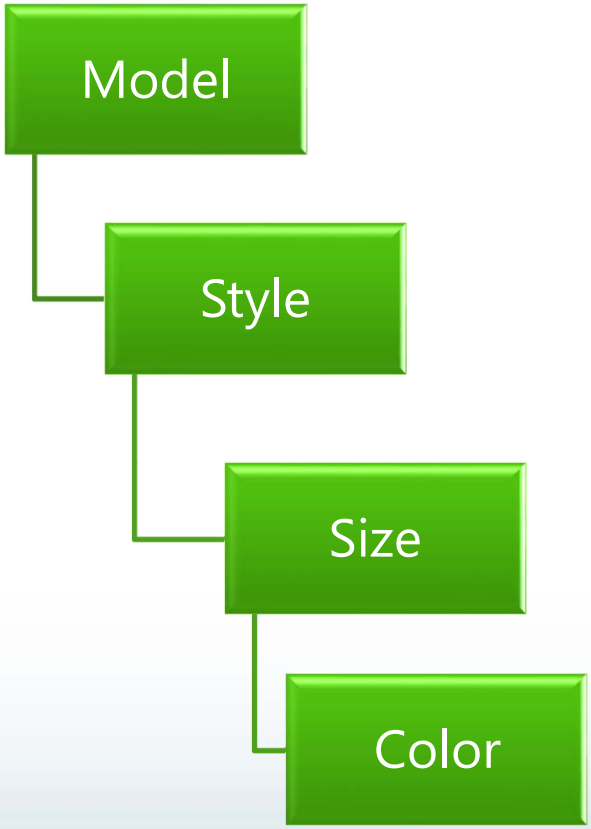

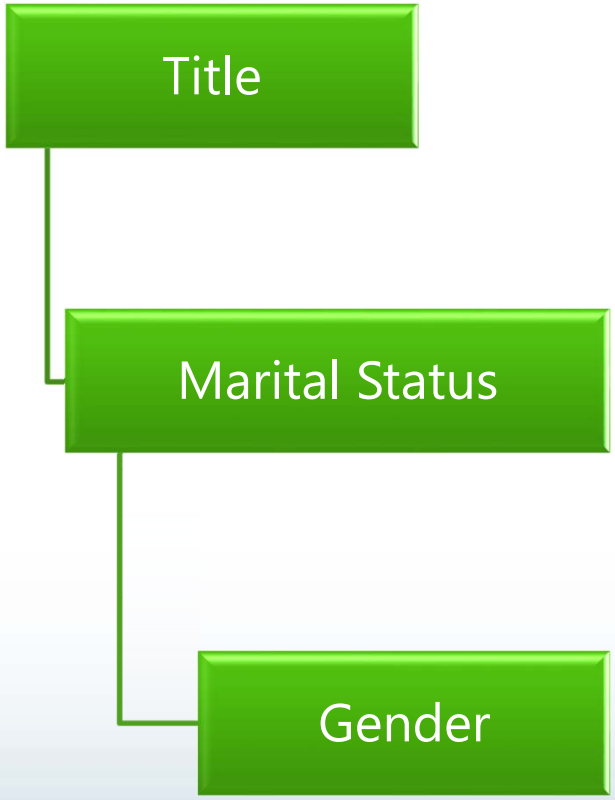

Natural &Unnatural Hierarchies

Natural	Unnatural
1. Hierarchy tree is materialized on disk in hierarchy stores	1. Not materialized on disk
2. Attributes automatically considered to be aggregation candidates.	2. Not considered as aggregation candidates.
3. Better performance	3. Perform poorly

Tip #1

Avoid unnatural hierarchies.

Balanced & Unbalanced Hierarchies

Balanced Hierarchy

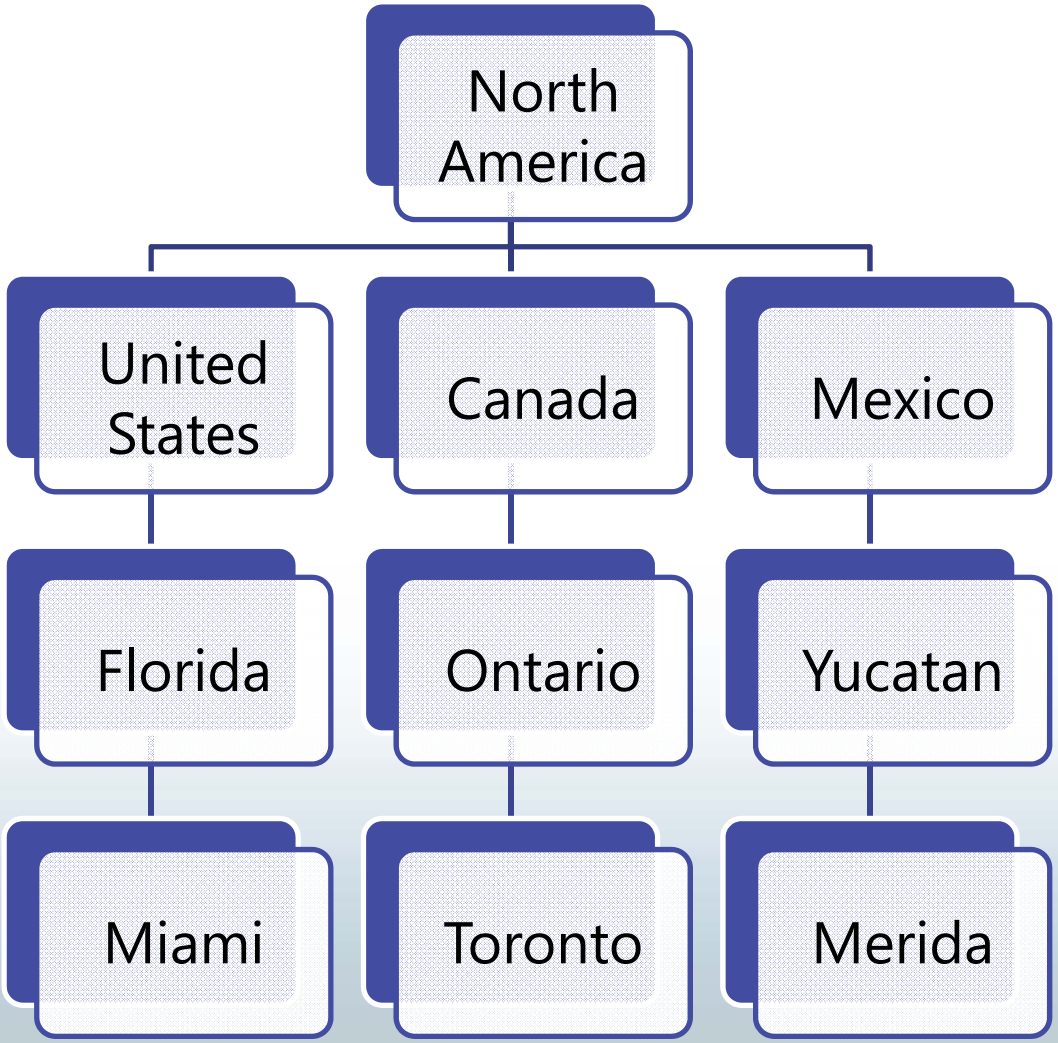

Unbalanced Hierarchy

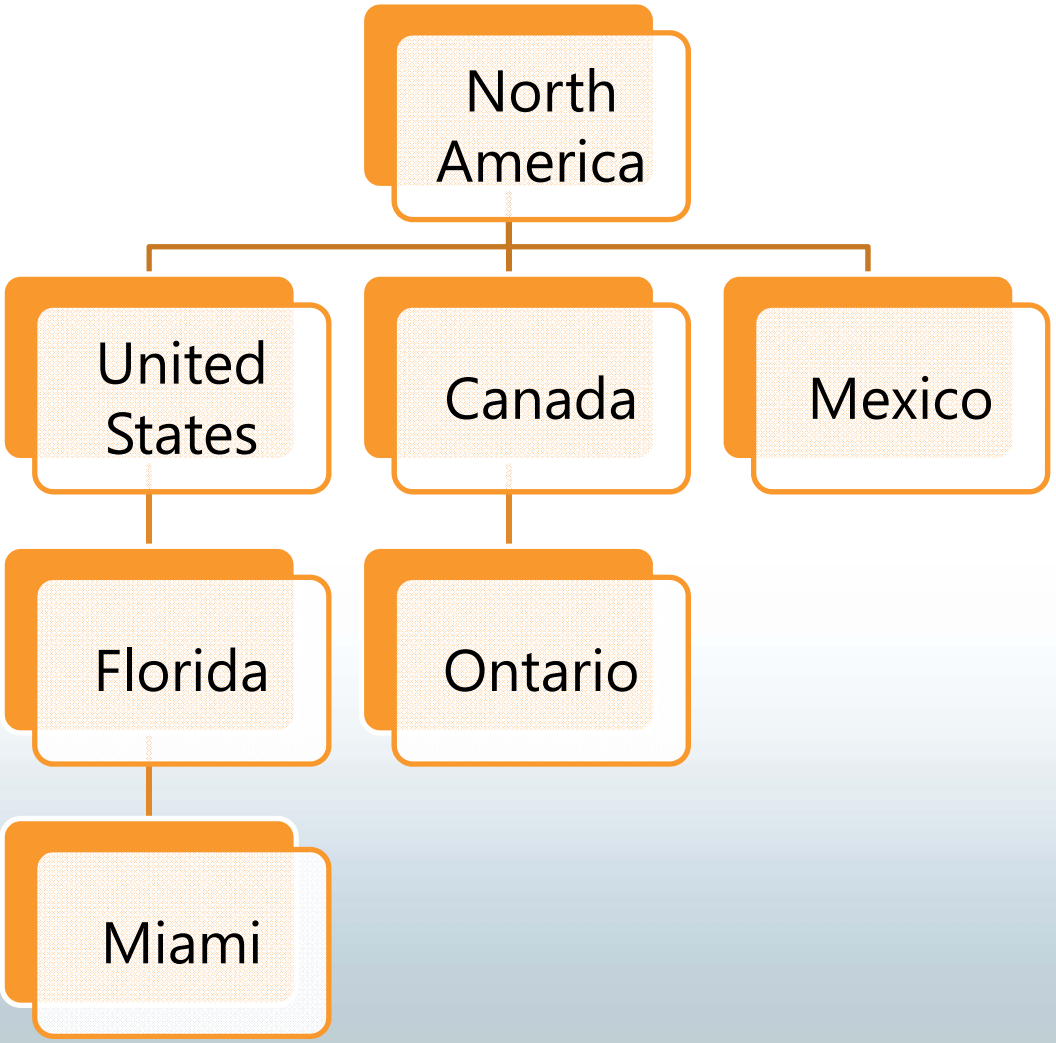

Ragged Hierarchy

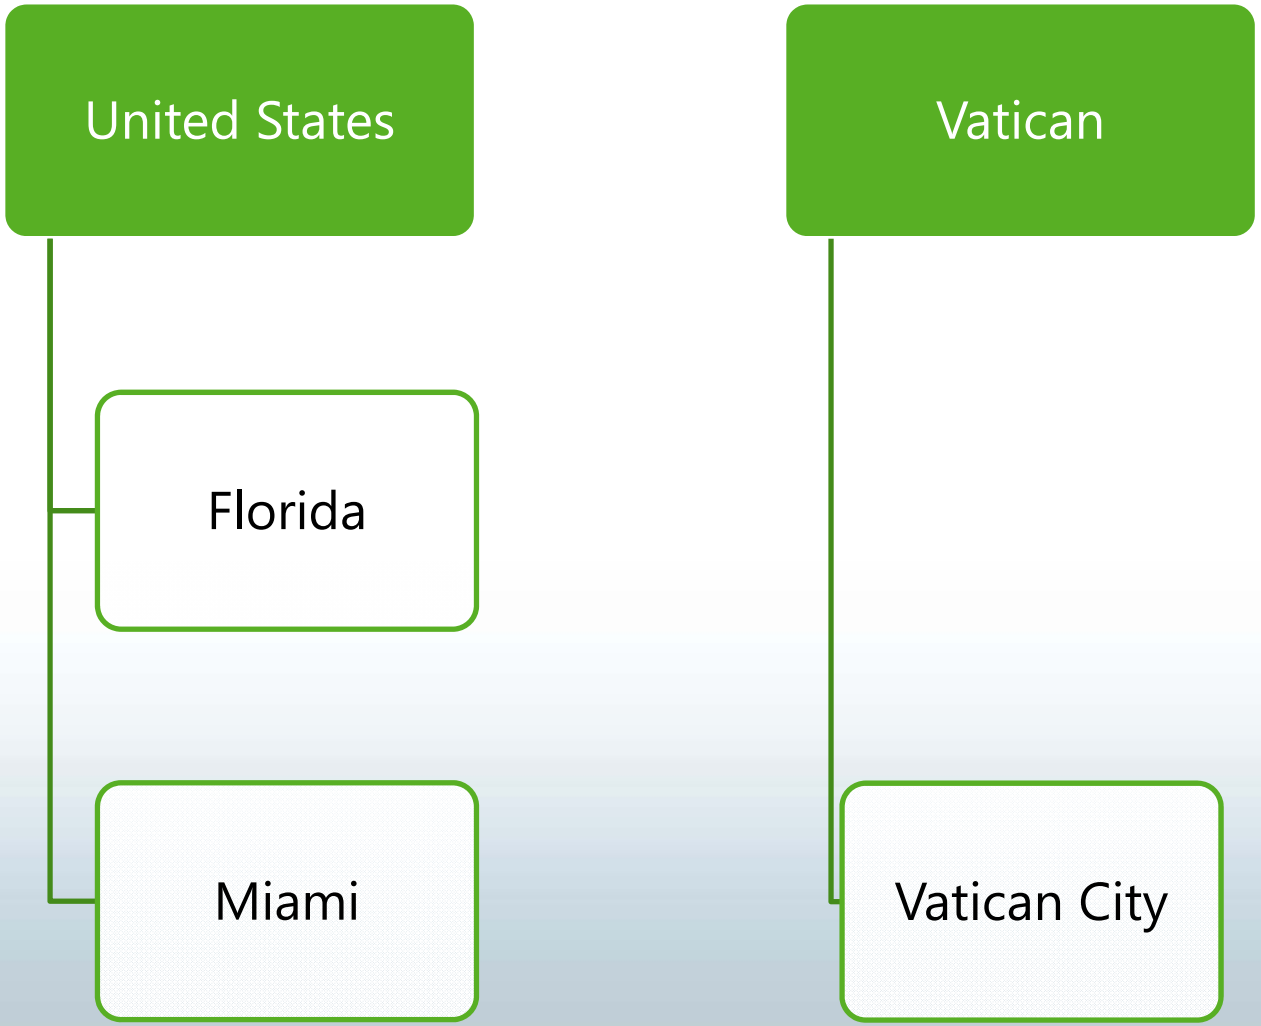

Dimension Designer

Attribute vs. User Hierarchies

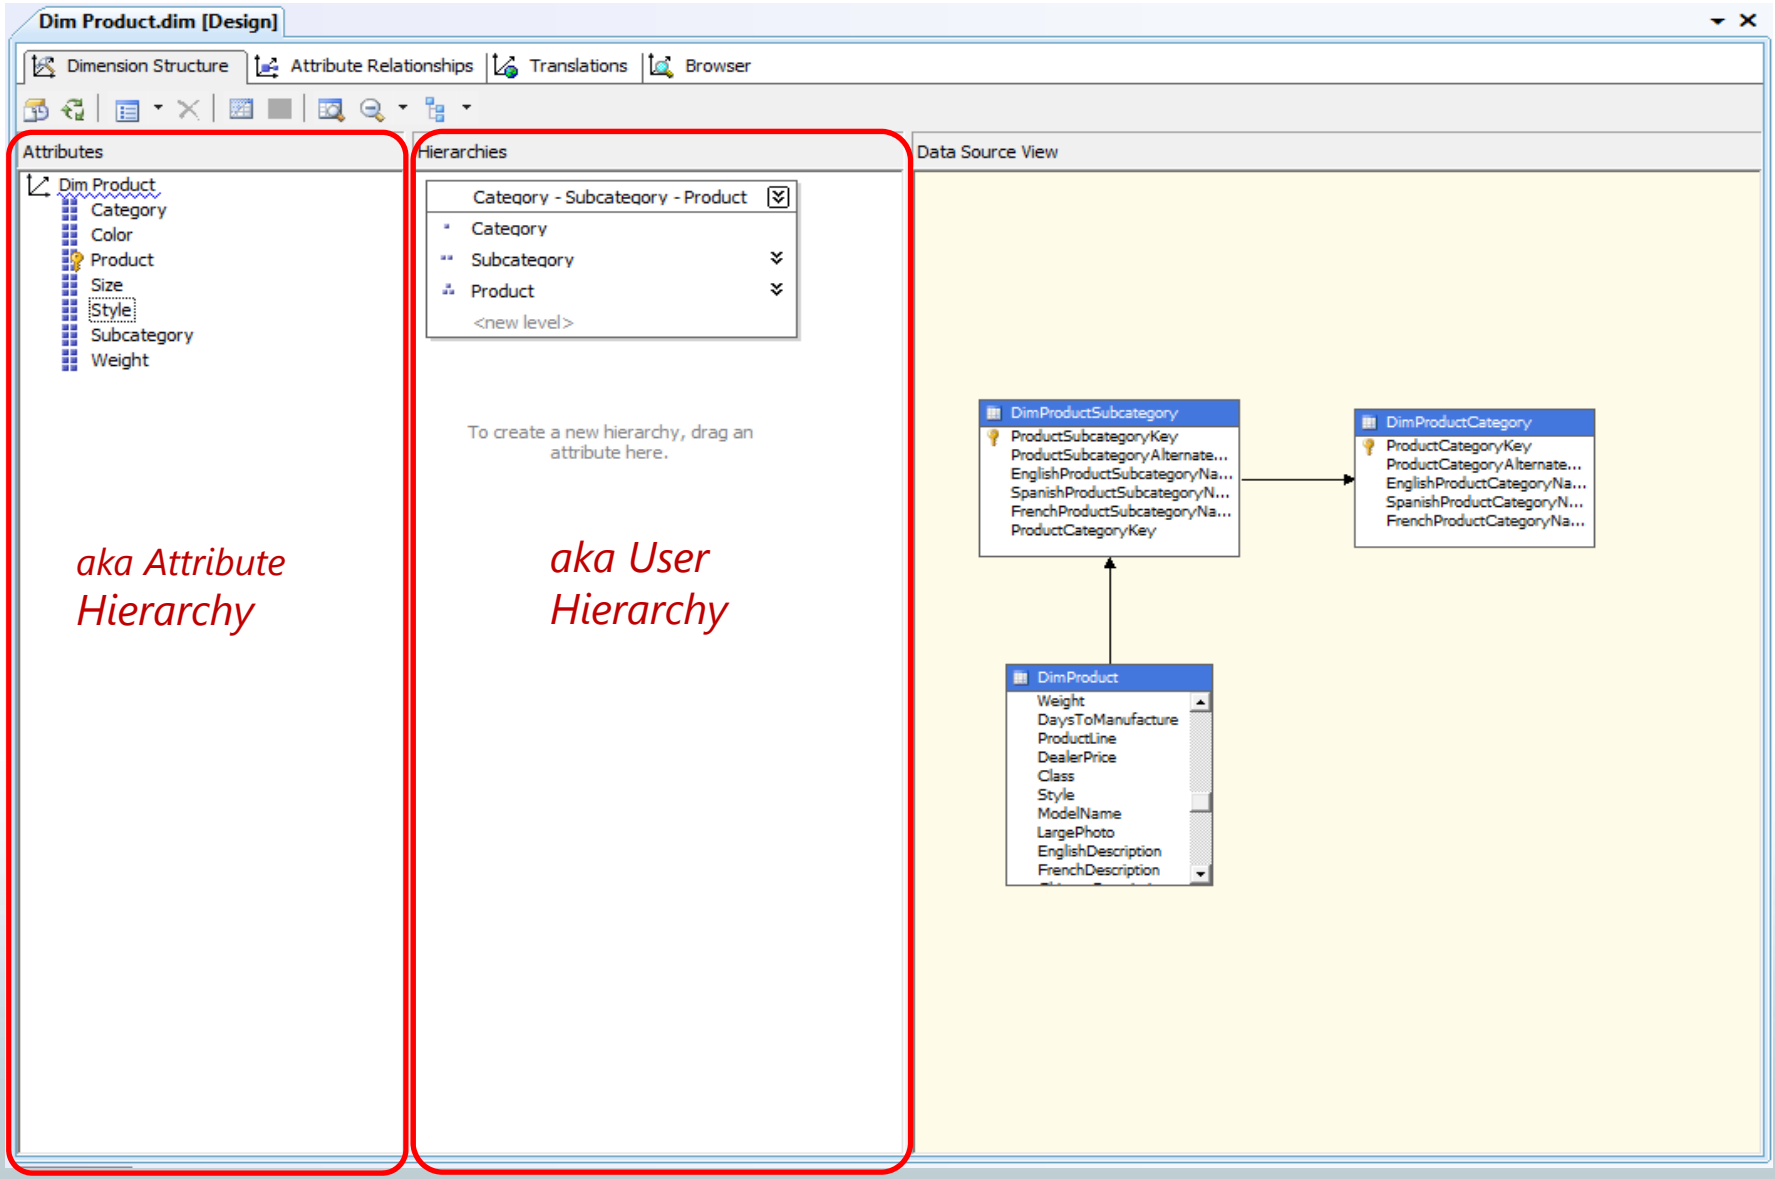

aka Attribute Hierarchy

aka User Hierarchy

Default Attribute Hierarchy – "ALL"

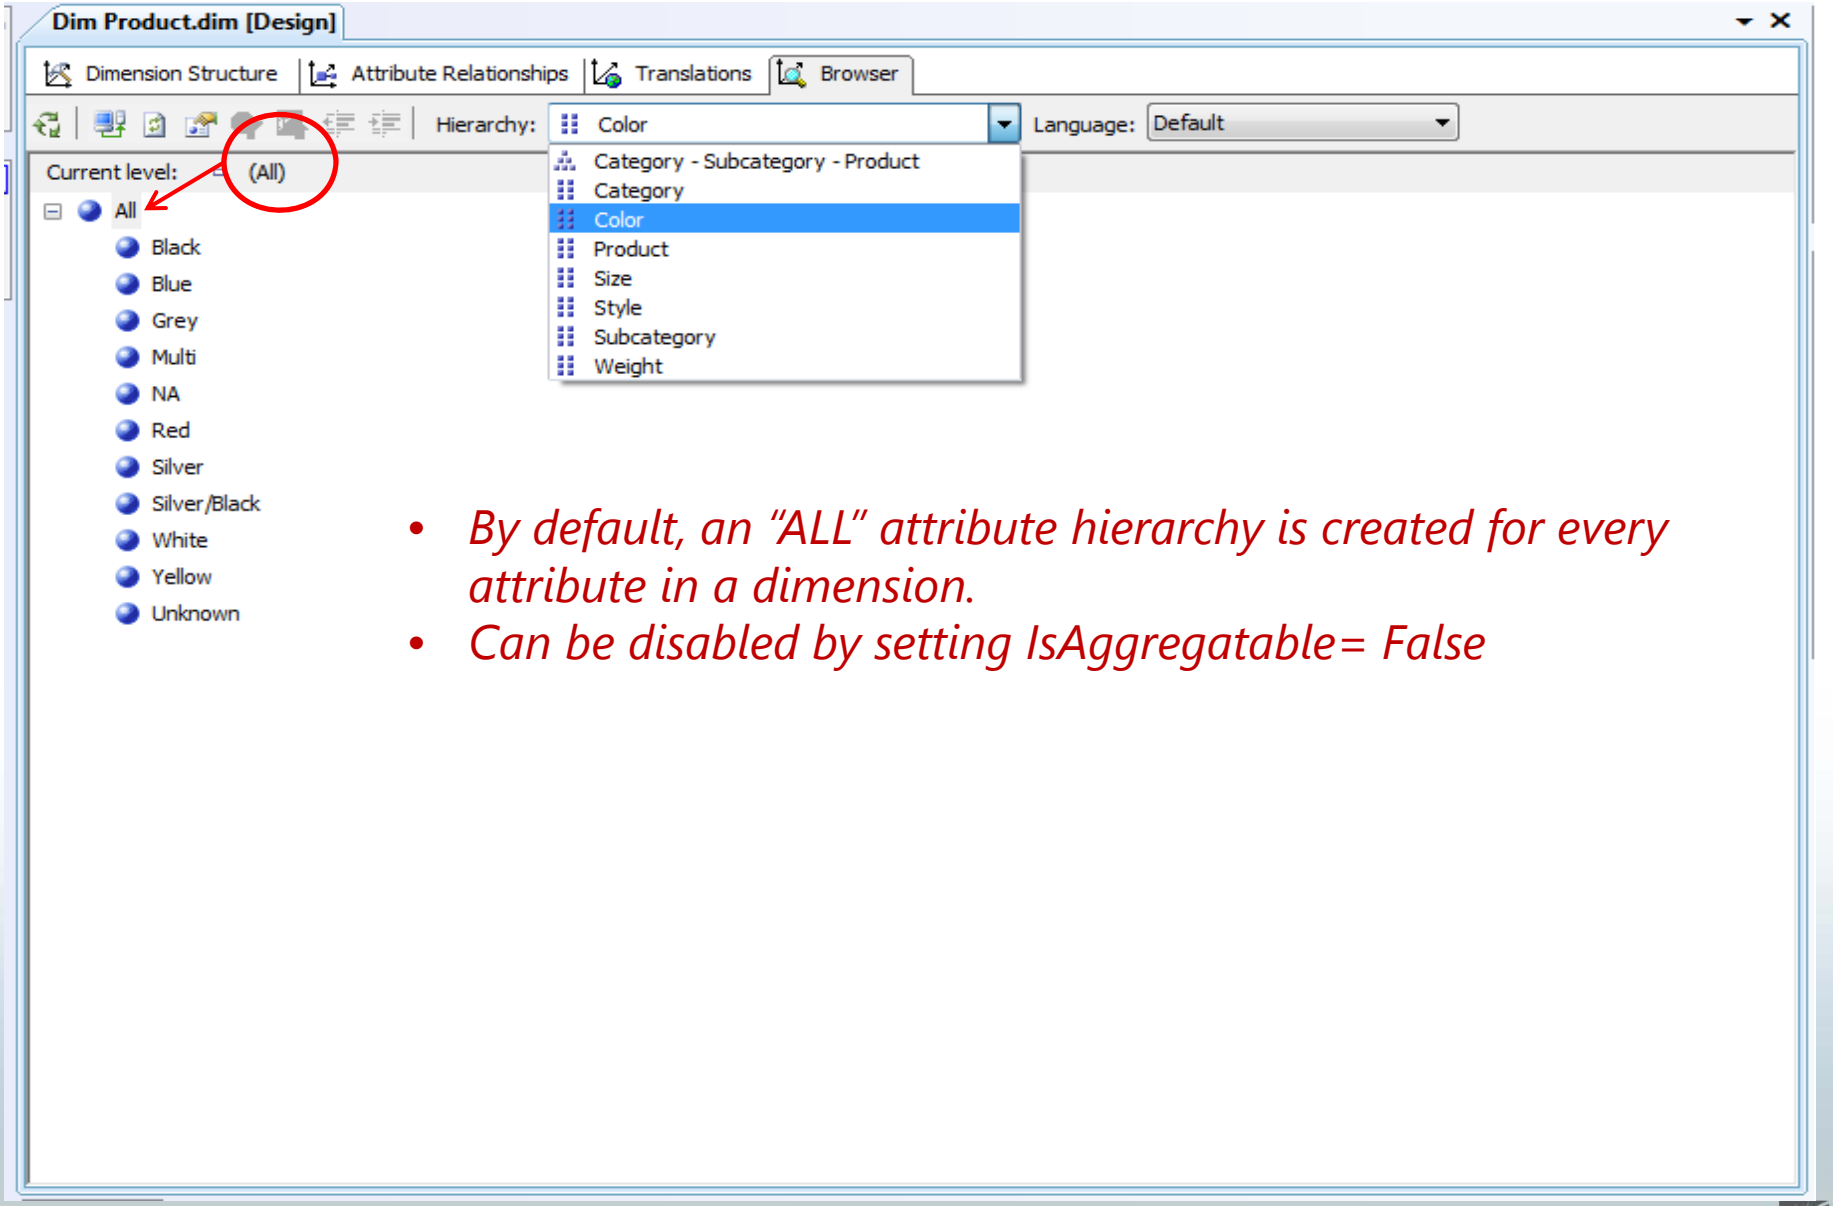

- *By default, an "ALL" attribute hierarchy is created for every attribute in a dimension.*
- *Can be disabled by setting IsAggregatable= False*

User hierarchies

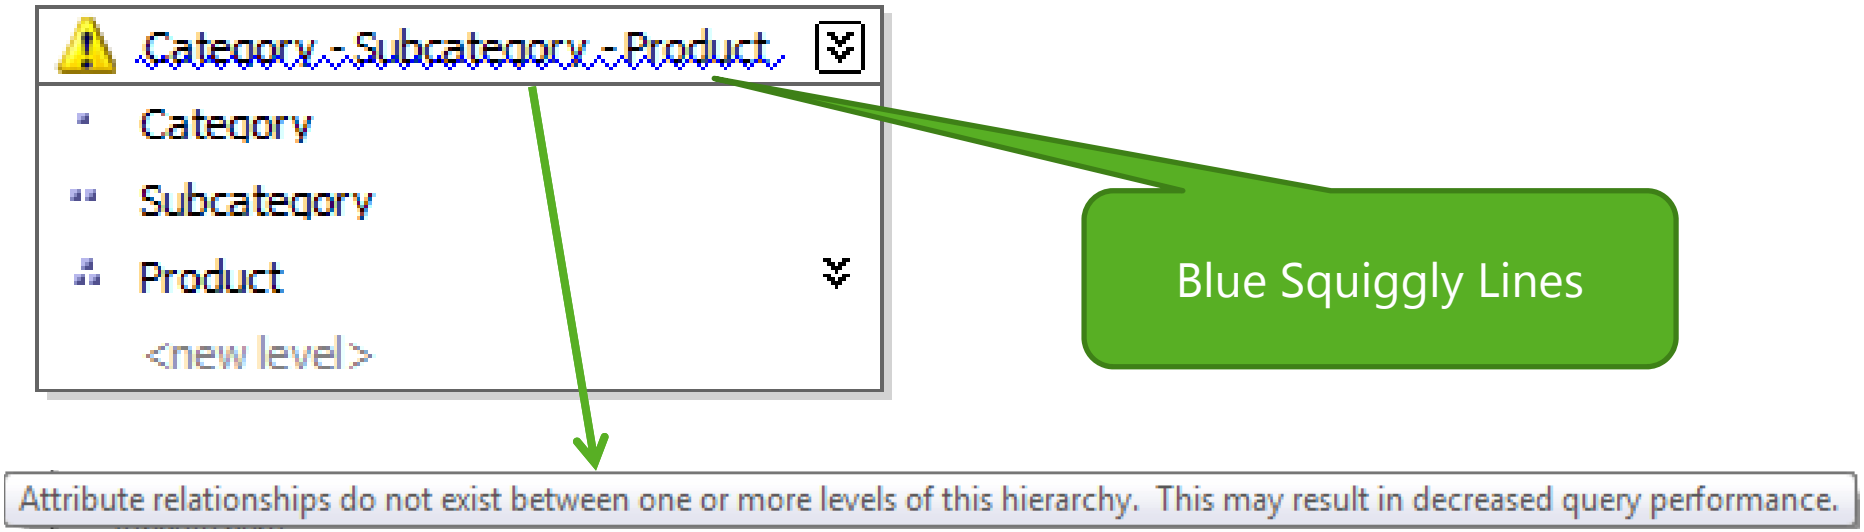

- *Attribute relationships are essential for better query performance.*

User hierarchies: Attribute Relationships

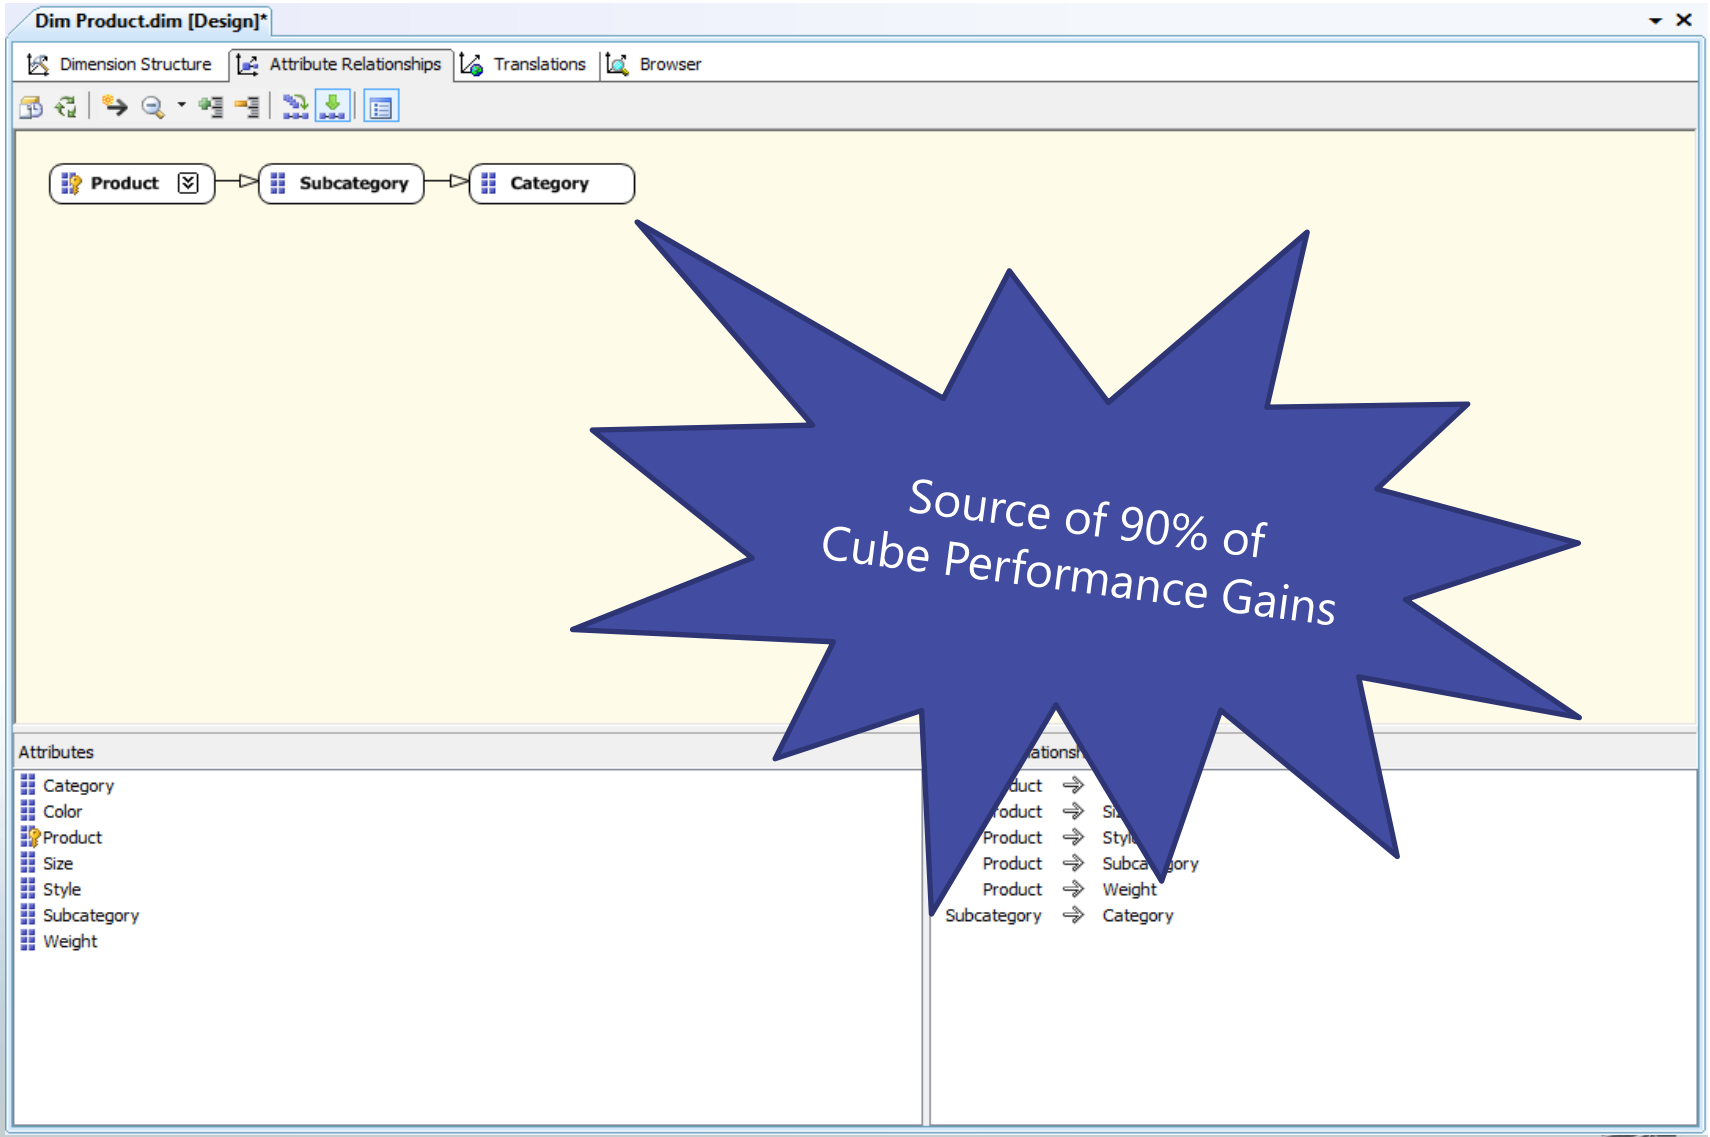

Tip #2

Always create attribute relationships in user hierarchies.

User hierarchies: Attribute Relationships

Inefficient

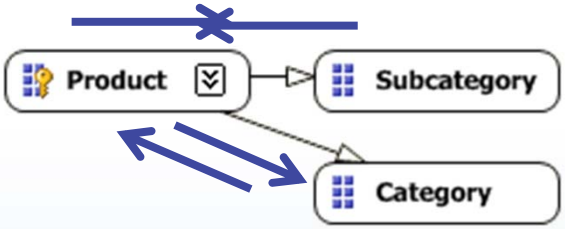

Efficient

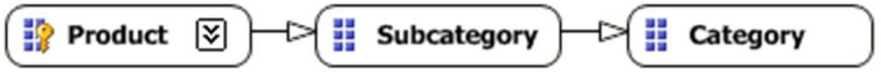

More Detail Less Detail

Exercise:
Follow my daddy!

Dimension & Attribute Properties

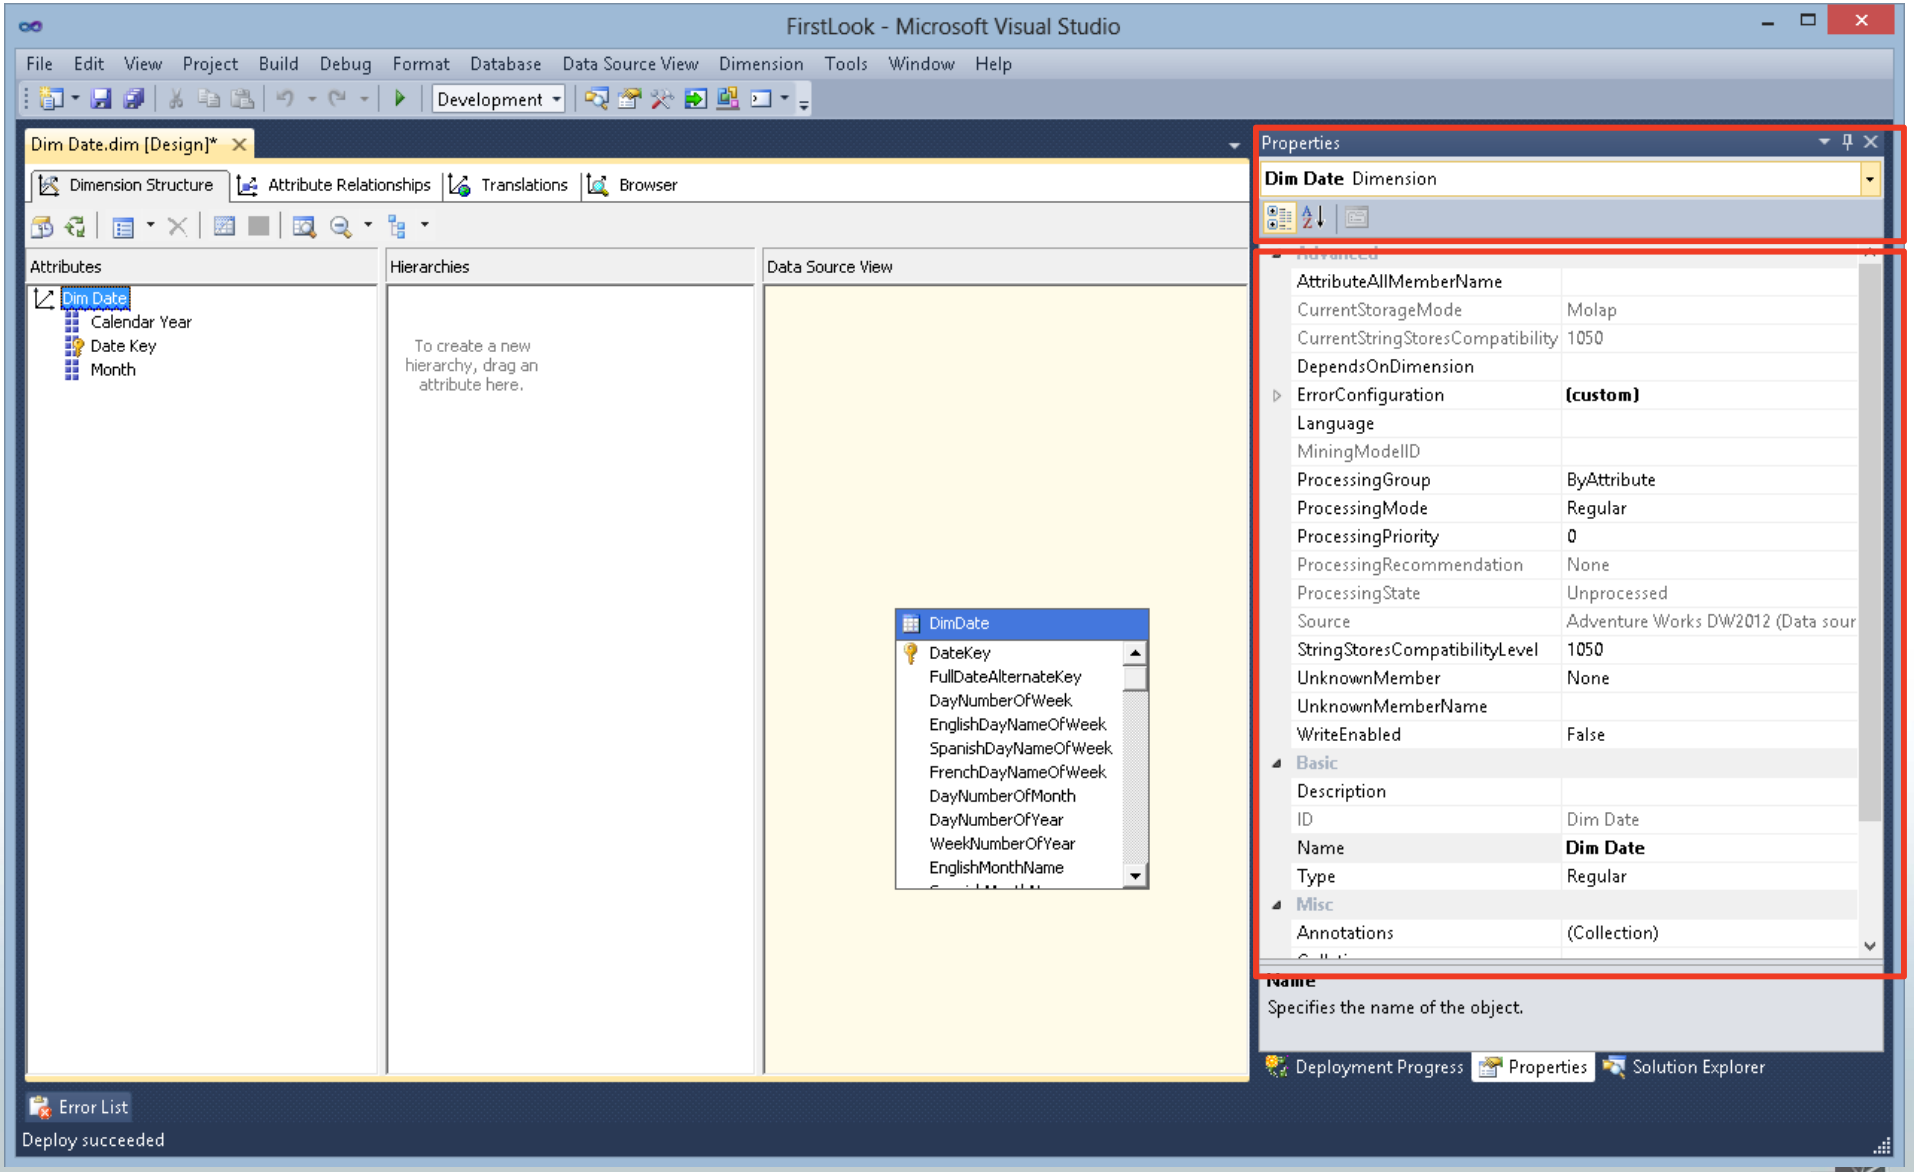

Product Hierarchy

Product Hierarchy

- Snowflake
  - SubCategory
  - Category
- Attribute Relationships
  - Key based
- Multiple Drilldowns
  - ProductLine – Model
- Attribute Discretization
  - Weight
  - Size Range

Demo:
Product Hierarchy

Reseller & Geography Hierarchy

Reseller & Geography Hierarchy

- Snowflake
  - Geography
- Reseller Dimension
  - Geography - Reseller Drilldown
- Geography Dimension
  - Referenced
  - Materialize

Demo:
Reseller & Geography Hierarchy

Employee Hierarchy

Employee Hierarchy

- Parent-Child
  - Self-join / Self-referencing
  - Read / write performance scalability issues
- Alternative Options:
  - Flat table / Naturalization (Adjacency list)
  - Hierarchy IDs (SQL Server 2008)
  - Many-to-Many
  - Snowflake

Employee Hierarchy

- Hierarchy Properties
  - HideMemberIf
  - Support for Ragged hierarchies

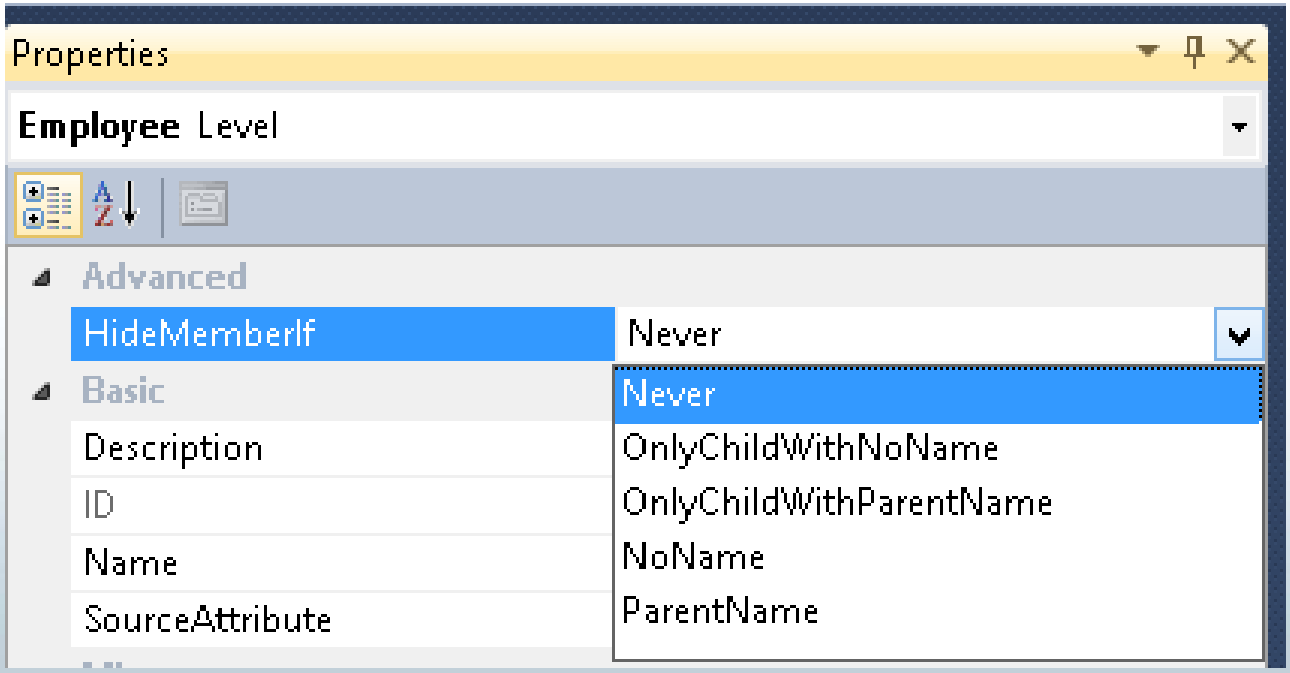

Employee Hierarchy

- Attribute Properties
- Hierarchy Properties
  - HideMemberIf

Employee Hierarchy

- Attribute Properties
  - MembersWithDataCaption: (free text)
    * (Parent) will add label to Parents showing as children of themselves.
  - MemberswithData: NonLeafDataHidden
    Hide Parents showing as children of themselves.
  - Naming Template: (free text)
    Employee Level * or CEO, VP, Manager, Associate
    Add custom label to each level

Employee Hierarchy

- Attribute Properties
  - RootMemberIf:

    Determines how the **root** or **topmost members** of a **parent-child hierarchy** are identified

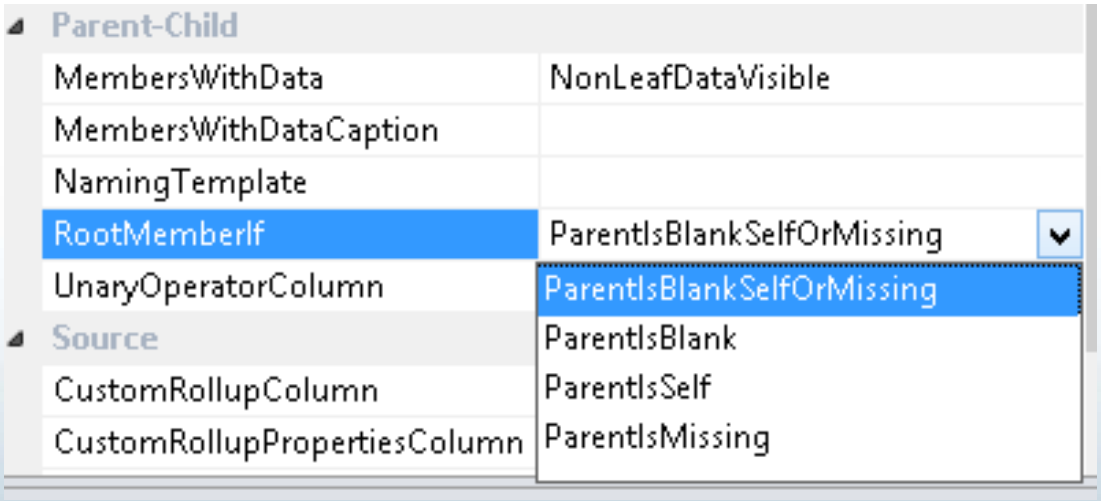

Employee Hierarchy

- Attribute Properties
  - RootMemberIf

    ParentIsBlank: Only keys in the dimension that have a parent key of null, zero, or empty string are the top most level in your parent-child hierarchy

    ParentIsMissing: Only keys in the dimension where no key can be found are the top most level in your parent-child hierarchy

    ParentIsSelf: Only keys in the dimension where the key is equal to itself are the top most level in your parent-child hierarchy

    By default it's set to "ParentIsBlankSelforMissing"," which means essentially "any of the above"

Source: http://www.databasejournal.com/features/mssql/article.php/10894_3749671_2/Dimension-Attributes--Introduction-and-Overview-Part-IV.htm

Demo:
Employee Hierarchy

Date Hierarchy

Date Hierarchy

- Multiple Calendars
  - Calendar Year Drilldown
  - Fiscal Year Drilldown
- Month Sorting
  - Key Columns: (Collection) Year, MonthNumberofYear
  - Orderby: Key
- Dimension & Attribute Properties
  - Dimension Type: Time
  - Attribute Type: Month, Quarter, Year,

Date Hierarchy

- Attribute Relationships
  - Rigid vs. Flexible

Demo:
Date Hierarchy

Dimension and Attribute Types

Dimension Types

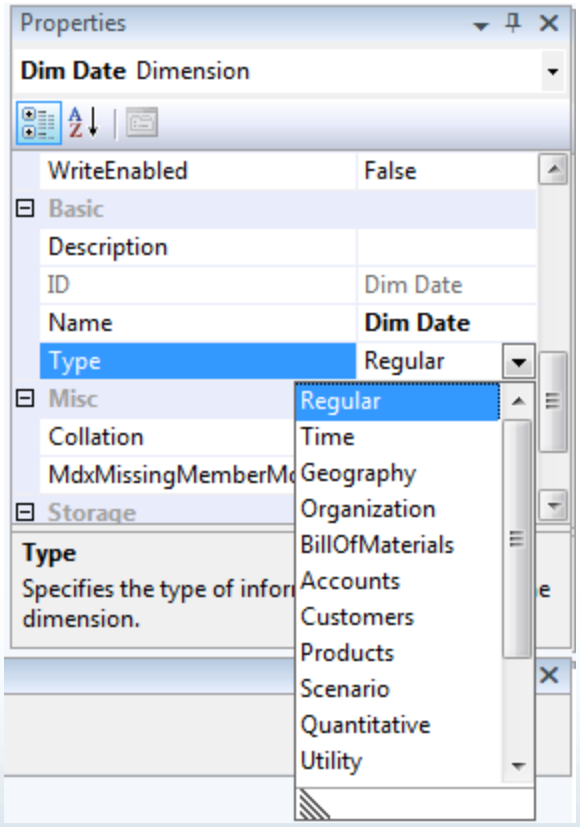

Attribute Types

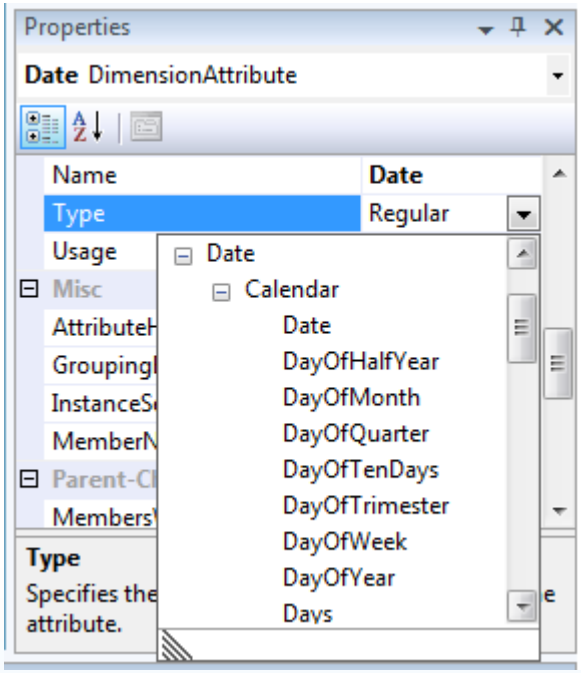

Parent-child hierarchies

Self-referencing relationship or self-join

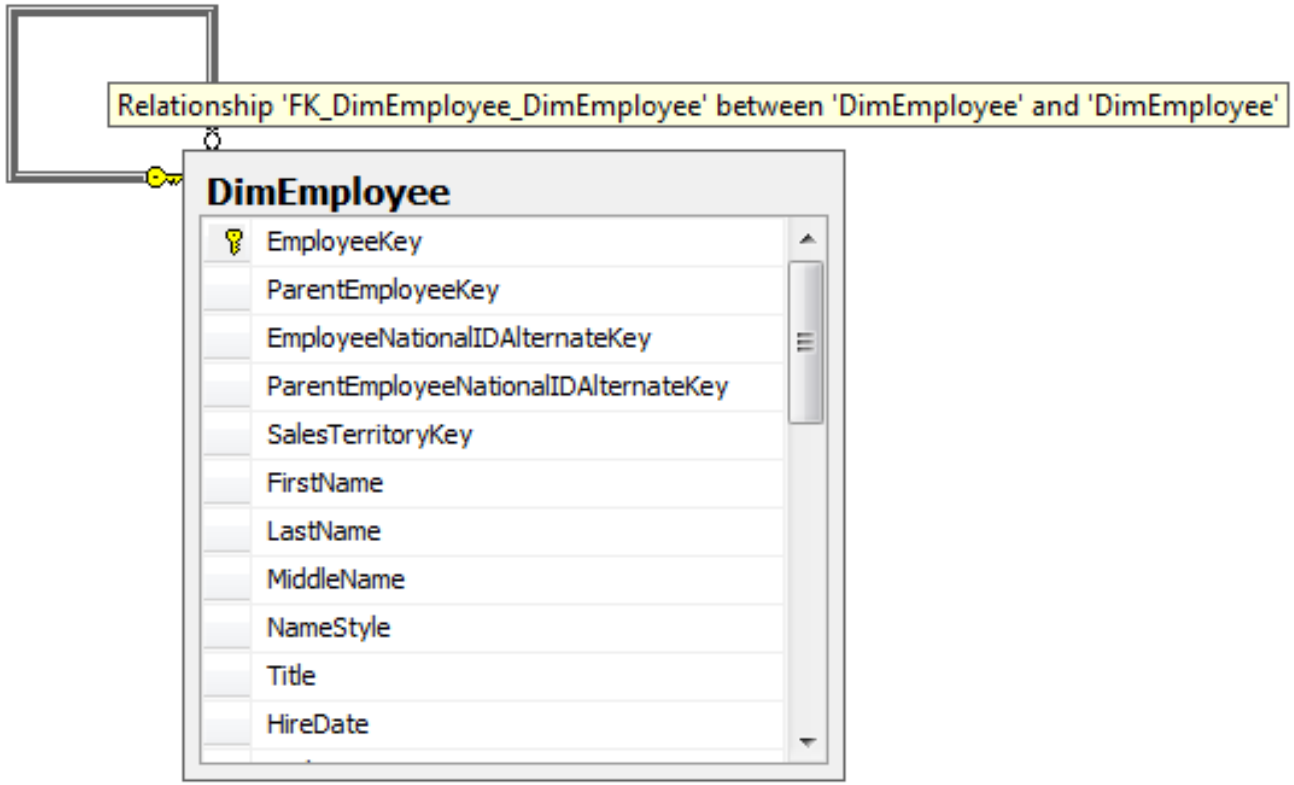

Performance Optimizations

- Aggregations
- Partitions

Aggregations

- Aggregation Design
  - Set of Aggregation Definitions
- Aggregation
  - Pre-aggregation of measure group data at certain granularity of related dimensions
  - Intersections between Dimension table and Fact table based on FK and references

Can be understood somewhat as indexes.

Aggregations

- Aggregation usage

Default

Select to set the aggregation usage setting for the attribute to Default. By using this setting, the designer applies a default rule based on the type of attribute and dimension.

Full

Select to set the aggregation usage setting for the attribute to Full. By using this setting, every aggregation for the cube must include this attribute or a related attribute that is lower in the attribute chain. The Full aggregation usage setting should be avoided when an attribute contains many members. If specified for multiple attributes or attributes that have many members, this setting might prevent aggregations from being designed because of excessive size.

None

Select to set the aggregation usage setting for the attribute to None. By using this setting, no aggregation for the cube can include this attribute.

Unrestricted

Select to set the aggregation usage setting for the attribute to Unrestricted. By using this setting, no restrictions are put on the aggregation designer; however, the attribute must still be evaluated to determine whether it is a valuable aggregation candidate.

Set All to Default

Select to set the aggregation usage settings for all attributes to Default.

Partitions

- Container for a measure group or part of it.
- Commonly partitioned by date, month, or year

Can be understood somewhat table partitioning.

Demo:
Adding Aggregations and Partitions

Considerations

- Query performance
- Storage
- Processing time

Data Types

Data Type	Query Perf	Storage	Processing
integer, bigint			
String, GUIDs			

Flexible vs. Rigid Attribute Relationships

Relationship Type	Processing Option	Query Perf	Storage	Processing
Flexible	Incremental			
Rigid	Incremental			

- Process Full will re-compute aggregations for both Flexible & Rigid

Attribute properties

Property	Query Perf	Storage	Processing
IsAggregatable = False *No default "All" attribute hierarchy*			
AttributeHierarchyOptimizedState = False *No Aggregations*			
AttributeHierarchyOrdered = False *No Sorting*			
Estimated Count <> 0 *Member count stats*			
OrderBy = Key (integer) *Order column*			
DiscretizationMethod <> None			

Aggregations

Query Perf	Storage	Processing
		

…too much of a good thing is not good.

Partitioning

Query Perf	Storage	Processing
		

…too much of a good thing is not good.

Summary

Summary

- Hierarchy Types & Classifications
  - Natural
  - Unnatural
  - Balanced
  - Ragged
  - Parent-Child
- Analysis Services
  - Attribute hierarchies
  - User hierarchies
- Hierarchy Performance Drivers
  - Attribute Relationships
  - Attribute Properties

Summary

- Performance optimizations
  - Aggregations
  - Partitions
- Optimizations Impact
  - Query
  - Processing
  - Storage

Contact Info

Jose Chinchilla

jchinchilla@agilebay.com

@sqljoe

http://linkedin.com/in/josechinchilla

http://www.agilebay.com

http://www.sqljoe.com

http://amazon.com/author/josechinchilla

PASS

Downloads

Downloads

Slides, Demos & Scripts:

http://bit.ly/18ckkiD

Q&A

Session Evaluations

Your feedback is important and valuable.

3 ways to access

Submit by 5pm Friday Oct. 18 to **WIN** prizes

Go to **passsummit/evals**

Download the GuideBook App and search: **PASS Summit 2013**

Follow the QR code link displayed on session signage throughout the conference venue and in the program guide

PASS

Thank you

for attending this session and the
2013 PASS Summit in Charlotte, NC