

# Introduction to SQL Server 2000 Analysis Services

**John H. Miller**

**Product Manager - SQL Server Group  
Business Intelligence Practices Team  
Microsoft Corporation**




# Agenda

- **Microsoft's Market Perspective and Fit Within Business Intelligence (BI)**
- **Overview of key OLAP & BI concepts**
- **Overview of Analysis Services**
  - **Intuitive Design**
  - **Storage modes & aggregation techniques**
  - **Use of partitions**
  - **Data Access**



# Vision - BI For Everyone


- Affordable
  - Scalable
  - Fast “Go To Market”
  - Ubiquitously accessible
  - Embeddable
    - Great Plains
    - Manugistics
    - Best Software
    - Web Trends
    - Visual Insights
    - Business Objects
    - Cognos
    - etc.
- \* More announcements forthcoming!



# Solution – BI Tools In The Box

- Design the data warehouse / data mart
  - SQL 2000 – Enterprise Manager
  - Visio 2000 – Enterprise Manager
- Populate the data warehouse / data mart
  - SQL 2000 – Data Transformation Services
- Create OLAP Cubes
  - SQL 2000 – Analysis Services
- Query The Data
  - SQL 2000 – English Query
  - Microsoft Office
    - Excel & Access
    - Office Web Components

# Scalability Proof Points





# AT&T Growth Markets



- **Opportunity**

- Increase market share by effectively targeting new opportunities in their business phone services

- **Solution**

- 1 TB SQL Server DB and OLAP solution designed to support strategic marketing initiatives within AT&T growth markets

- **Benefits:**

- Accurate profiling, customer insight leading to increases in both revenues and sales productivity

- **Results:**

- 300% ROI in six months!!!

# PCS Health Systems



- **Opportunity:**

- Provide information to finance group coupling claim information with cost and revenue information

- **Solution:**

- Use SQL Server (DTS, OLAP & RDBMS) to offload queries from DB/2 data warehouse to a 1.3 TB claims data mart

- **Benefits:**

- Instead of waiting days for their answer, finance users now get results in less than five minutes

- **Results:**

- IT able to provide better customer service. Improved level of service has led to approval of 6 additional data mart projects

# Redefining OLAP Scalability

## T<sup>3</sup> (pronounced T-Cubed)

### Goal:

- Demonstrate the scalability of Analysis Services
  - Build cube from 1TB+ of source data
  - Describe the techniques used to operate at that scale
- Use the cube to demonstrate high speed queries
- Proof-of-concept system

**Address real business issues:  
schema, data, goals**

Results audited by Winter Corp.

see <http://www.microsoft.com/SQL/techinfo/terabytecube.htm>



# T<sup>3</sup> Storage Requirements

## Table storage (relational)

Table	Million Rows	GBytes
detail_brand_*	1,030	163.4
detail_prodmod_*	20	3.2
detail_subcat_*	11	1.7
detail_upc_*	4,881	793.8
month_brand_*	295	47.1
month_prodmod_*	5	0.8
month_subcat_*	3	0.4
month_upc_*	1,429	225.3
<i>Total</i>	7,674	1,236

**7.7 Billion**

## Cube storage

Cubes	GBytes
Week_Brand	17.8
Week_Class	0.2
Week_Subgroup	0.0
Week_Item	434.7
Month_Brand	4.9
Month_Class	0.1
Month_Subgroup	0.0
Month_Item	24.5
Market Research	482

**1.2 TB**

**471 GB**

**39%**

# T<sup>3</sup> Performance

- Processing

- 7.7 billion rows, 50 hours
- **153 million rows/hr**
- 42K rows/sec
- 60-70% CPU utilization

- Querying

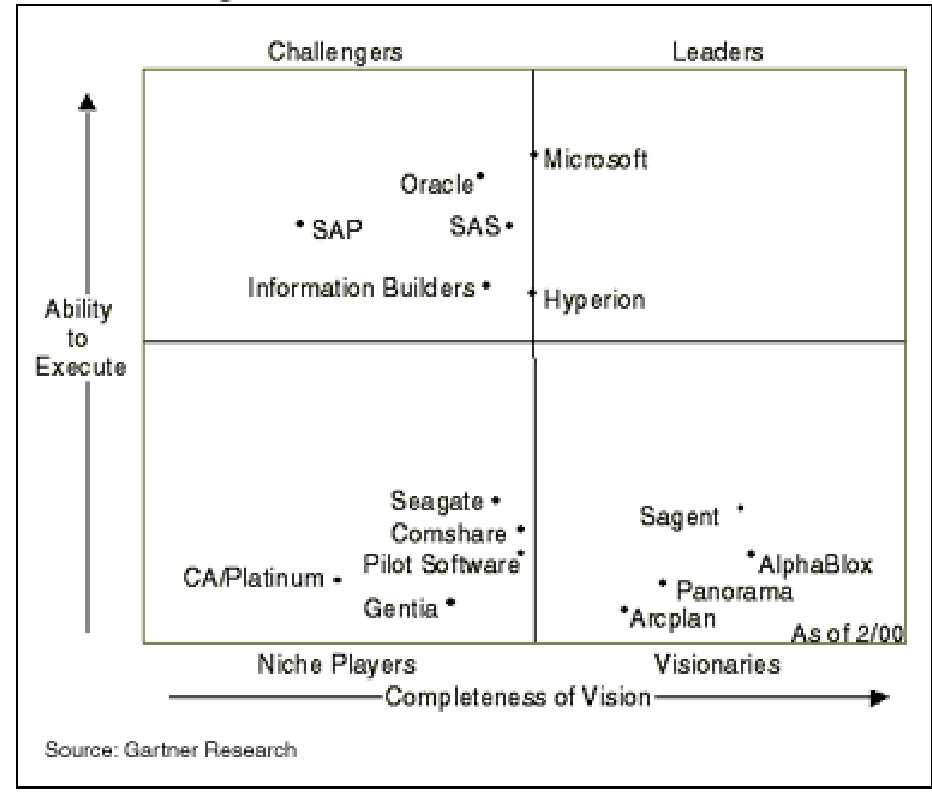
- 50-user workload, 1350 queries, 30-sec think time
- Cold cache
- **Median response 0.08 sec, mean 1.2 sec**
- **Low CPU load - didn't have enough queries running simultaneously!**



# Microsoft Dedication To BI

- Recognized industry leader
- Committed, financially stable vendor
- Significant, ongoing investment in BI

Figure 4  
BI Platform Magic Quadrant



# Agenda

- **Microsoft's Market Perspective and Fit Within Business Intelligence (BI)**
- **Overview of key BI concepts**
- **Overview of Analysis Services**
  - **Intuitive Design**
  - **Storage modes & aggregation techniques**
  - **Use of partitions**
  - **Data Access**





# OLTP Differs From BI

- **OLTP supports**

- Streamlining operations
- Real time production systems
- Current, changing data
- Granular – Transactional

- **BI supports**

- Analyzing Operations / Improving Decision Making
- Consistent, heterogeneous data
- Voluminous, historical, stable data
- Summarized data

- **BI has different design & storage requirements from OLTP**

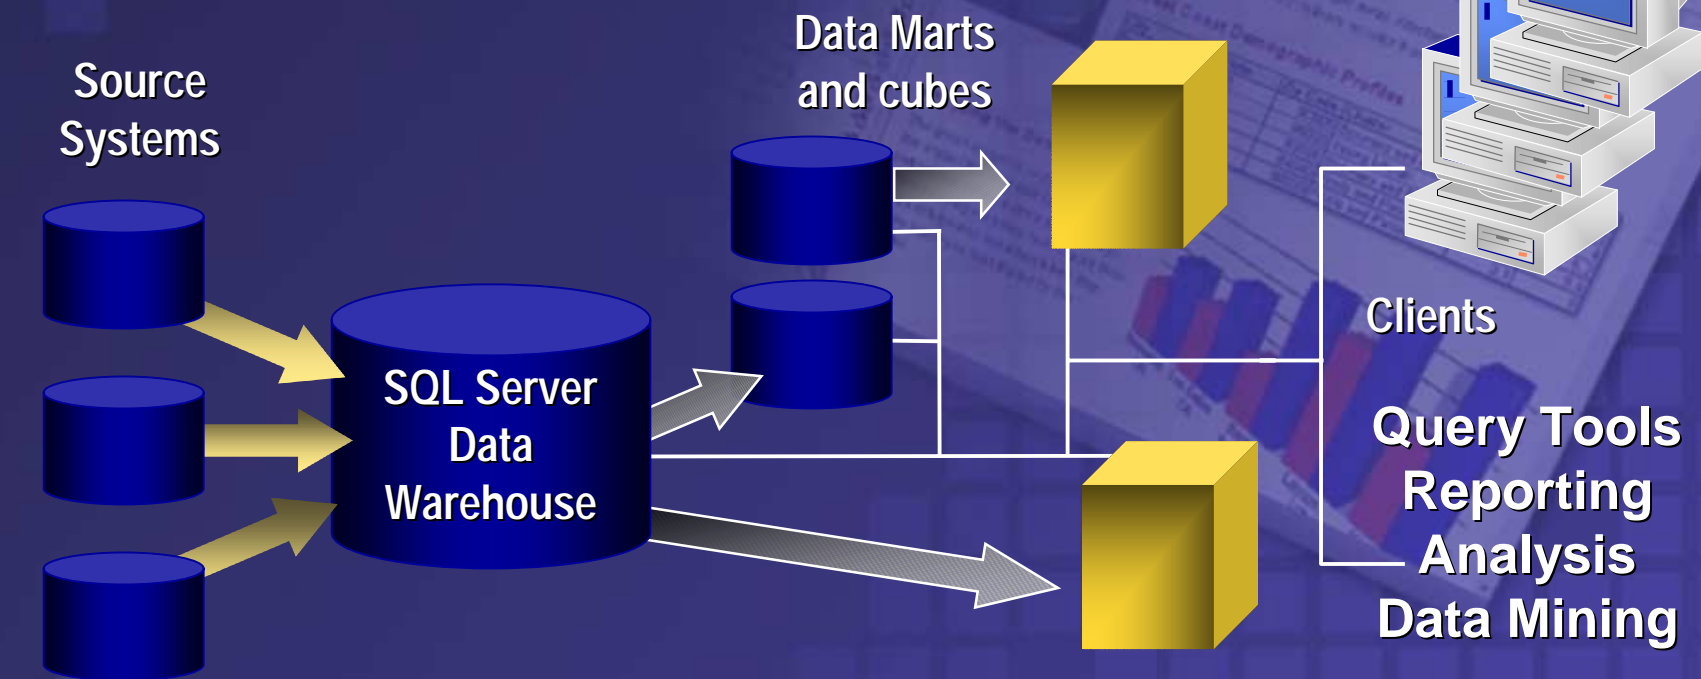
**Operating  
Business**



**Managing  
Business**

# Classical BI Architecture

## Elements of the process



Design the Data Warehouse → Populate Data Warehouse → Create OLAP Cubes → Query The Data

# BI Architectural Goals

- Provide platform to deliver a great user-analyst experience
  - With data that is consistent, centralized and easily accessible
  - Without getting in the way of OLTP systems
- Ability to incorporate data from internal or external sources - regardless of format or platform
- Agile so that it can adapt to changes in the business

# BI User Requirements

- **Support decision making – about managing & planning**
  - How/what/when/why/where of a business
  - Facilitate queries without hindering operational systems performance or having to change the design
  - Provide centralized repository of consistent data
  - Answer complex queries quickly
  - Enable data mining to discover new relationships in data
- **Provide different levels of analysis**
  - View data from many perspectives
  - Easily navigate from summary to detail
- **End user acceptance and usage is the true measure of success**



# What is OLAP?

- **OnLine Analytical Processing**

- It's a cool way of cheating that enables you to get queries answered incredibly fast
- OLAP aggregates data (it pre-summarizes data) across all dimensions

Example:


by MO, QTR, YR

by Country, State, City ...

etc...

- **Basic argument:**

“Why read through each and every detailed transaction to get an answer when the question can be answered more quickly using summary level data”




# Why Use OLAP With DWHS?

- **OLAP is an enabling technology that supports dynamic analysis**
  - **Intuitive multidimensional model**
  - **Fast response times against huge databases**
  - **Offers complete syntax for expressing analytical queries and business logic**
  - **Optimizes the use of network resources as well as Internet/Intranet deployments**

# Data In The Hands Of Users

- **Pivoting**
  - Swapping page/row/column layout
- **Slicing**
  - Select specific dimension members on an axis
- **Drilldown**
  - Navigate from summary to detail data
- **Drill through**
  - Retrieve granular data from Fact Table
- **Calculations**
  - Adding derived dimension members or measures
- **Visualization**
  - Charting, mapping, etc.



# Agenda

- **Microsoft's Market Perspective and Fit Within Business Intelligence (BI)**
- **Overview of key BI concepts**
- **Overview of Analysis Services**
  - **Intuitive Design**
  - **Storage modes & aggregation techniques**
  - **Use of partitions**
  - **Data Access**





# **OLAP – Enabling Analysis**

## **Data Sources**

- **Analysis Services works as well with data coming from Oracle, DB/2 and others as it does with SQL**
- **Analysis Services supports all data sources accessible via ODBC (including flat files)**
- **Analysis Services also supports newer generation OLE-DB data sources**
- **DTS and SQL capable of accepting XML**

# OLAP – Enabling Analysis

## Mapping warehouse to cube

- **Dimension tables used to build OLAP dimensions**
  - Shared dimensions enable analysis to take place across cubes
  - Private dimensions support different hierarchies and aggregation rollups
- **Virtual cubes join two or more physical cubes into one logical cube**
  - Requires one or more shared dimensions

# OLAP – Enabling Analysis

## Dimensional design

- **Dimensional modeling easy to accomplish in Analysis Services**
  - Intuitive design palette enables users to model their data and view the results from the same pane
  - Wizards exist to help guide users through common and complex tasks
- **On-line tutorial exists within product to guide new users through basic and advanced concepts**



**Designing dimensions**

# OLAP – Enabling Analysis

## Cube design

- **Public / Private Dimensions available**
- **Quantitative columns in fact table become measures**
- **Calculated members support creation of:**
  - **Derived entities in a dimension**
    - New products or regions
  - **Derived measures**
    - $\text{Income} = \text{revenue} - \text{expense}$






# Designing Cubes

# OLAP – Enabling Analysis

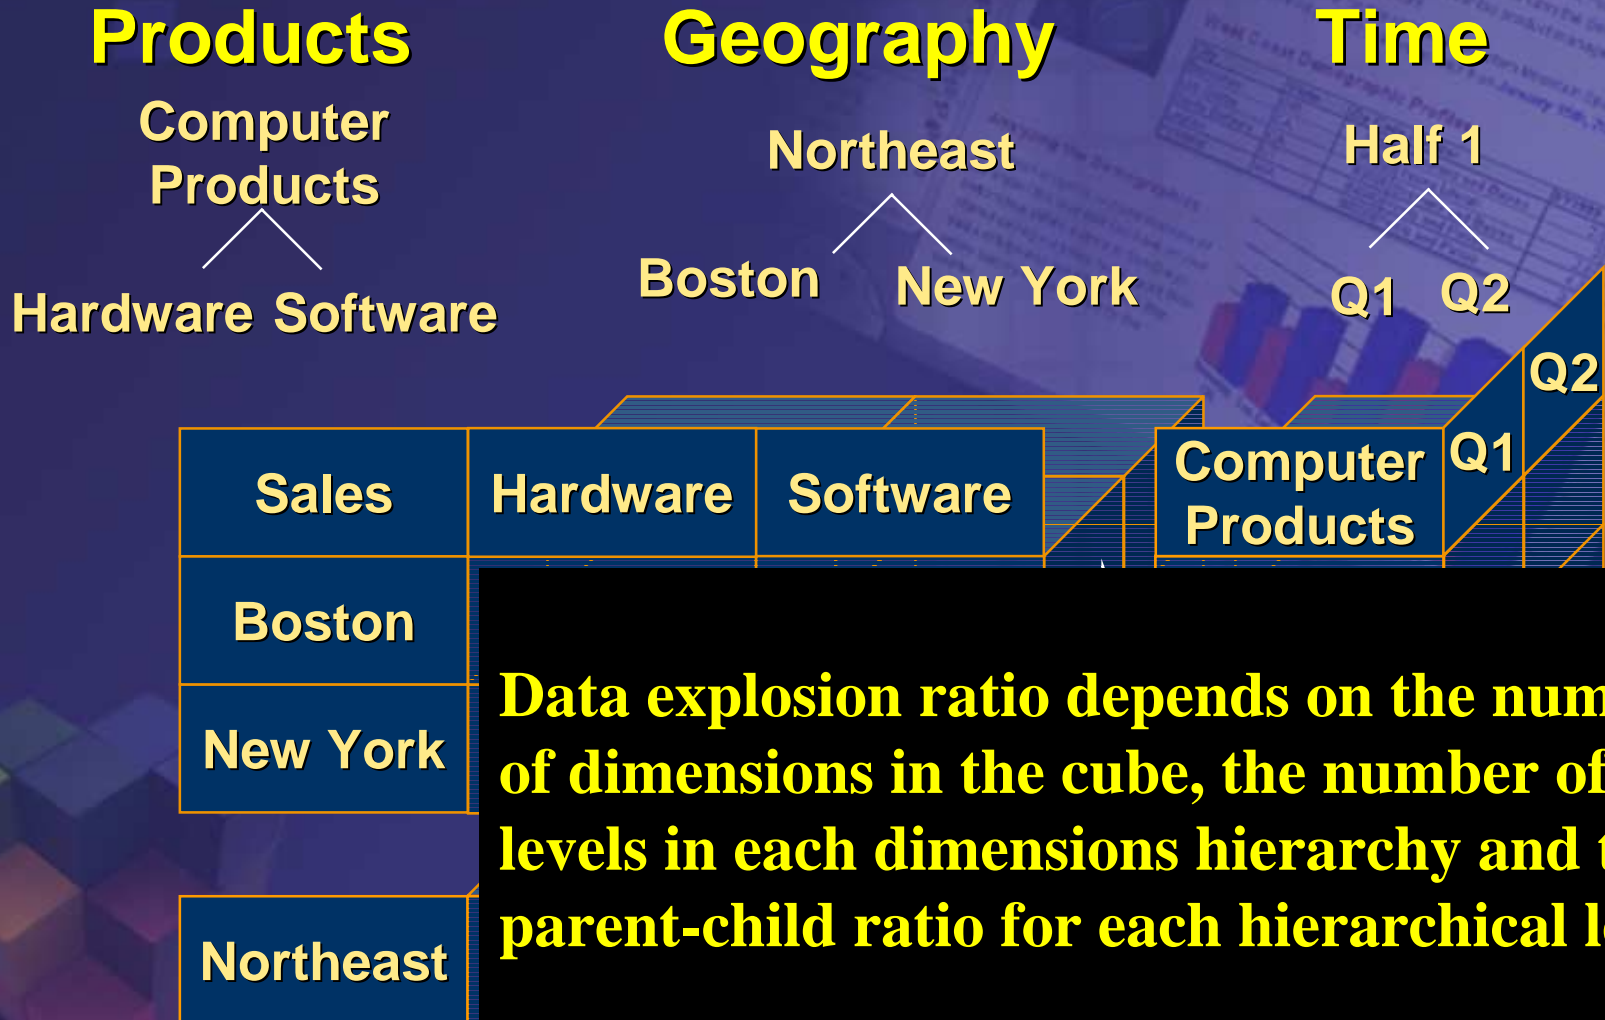
## Flexible Storage

- Debates between MOLAP and ROLAP vendors obscure customer needs
- Analysis Services supports MOLAP, ROLAP, and HOLAP and offers seamless integration of all three
- Users & applications see only cubes



# OLAP – Enabling Analysis

## Dealing with data EXPLOSION!



# OLAP – Enabling Analysis

Dealing with data EXPLOSION!

**Persist only base aggregates**

- Min
- Max
- Sum
- Count
- Distinct Count

\* other, more complex aggregates are derived from base aggregates

# OLAP – Enabling Analysis

## Dealing with data EXPLOSION!

Automagically eliminate **all** sparse data from the cube

	Furnaces	Air Cond.	HVAC Products
Phoenix		150	250
Anchorage	250		350
USA	350	250	600

Data cells: 4

Sparce cells: 2

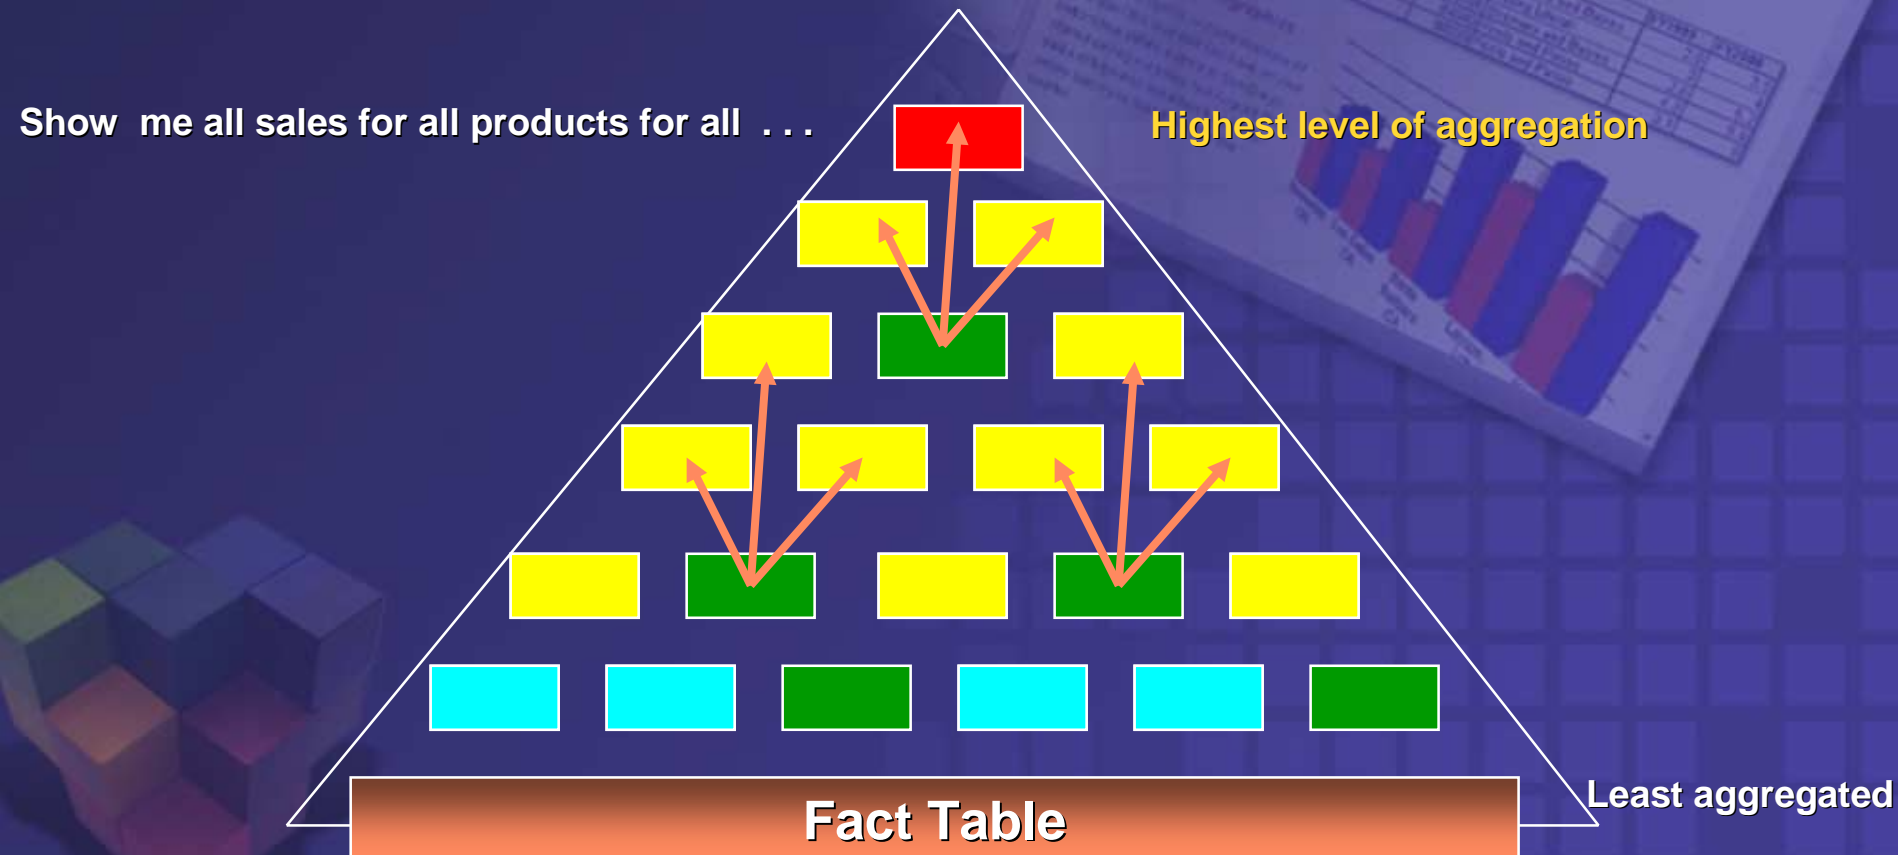
Ratio: 0.50



# OLAP – Enabling Analysis

## Dealing with data EXPLOSION!

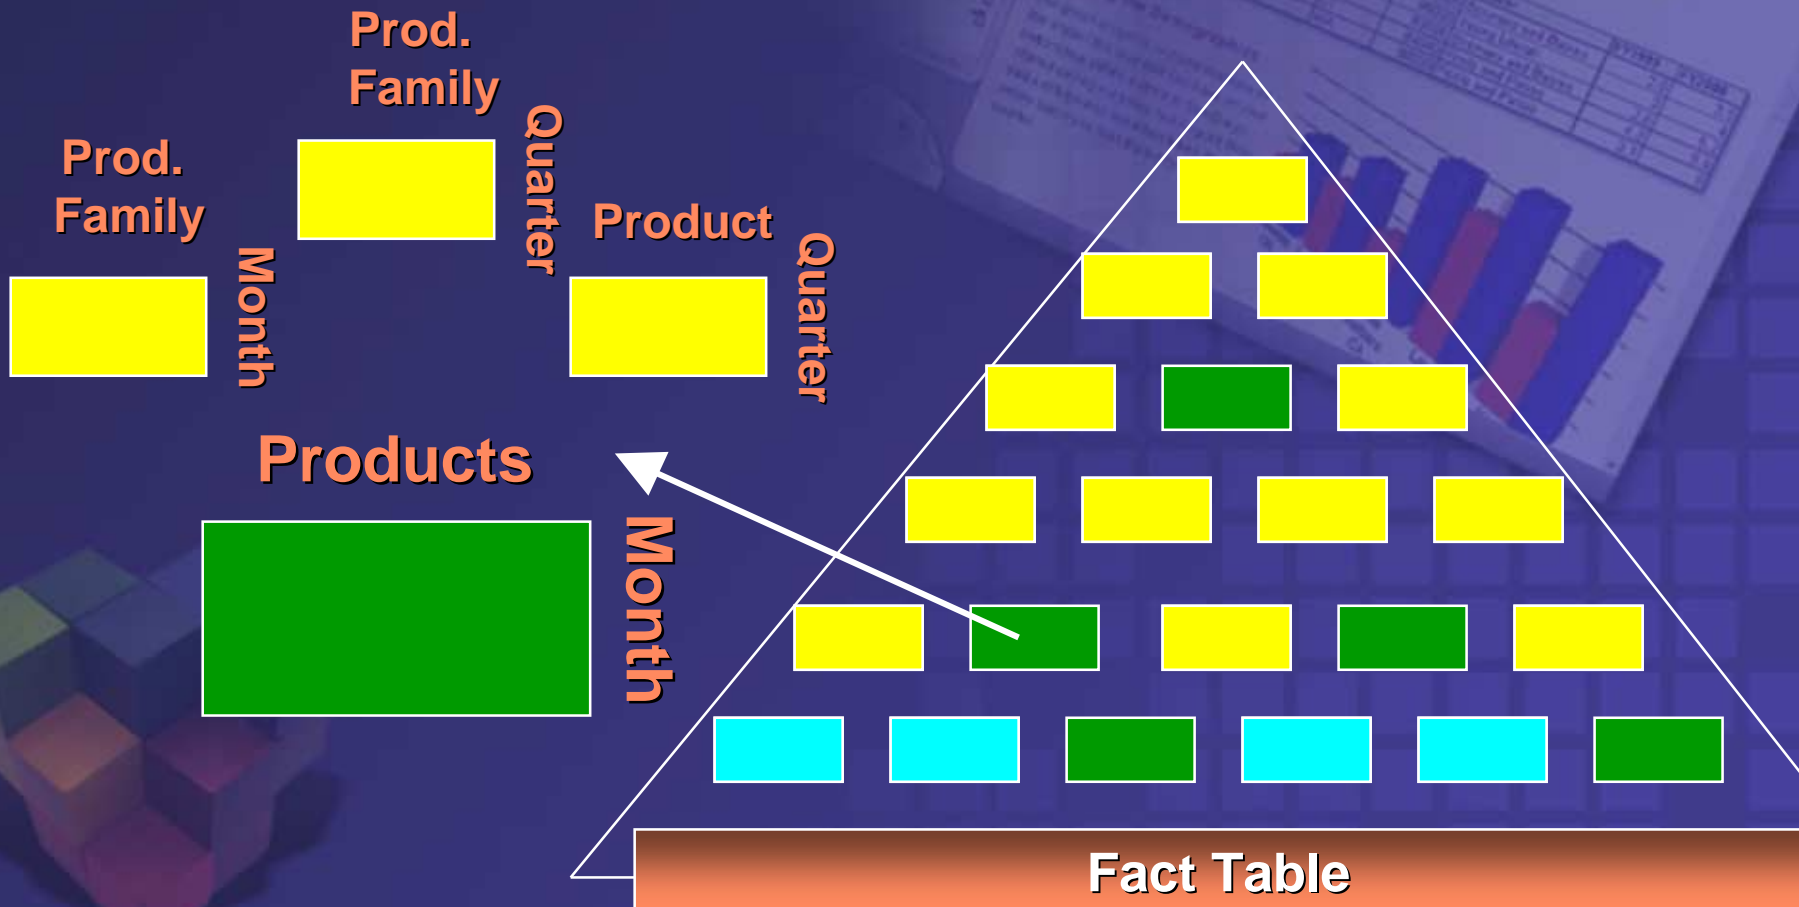
Do only partial pre-aggregation



# OLAP – Enabling Analysis

## Dealing with data EXPLOSION!

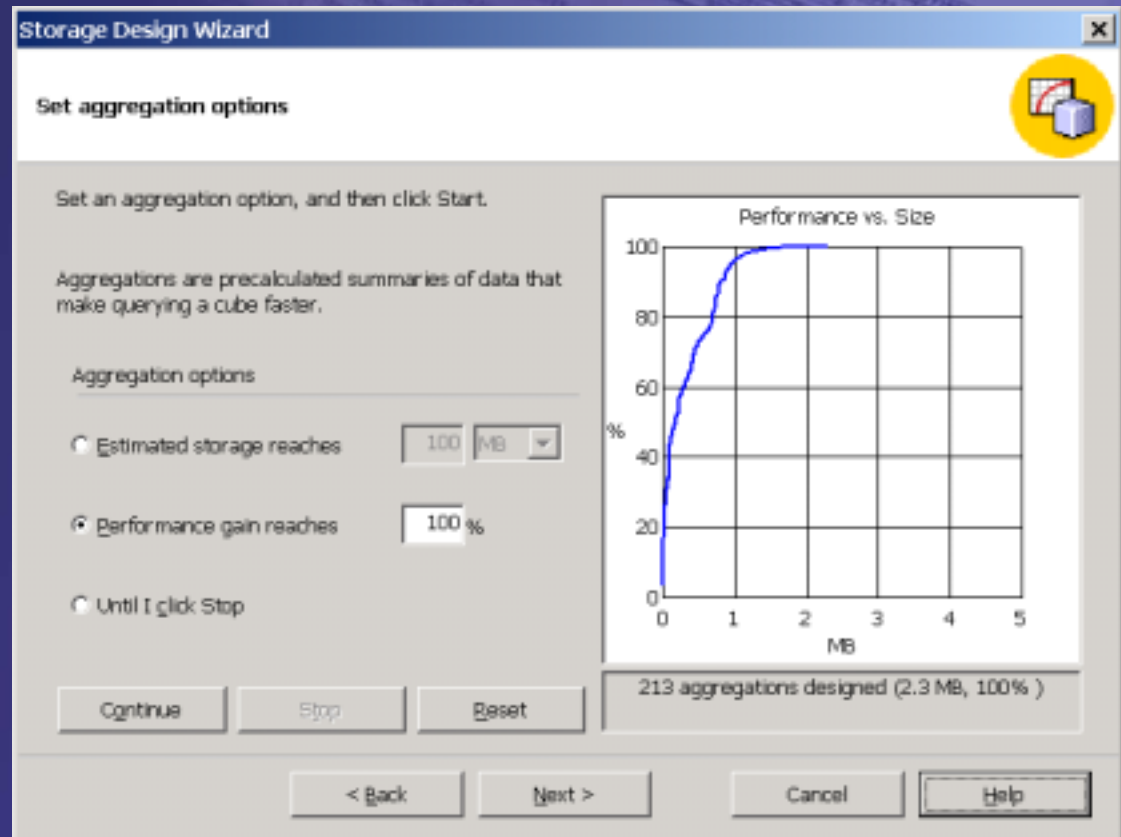
Do only partial pre-aggregation



# OLAP – Enabling Analysis

## Dealing with data EXPLOSION!

- Storage Design wizard finds the “80-20” rule in the data
  - The 20 percent of all possible pre-aggregations that provide 80 percent of the performance gain
  - Analyses the level counts for each dimensions as well as the parent-child ratios for each level



The screenshot shows the 'Storage Design Wizard' dialog box with the 'Set aggregation options' step. The 'Performance gain reaches' option is selected, set to 100%. A graph titled 'Performance vs. Size' shows a curve that rises steeply from 0% at 0 MB and levels off at 100% performance around 2.3 MB. Below the graph, it states '213 aggregations designed (2.3 MB, 100%)'. The dialog includes 'Continue', 'Stop', and 'Reset' buttons, and a navigation bar at the bottom with '< Back', 'Next >', 'Cancel', and 'Help' buttons.

Storage Design Wizard

Set aggregation options

Set an aggregation option, and then click Start.

Aggregations are precalculated summaries of data that make querying a cube faster.

Aggregation options

Estimated storage reaches  MB

Performance gain reaches  %

Until I click Stop

Continue Stop Reset

Performance vs. Size


213 aggregations designed (2.3 MB, 100%)

< Back Next > Cancel Help

# OLAP – Enabling Analysis

## Dealing with data EXPLOSION!

### Final Cube Compression



When using either MOLAP or HOLAP storage modes, Analysis Services is capable of compressing the final cube size



# Designing Storage




# OLAP – Enabling Analysis

## Cube Partitioning

One logical cube

**Product**

Groceries  
Electronics  
Clothing  
Garden  
Automotive



**Geog**

North

South

East

West

Q1

Q2

Q3

Q4

**Time**


# OLAP – Enabling Analysis

## Cube Partitioning

One logical cube

**Product**

Groceries  
Electronics  
Clothing  
Garden  
Automotive



**Geog**

North

South

East

West

Q1

Q2

Q3

Q4

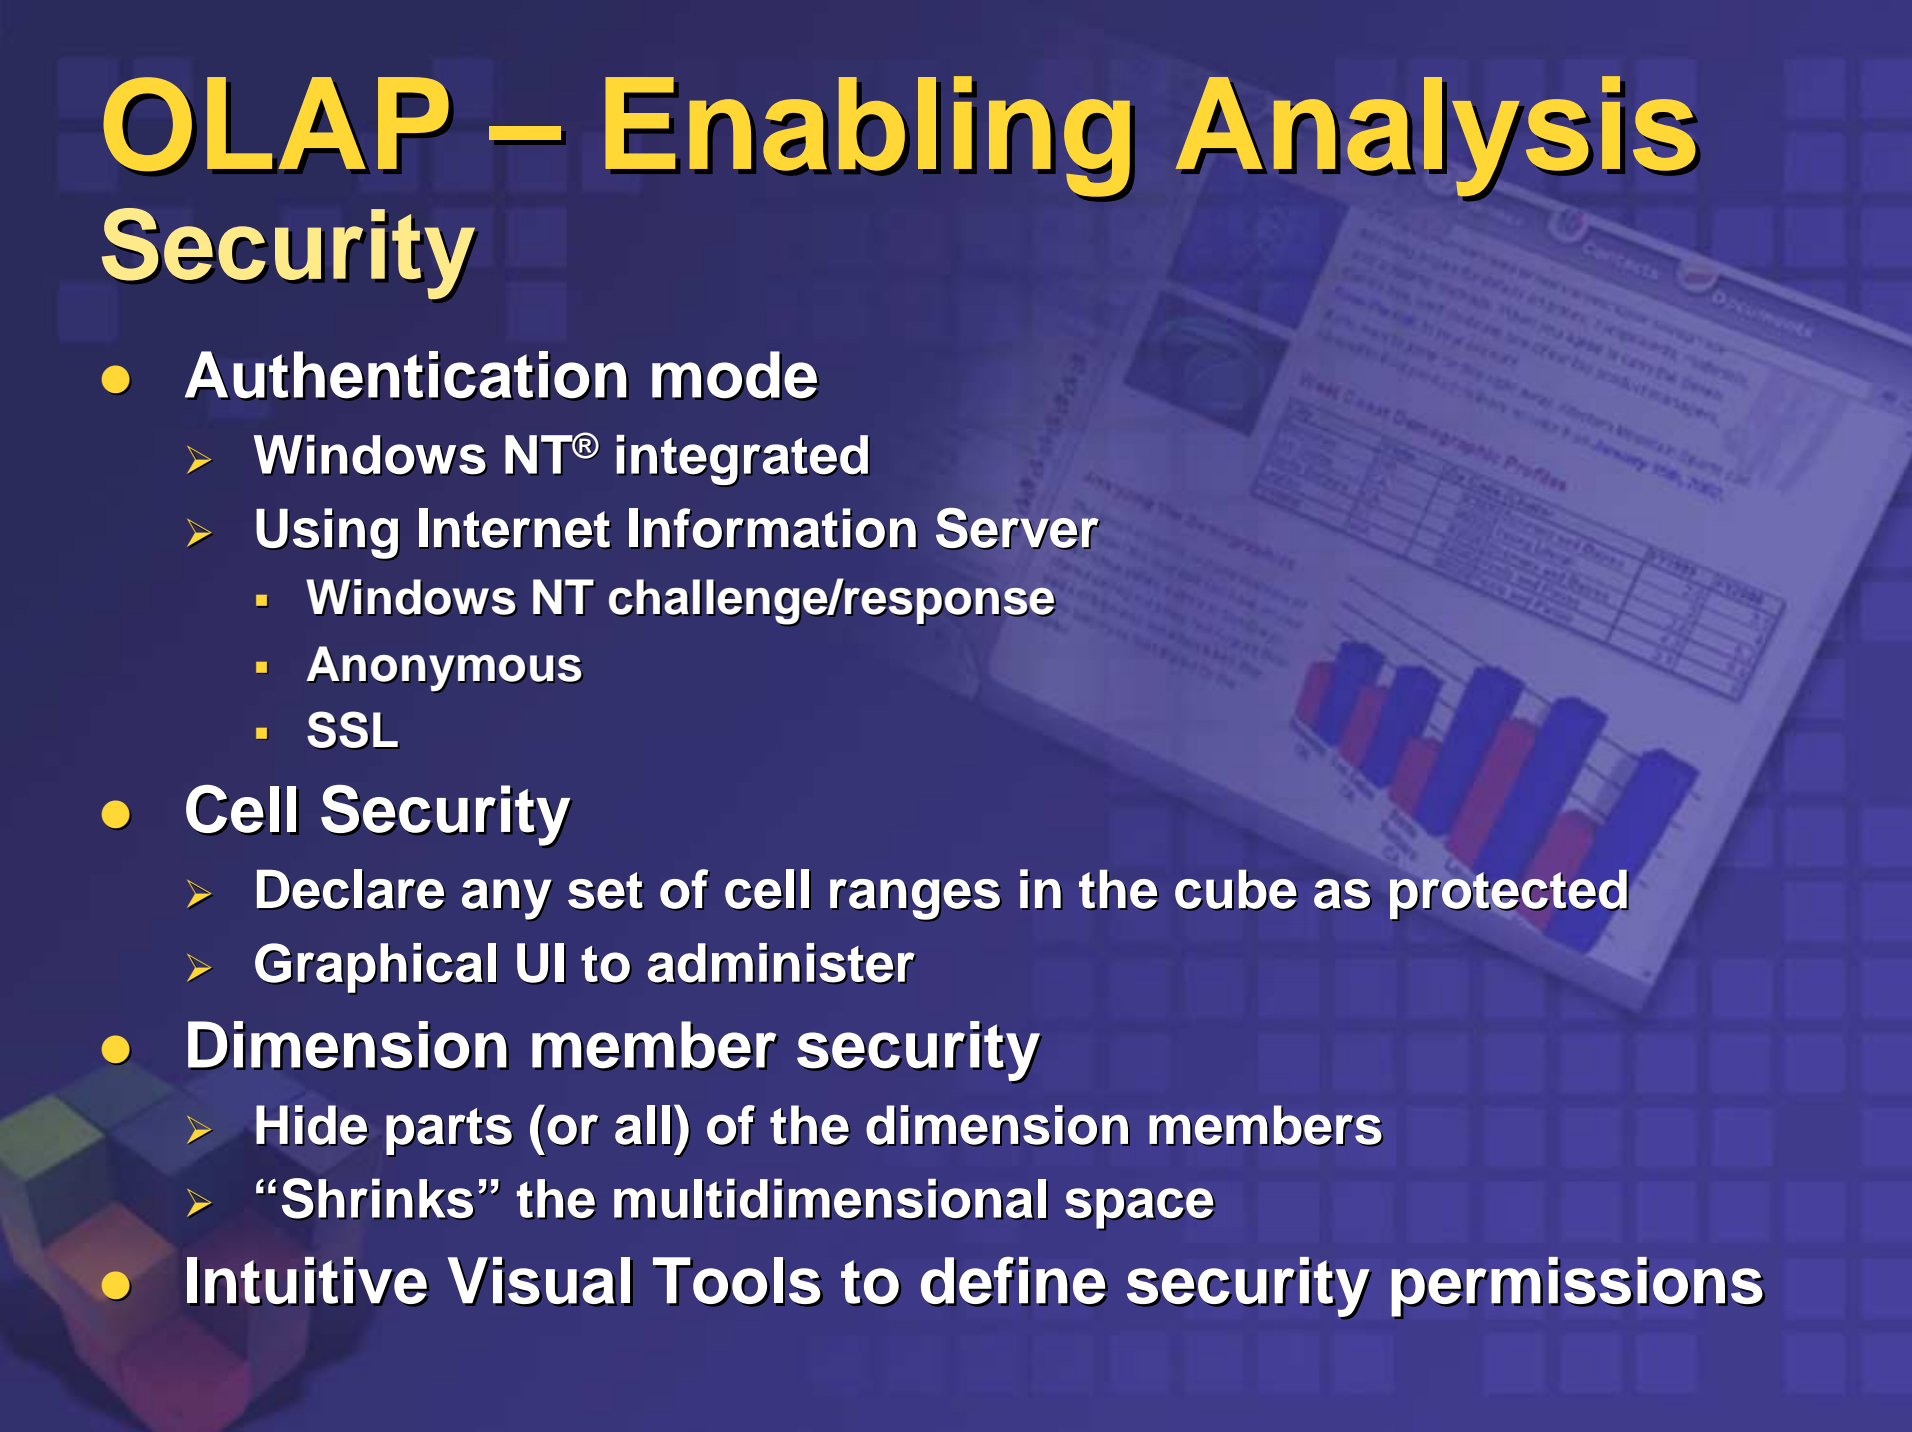
**Time**

# OLAP – Enabling Analysis

## Importance of Partitions

- **Central to Analysis Services**
- **Data Management**
  - Incremental updates
  - Adding/dropping time periods
  - Granular control over aggregation design
  - Enable hybrid cube design (e.g. combination of MOLAP and ROLAP partitions in the same logical cube)
- **Basis for advanced features like Write-back**
- **Important scalability & performance implications**
  - Single server parallelism (partition processing)
  - Multi-server parallel processing
  - More efficient query processing

# OLAP – Enabling Analysis Security

- **Authentication mode**
    - Windows NT® integrated
    - Using Internet Information Server
      - Windows NT challenge/response
      - Anonymous
      - SSL
  - **Cell Security**
    - Declare any set of cell ranges in the cube as protected
    - Graphical UI to administer
  - **Dimension member security**
    - Hide parts (or all) of the dimension members
    - “Shrinks” the multidimensional space
  - **Intuitive Visual Tools to define security permissions**
- 
- The background features a semi-transparent blue-tinted image of a computer monitor displaying a data analysis interface. The screen shows a 'Weeks' section with a table of data, a 'Weeks Demographic Profile' table, and a 3D bar chart with blue and red bars. In the bottom left corner, there is a 3D cube graphic with various colored faces (green, blue, red, yellow).

# OLAP – Enabling Analysis Security

- Security based on MDX expressions
- Cell-level, member-level
- Extends roles concept, MDX filters

## Cell-level security

	Employees	Total salary	Sq Ft
East	50	N/A	16,500
Central	75	\$7,600K	20,625
West	35	N/A	10,500

## Member-level security

	Employees	Sq Ft
East	50	16,500
Central	75	20,625
West	35	10,500



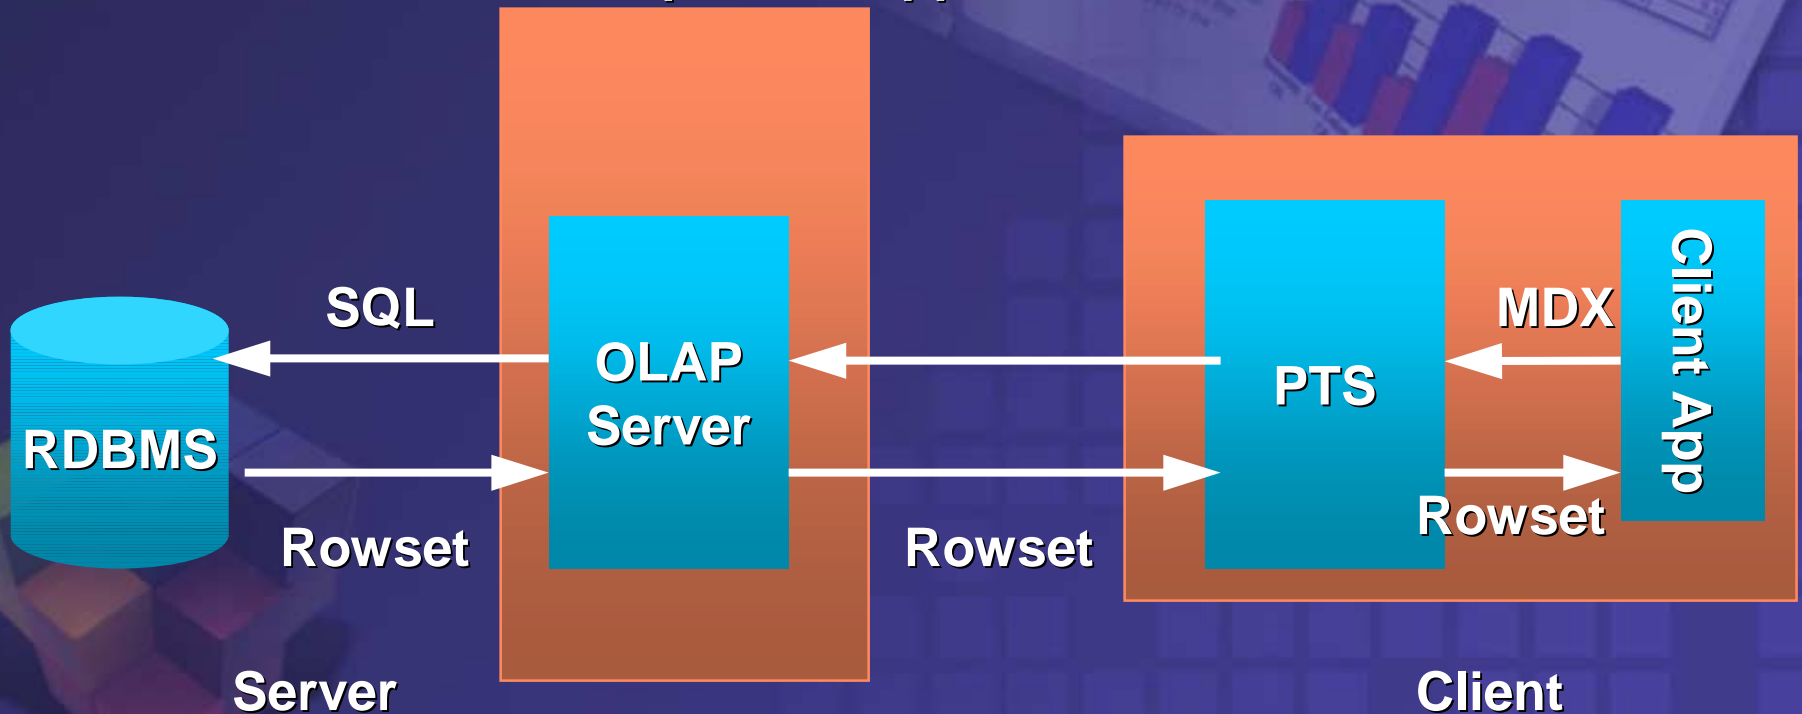


**Security**

# OLAP – Enabling Analysis

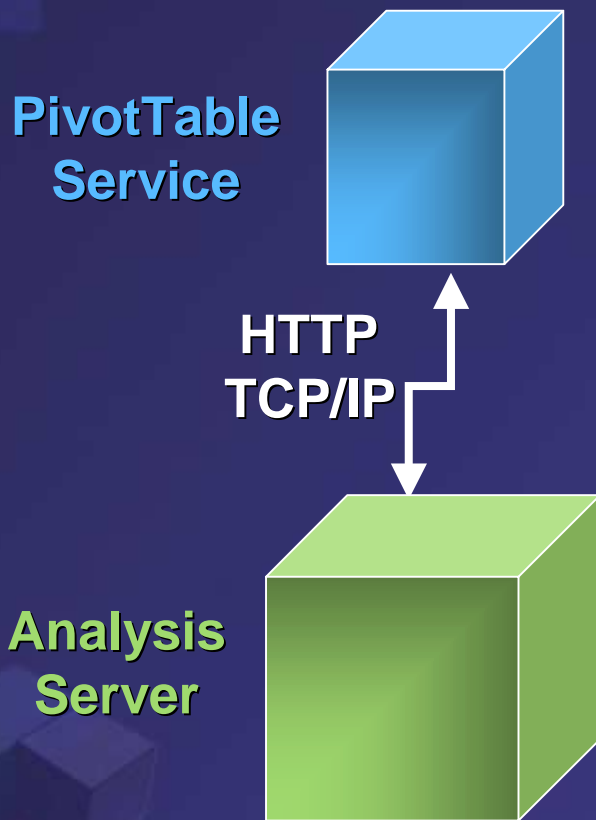
## Drill Through

- View the underlying fact table rows for a cell or aggregation
- Ease of use...
  - Just check the box!
  - Native and transparent support



# OLAP – Enabling Analysis

## Efficient Network Utilization



### Client Features

- multidimensional calc engine
- data & metadata caching (session)
- query management
- client OLEDB for OLAP interfaces

### Server Features

- multidimensional calc engine
- data & metadata caching (multi-user)
- query management
- server DSO admin interfaces
- security
- data refresh
- aggregation management
- ...

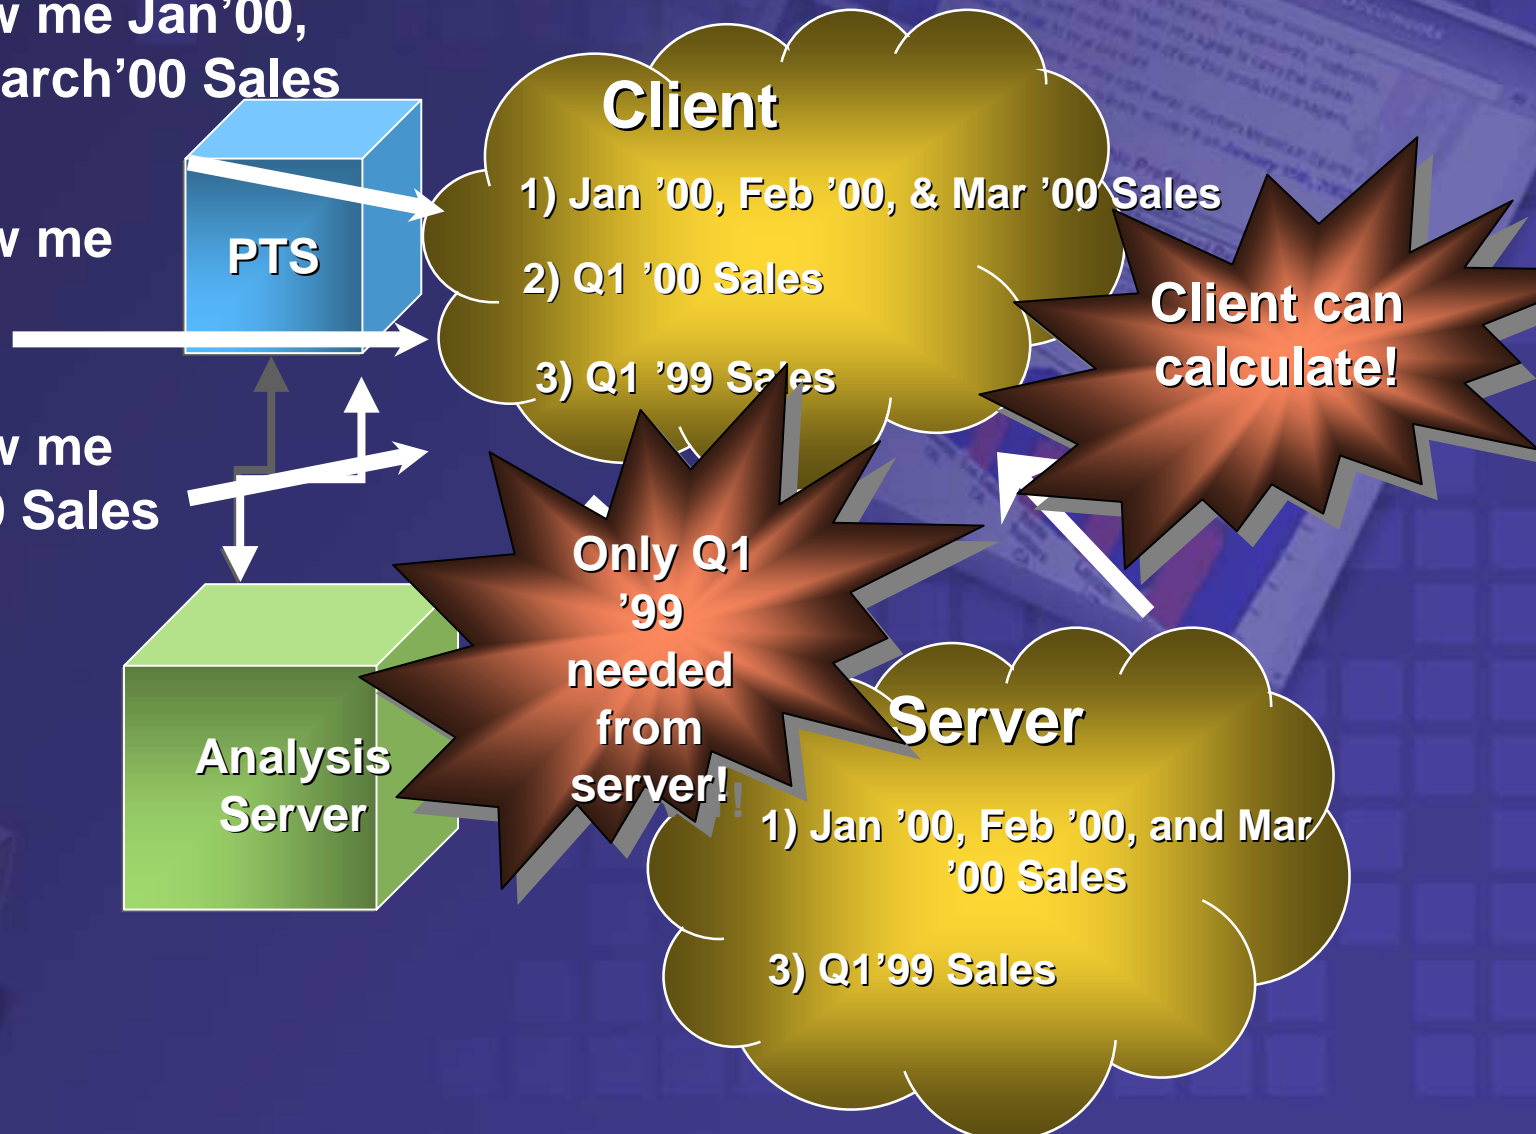
# OLAP – Enabling Analysis

## Cooperative Caching

**Query 1)** Show me Jan'00, Feb'00, and March'00 Sales

**Query 2)** Show me Q1'00 Sales

**Query 3)** Show me Q1'00 & Q1'99 Sales



# **OLAP – Enabling Analysis**

## **Mobile & Internet Access To Data**

- **Ability to take cube with you on the road for mobile analysis**
- **Uses HTTP to pass through firewalls**
- **Uses IIS to provide authentication over the Internet**
- **XML for Analysis – specification and SDK available via web download**
- **Support for PocketPC (Window CE devices)**





# Office & Office Web Components

# Features we didn't have time to cover ...

- Member properties & virtual dimensions (same as attributes in other products)
- Write-back & drill through
- Virtual cubes
- Data Mining
- Actions
- Custom rollups
- Calculated cells
- Dimension architecture
- Usage based optimization

# Solve Prickly OLAP Problems That Plague Other Vendors

- Cube explosion – not uncommon to experience exponential data explosion
- Time needed to reprocess/refresh cube(s)
- Confusion over storage options (MOLAP, ROLAP, HOLAP)
- Inability to handle really complex calculations
- Cost prohibitive to roll out across enterprise

**Analysis Services effectively solves all these problems**

# OLAP Advantages In Summary

<b>Intelligent Aggregations</b>	<ul style="list-style-type: none"><li>◆ Significantly smaller databases</li><li>◆ Faster initial and incremental processing</li></ul>
<b>Flexible Storage Architecture</b>	<ul style="list-style-type: none"><li>◆ Supports MOLAP, ROLAP, and HOLAP equally well</li><li>◆ Application requirements determine storage, not vendor</li></ul>
<b>Ease-of-Use Ease-of-Mgmt</b>	<ul style="list-style-type: none"><li>◆ Fast “To Market” Solutions</li><li>◆ Lower TCO</li><li>◆ Broad accessibility to data</li></ul>
<b>PivotTable Service</b>	<ul style="list-style-type: none"><li>◆ Client-side cache -- improved performance / efficiency</li><li>◆ Mobile/Disconnected analytical support</li><li>◆ Web enabled access</li></ul>
<b>Integration</b>	<ul style="list-style-type: none"><li>◆ Office, SQL Server, .Net Servers, 3<sup>rd</sup> Parties</li><li>◆ Compelling bridge between Excel and SQL Server</li><li>◆ End-to-End data storage, transformation, and analysis</li></ul>

# Microsoft Software for an Agile Business

Fast. Scalable. Flexible. Powerful. Reliable. Enterprise-ready.



knosys®



Microsoft®  
**SQL Server**™



**D A M A N**  
CONSULTING