Building a 10 Terabyte SQL Server Data Warehouse

Pradeep Arora

Solutions Architect, MCSE 2003, MCDBA
Hewlett-Packard

How we use data?

Basic Data Usage Patterns

- Short transactions - OLTP
- Long transactions - Extract Transform Load / Batch Processing
- OLAP / cube processing
- Reporting / Queries
  - Short running – Standard
  - Short running – Cube access
  - Long running - Ad Hoc queries
- Data Mining

Data In The Hands Of Users
End user analysis activities

- 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.
- Data Mining
  - Finding patterns in your data

HP Technologies

Microsoft Technologies

Microsoft's BI Platform

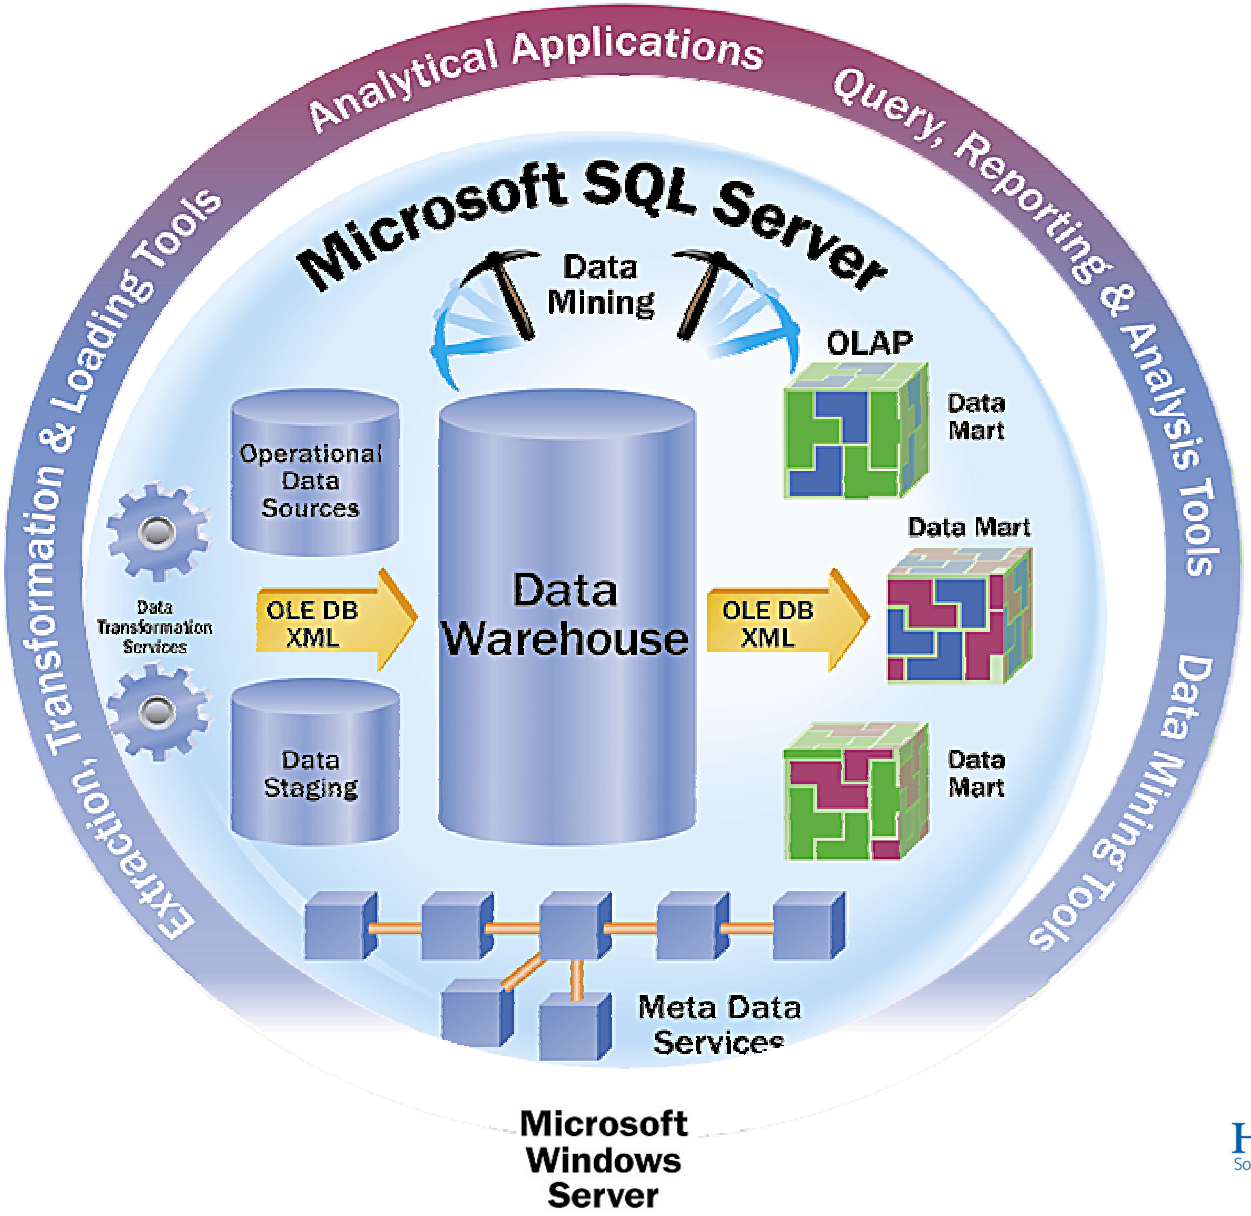

BI Data Flow
Elements of the process

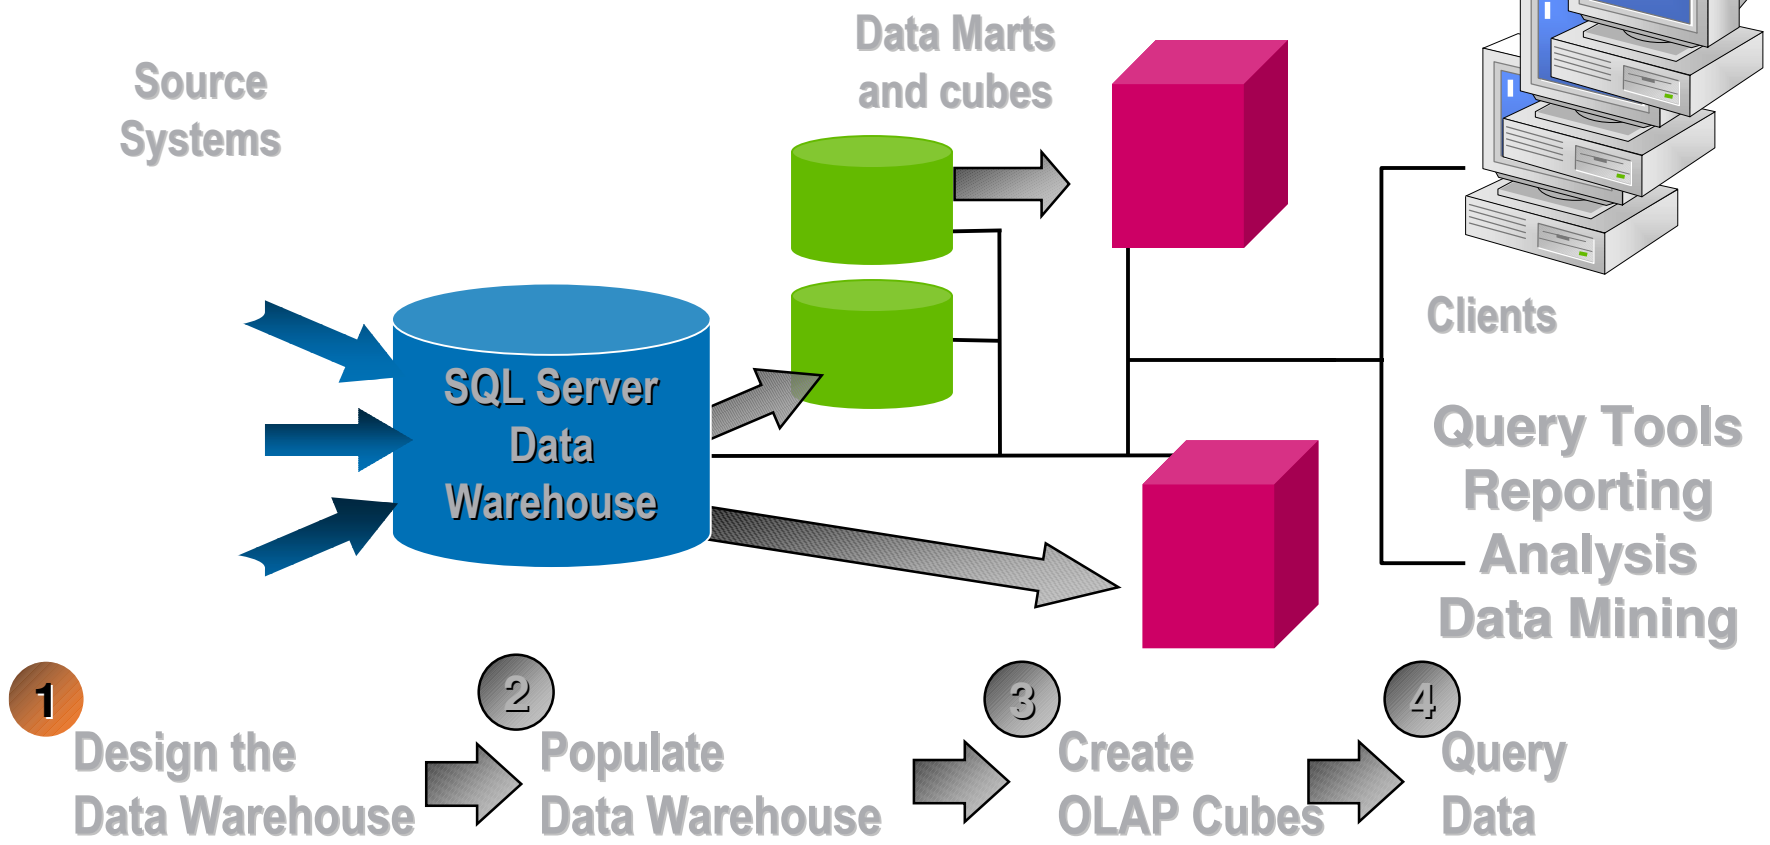

SQL Server 2000 Analysis Services
What are Analysis Services?

- Middle tier for OLAP and Data Mining
- OLAP – Access to the data
  - Easy to select, navigate, and explore data

SQL Server 2000 Analysis Services
Cube Structure – Schema and Tables

- Schema
  - Joined tables in the data warehouse
  - Provides data for the cube
- Fact Table
  - Central table in the schema
  - Numerical data (facts)
  - Provides historical information about operations
- Dimension Tables
  - Additional tables joined to the fact table

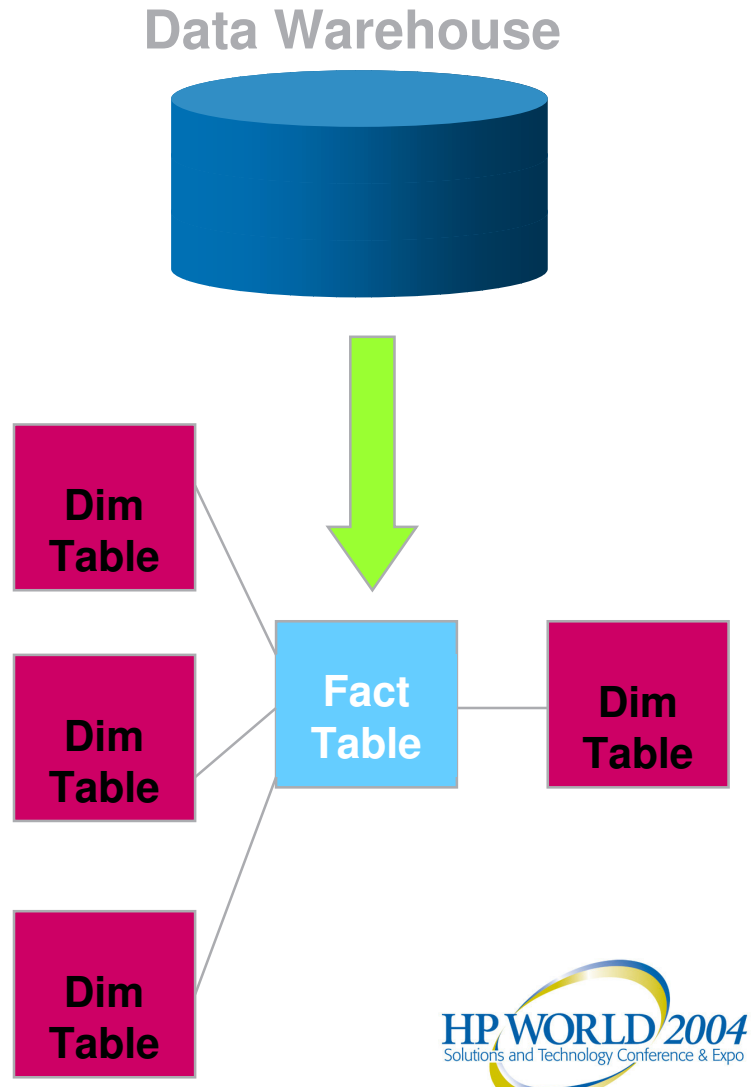

SQL Server 2000 Analysis Services
Cube Structure – Dimensions and Measures

- Measure
  - Values from Fact Table
  - Values that are analyzed and aggregated
- Dimension
  - Values from Dimension Tables
  - Describe a similar set of members to be analyzed
- Levels
  - Hierarchy within a dimension
  - Categorical breakdowns
  - Time levels might be Year, Quarter, Month

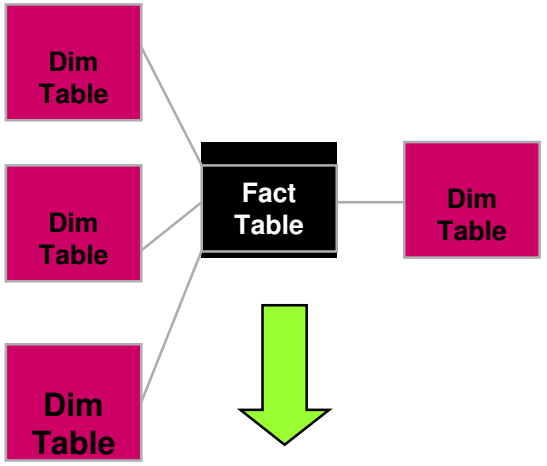

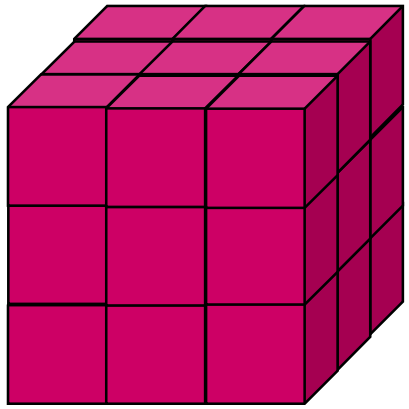

Analysis Services Design
Client/Server Features

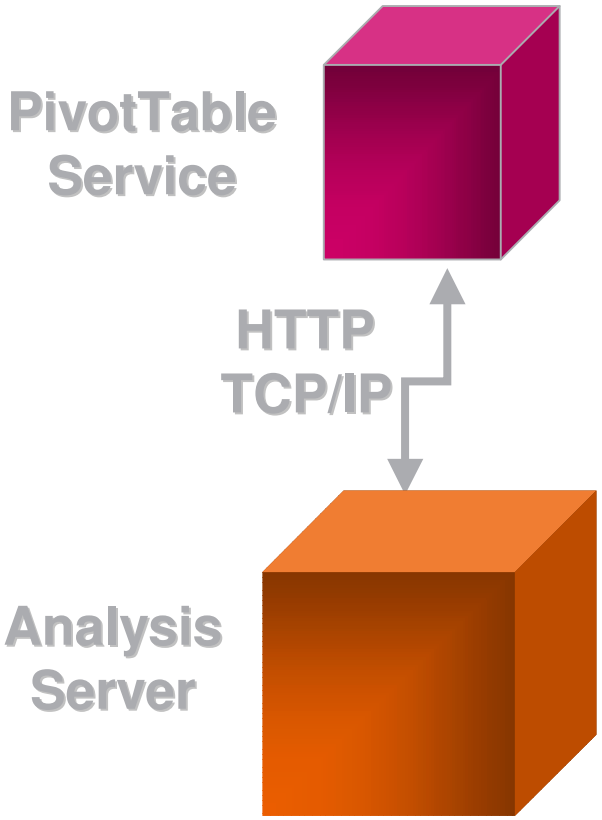

Client Features

- 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
- . . .

Analysis Services Design
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

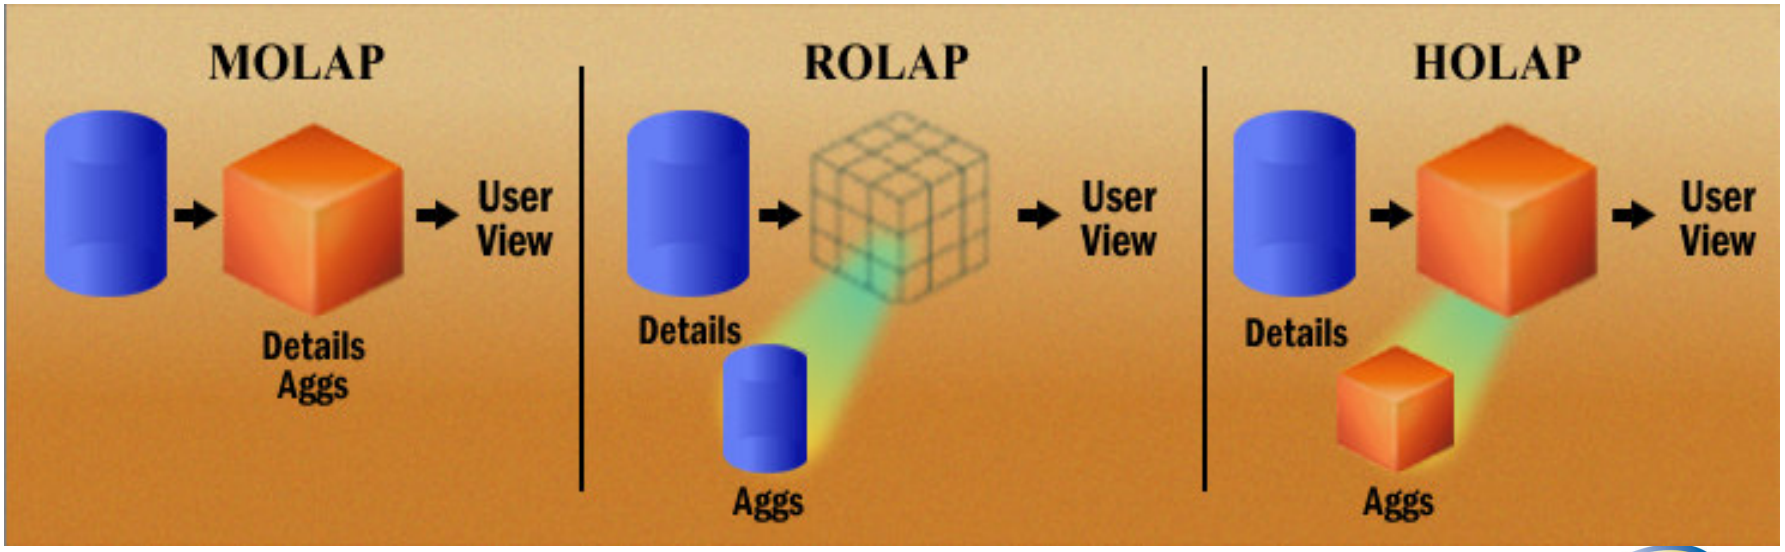

Analysis Services Design Partitioning

- One logical cube

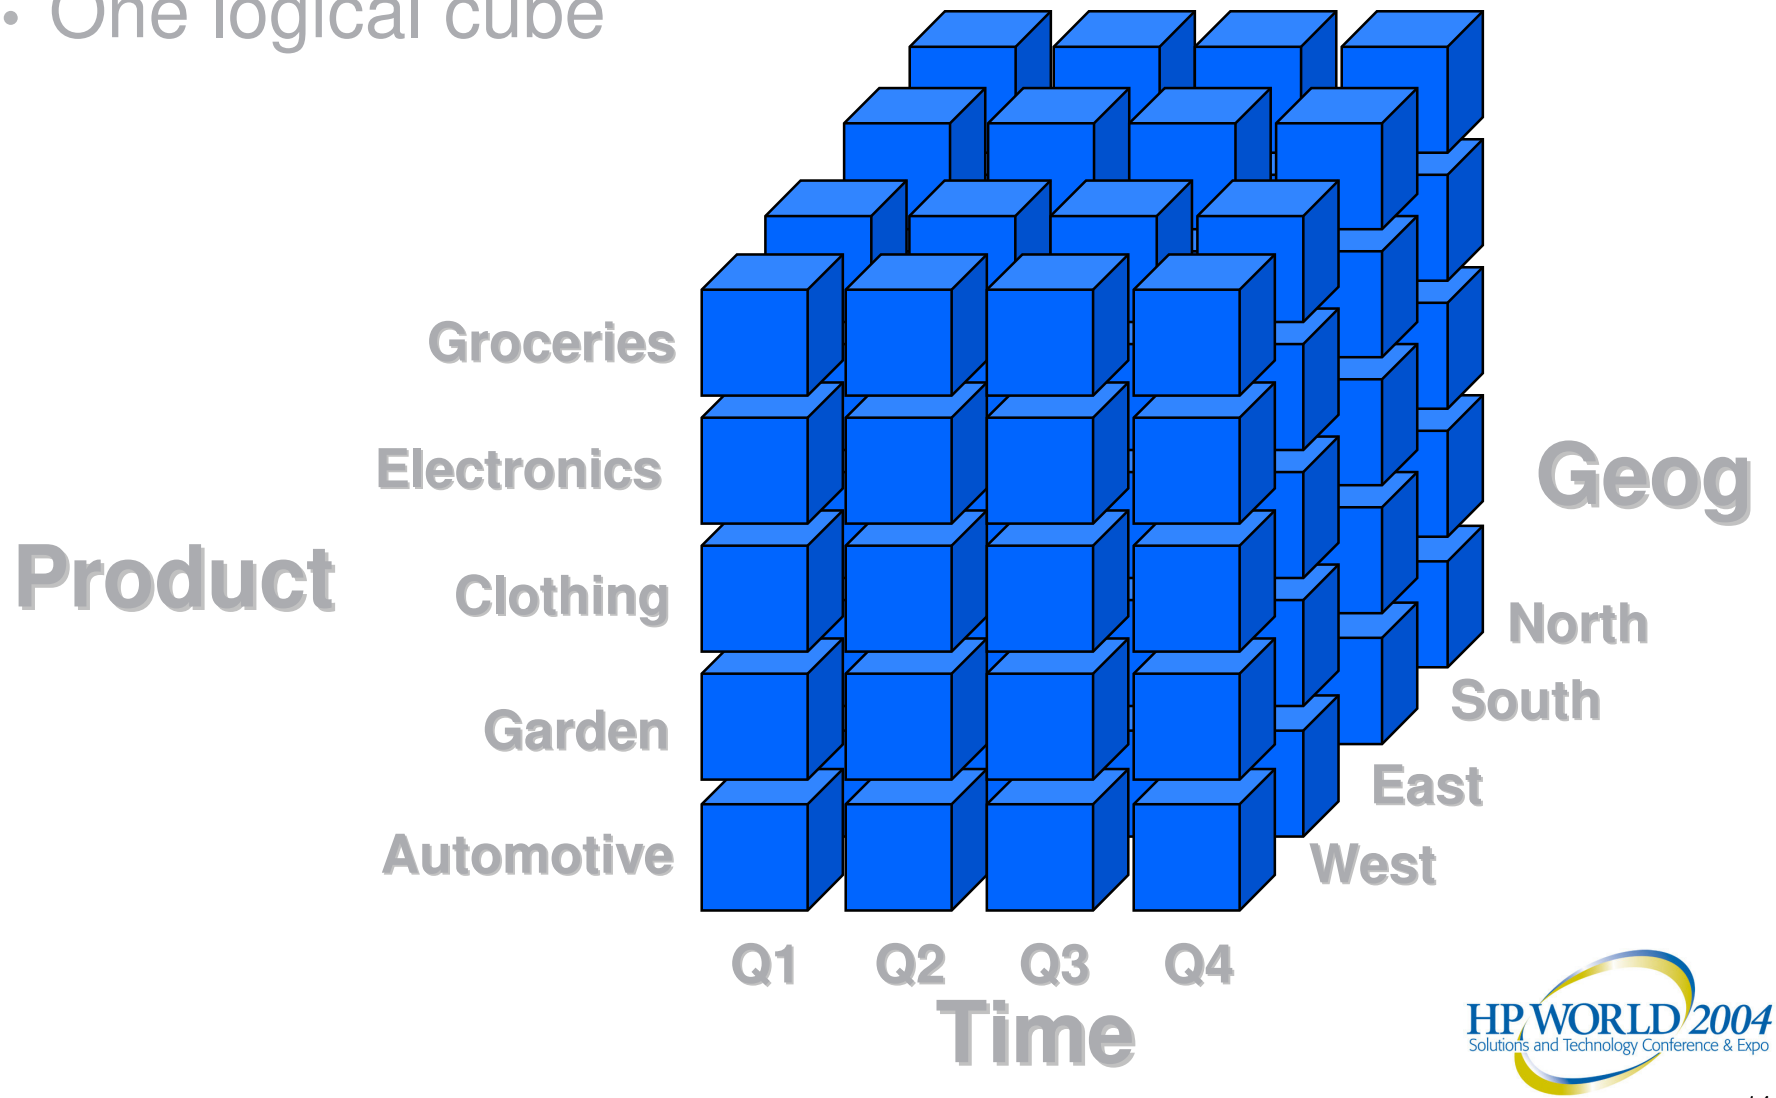

Analysis Services Design Partitioning

- One logical cube

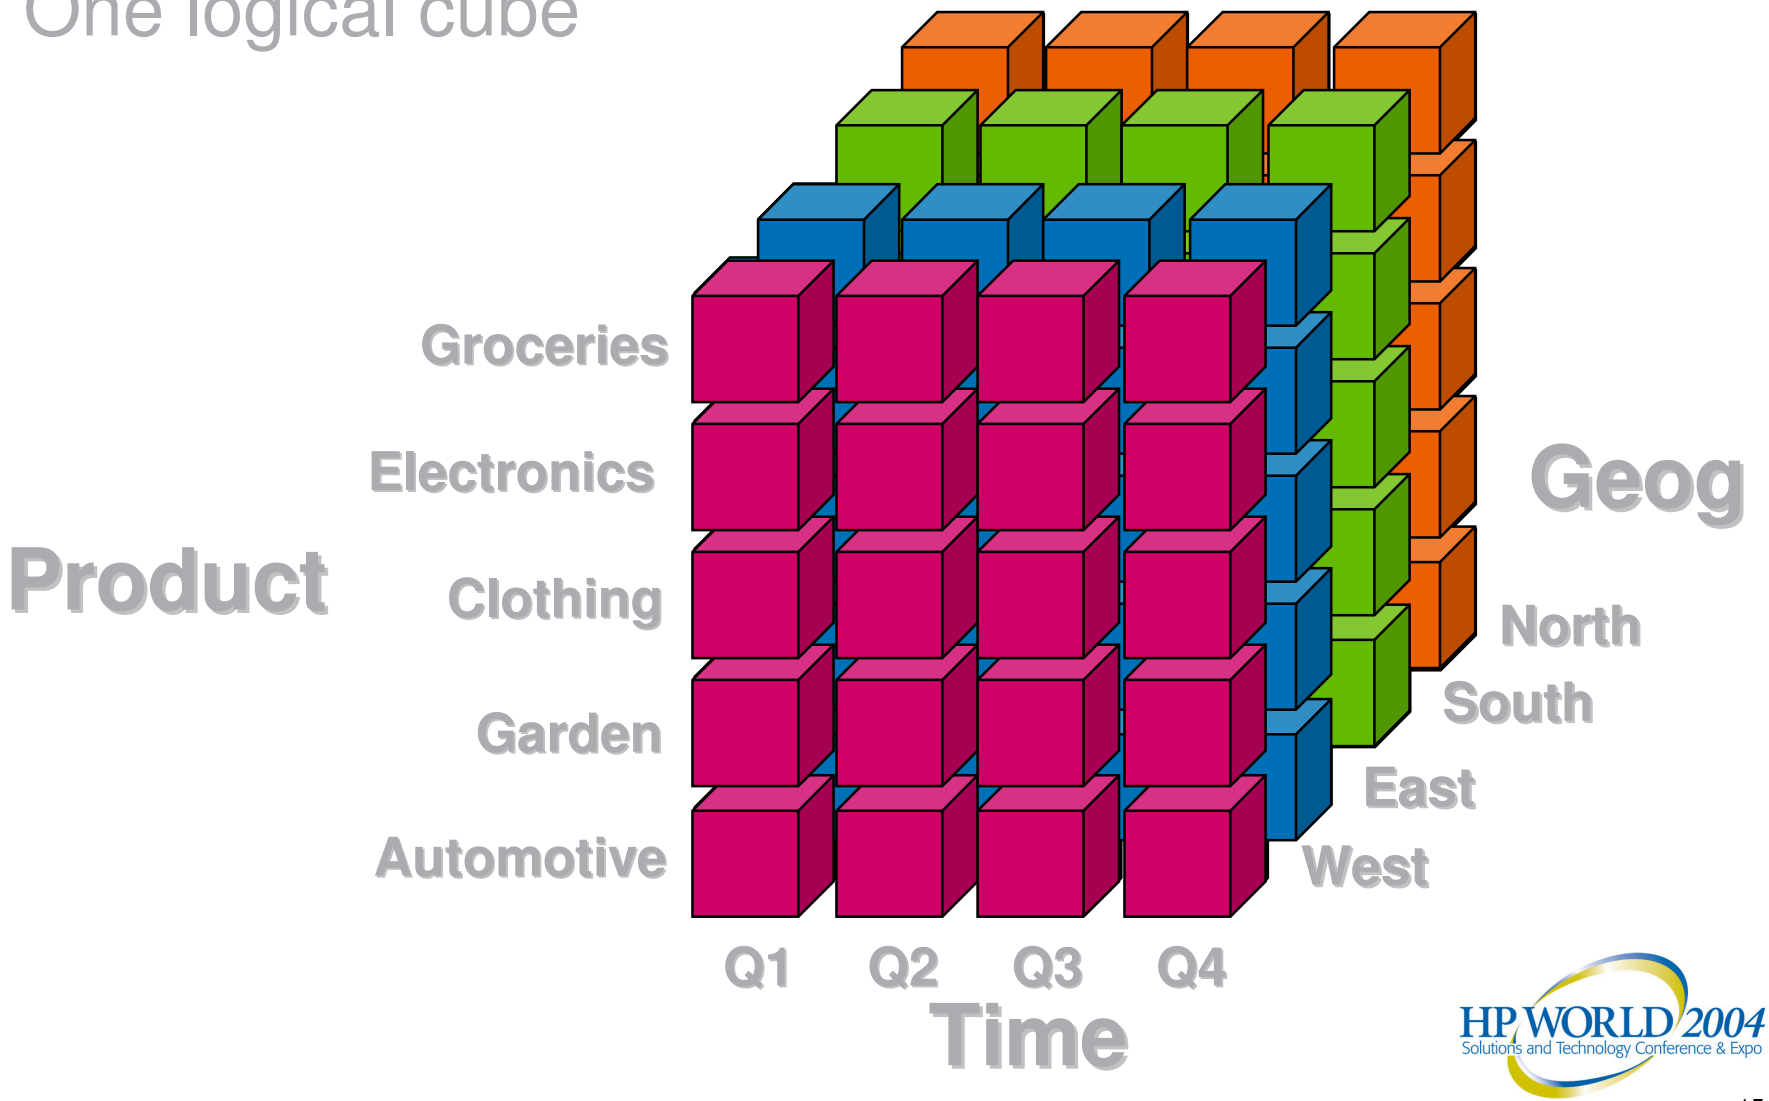

Analysis Services Design
Partition Summary

- 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
  - Write-back
- Important scalability implications
  - Multi-server parallel processing
  - More efficient queries
    - Either single server or multiple server solutions

Query Optimizer

- Star schema aware
- Snowflake schema aware
- Improved cardinality estimation
  - Exploit new statistics
    - captures spikes, skews, & frequent values
  - Improve many computations
- More granular statistics
  - Dynamic bucket boundaries

Partitioning Files & FileGroups

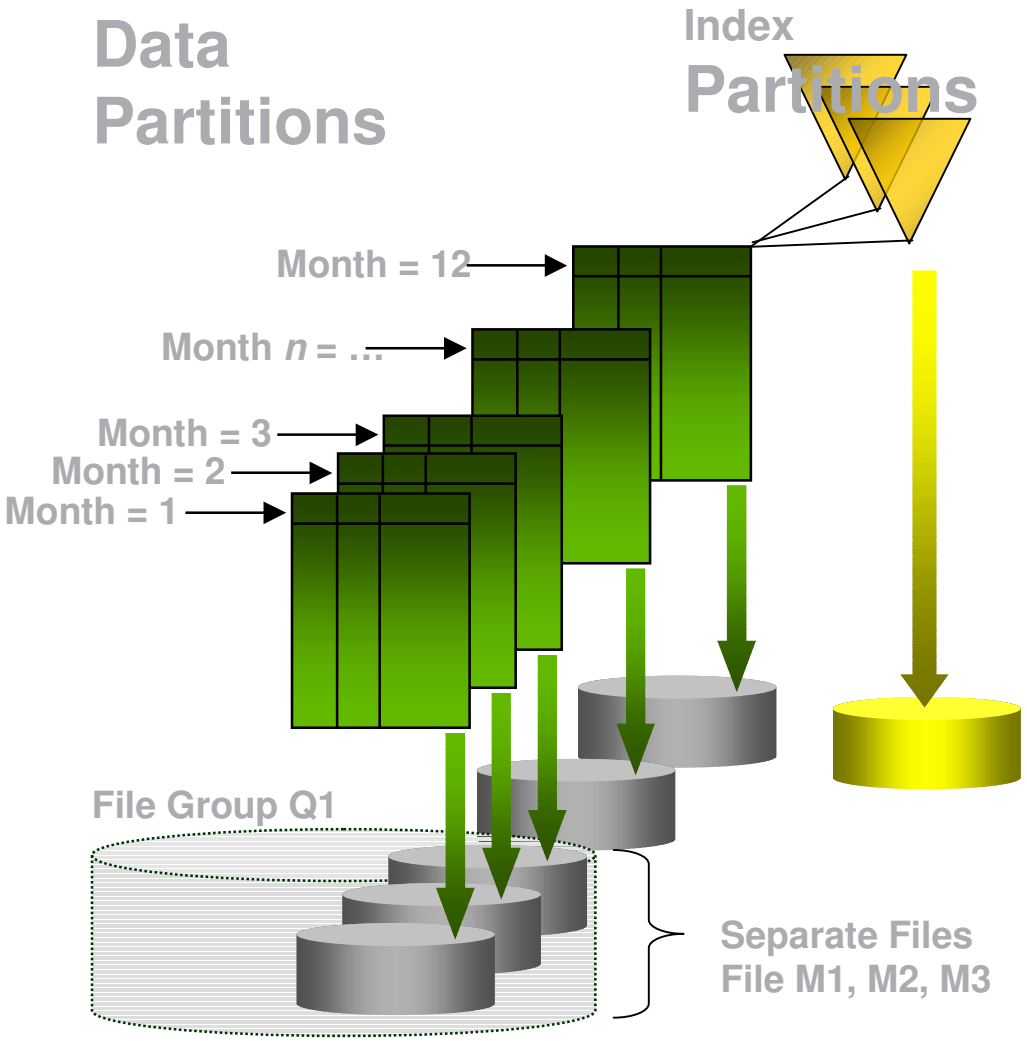

- Placement of data on individual files or file groups
- Granularity improves loading / purging / backup & restore
- Separation improves I/O throughput (parallelism) and reduces contention
- Clean implementation prevents databases from sharing same files

Data Mining

What is DM?

- *Deducing knowledge by examining data*
  - "Knowledge": Patterns, Clusters, Rules, Decision trees, Neural networks, Equations, Association rules…
  - "Examining data" – scanning samples of known facts about "cases."
- Once the "knowledge" is extracted it:
  - Can be browsed
    - Provides very useful insight on the cases behavior
  - Can be used to predict values of other cases
    - Can serve as a key element in closed loop analysis

SQL Server 2000 Analysis Services
Analysis Services – Data Mining Models

- Data Mining Models - Clustering
  - Analyze data for patterns
  - Make predictions
  - Store output
    - Tabular column, cube dimension, Mining model diagram

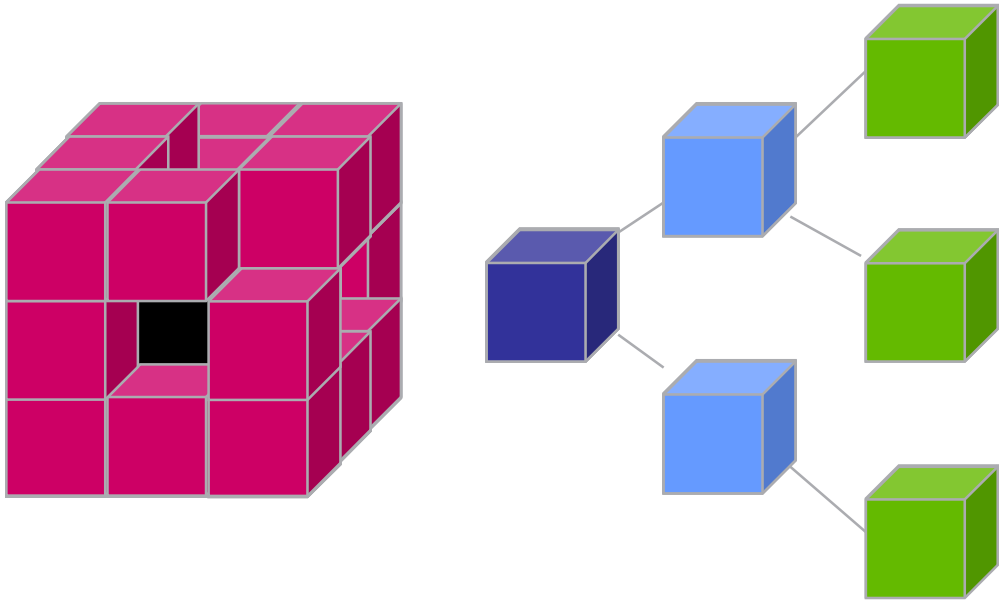

Decision Trees

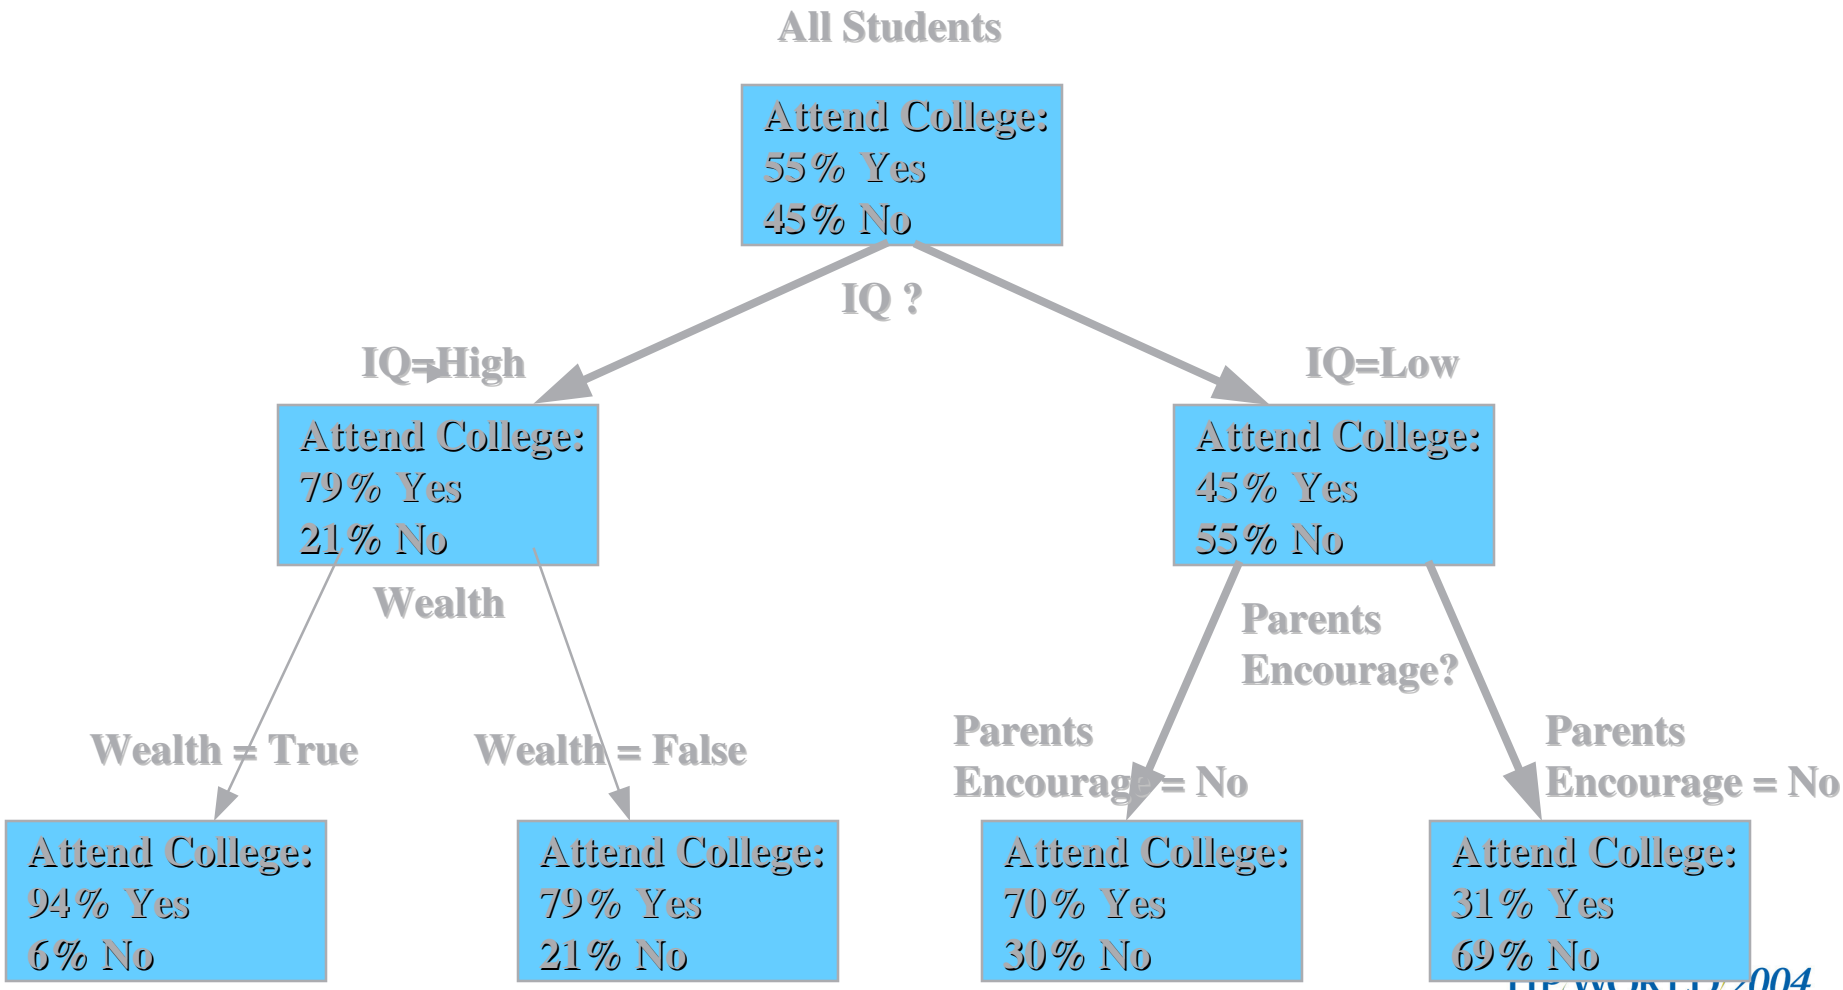

Business Oriented DM Problems

- Targeted ads
  - "What banner should I display to this visitor?"
- Cross sells
  - "What other products is this customer likely to buy?"
- Fraud detection
  - "Is this insurance claim a fraud?"
- Pricing
  - "How deep a discount should I offer to the customer?"
- Risk Management
  - "Should I approve the loan to this customer?"

Applications for Data Mining

- Finance
  - Risk analysis, profitability prediction
- Health
  - Treatment efficacy, fraud detection
- Retail
  - Cross sell, basket analysis
- eCommerce
  - Suggestion list
- Tax
  - Anomaly detection, review selection
- Sport
  - What combination increases goal percentages?

Data Flow

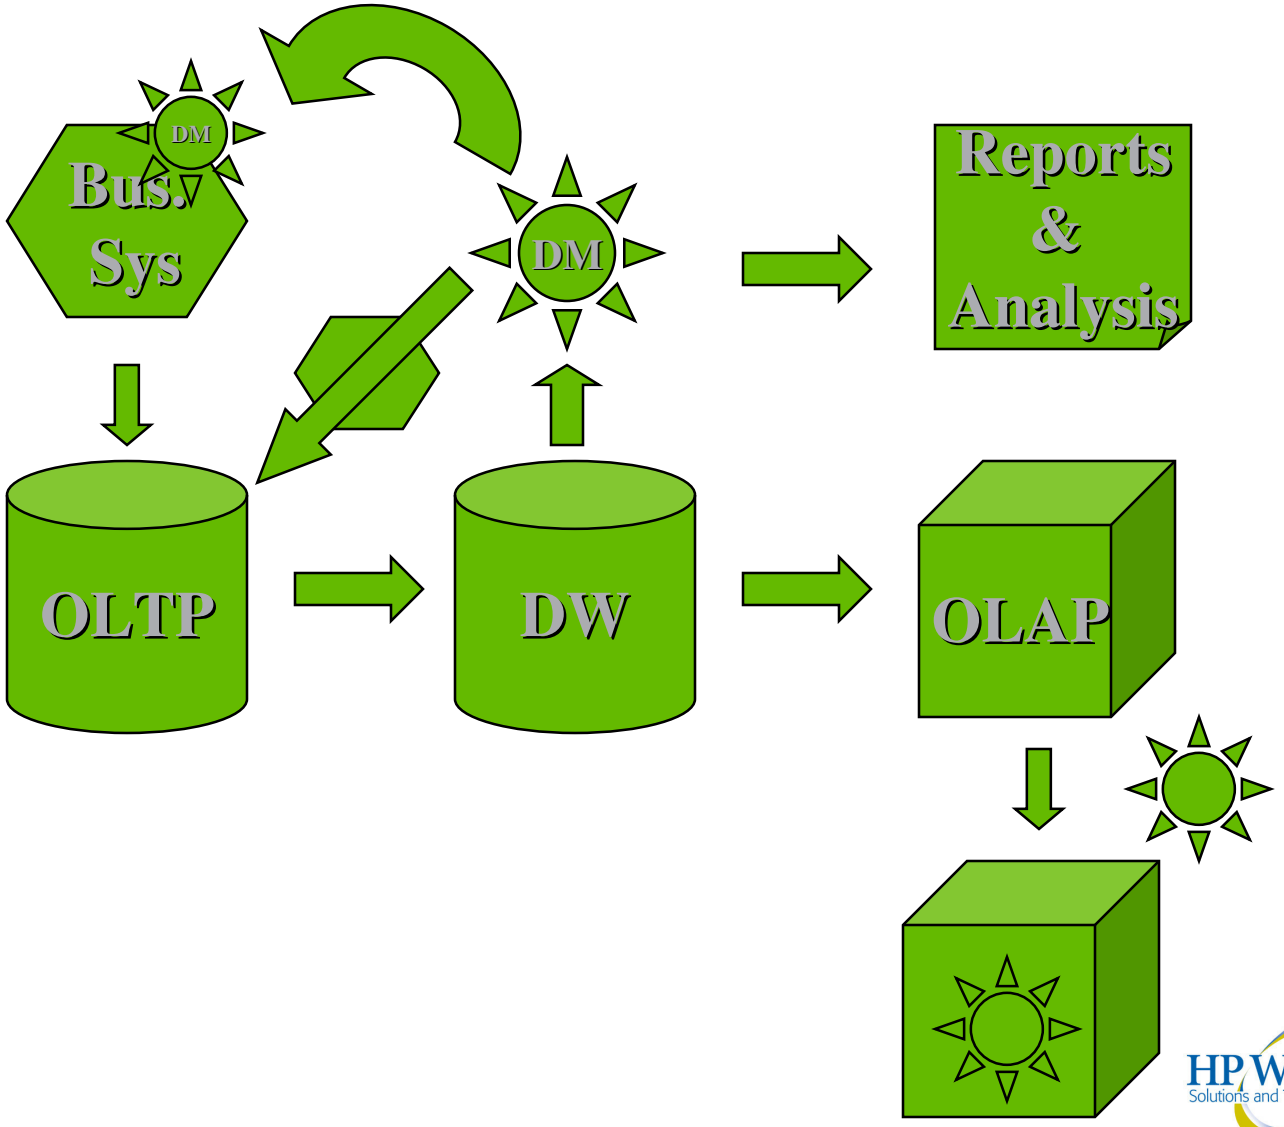

64 Bit
Microsoft SQL Server

SQL Server 2000 64-bit
Advantages

- **Overall 64-bit**
- **Significantly larger direct-addressable memory space than 32-bit.**
- **Improved chip architecture designed to handle parallel processing tasks more effectively.**
- **Better on-chip cache management.**
- **Improved multiprocessor performance and enhanced on-processor parallelism.**
- **Increased bus and I/O bandwidth for faster and wider throughput.**

SQL Server 2000 64-bit
Advantages- Relational Engine

- **Eliminates limitations associated with Address Windowing Extensions (AWE).**
- **Provides all database operations with larger memory address space.**
- **Provides improved parallel processing capabilities and multiprocessor support.**
- **Provides increased bus and I/O bandwidth for faster and wider throughput.**

SQL Server 2000 64-bit
Advantages- Analysis Services

- **Enables loading of very large dimensions into memory.**
- **Allows for a significantly larger query cache.**
- **Allows for a significantly larger aggregate cache during partition processing.**
- **Eliminates issues associated with placing dimension-level security on large dimensions.**
- **Provides enhanced capabilities for processing very large dimensions/large partitions.**

Key Product Information

- Leverages large memory support capabilities of 64-bit version of Windows Server 2003 and hardware enhancements of IPF
  - Theoretical limit up to 18 Exabytes of addressable memory
    - 64 GB of RAM on Enterprise, 512 GB of RAM on Datacenter
    - Virtual memory currently supported up to 8TB
    - Currently testing at 512 GB
- SQL Server 2000 (64-bit)
  - Enterprise and Developer editions
  - SP3 code base
  - Includes all security fixes
- SQL Agent
- Analysis Services (64-bit)
- Data Mining
- Books On Line
- Server Network Utility
- Service Manager

Key Product Information
What's ***NOT*** Included?

- Data Transformation Services
  - Copy database wizard not available
  - However, you can store packages on 64-bit servers
- Most tools
  - Enterprise Manager/Analysis Manager
  - Query Analyzer
  - Profiler
- English Query & Meta Data Services
- SQL Mail
  - SQL Agent Mail can be configured on 64-bit.
- Use a 32-bit workstation to administer a 64-bit server.
- Do not try emulation mode…it is ***not*** supported.
- MSDE

Object Compatibility
SQL Server and Analysis Services

- On-disk format is the same for both 32-bit and 64-bit.
- SQL Server
  - Detach a 32-bit database and attach it to 64-bit server. It ***IS*** that easy.
  - Backup and restore supported across 32-bit and 64-bit.
- Analysis Services
  - Archive the 32-bit database and restore it on 64-bit.
  - For large AS databases, copy and paste and process on 64-bit.
- All T-SQL and MDX "just runs"

Advantages Of 64-Bit To RDBMS

- Larger direct addressable memory
  - Eliminates the need for the AWE layer
  - ALL database operations benefit
    - Not just data cache operations
  - Increased support of concurrent users (cursors, transactions, locks, etc.)
  - Larger query cache, procedure cache, sort heap, buffer pool, open DB objects, intermediate results *(complex query plans)*, etc.
- Processor architecture handles parallel processing better
  - Good for DB parallel processing tasks
  - More work done in a single clock cycle
- Increased bus and I/O bandwidth throughput
  - On-die cache
  - 128-bit bus width

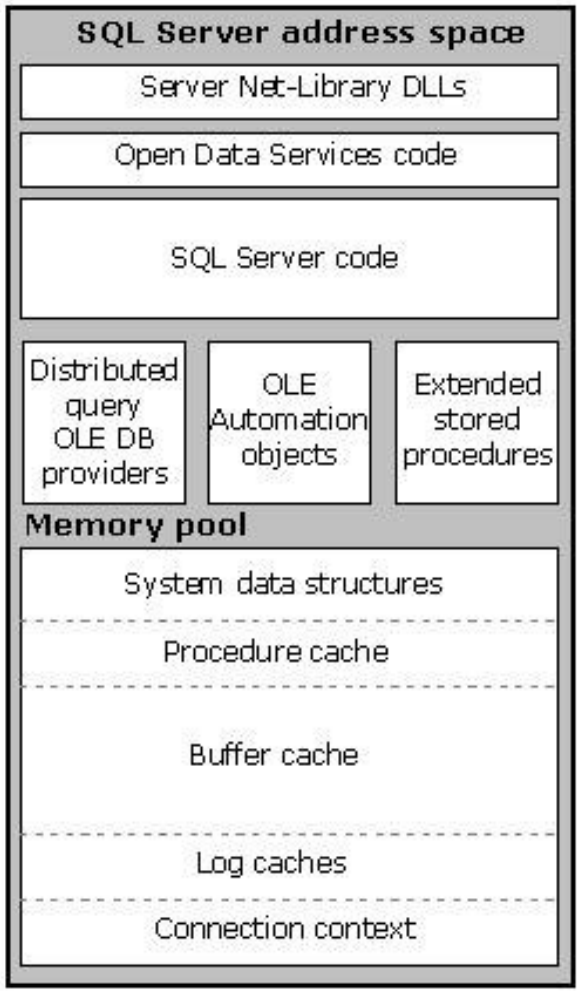

HP Internal Large Data Warehouse Implementation

HP Analytics & Research Architecture

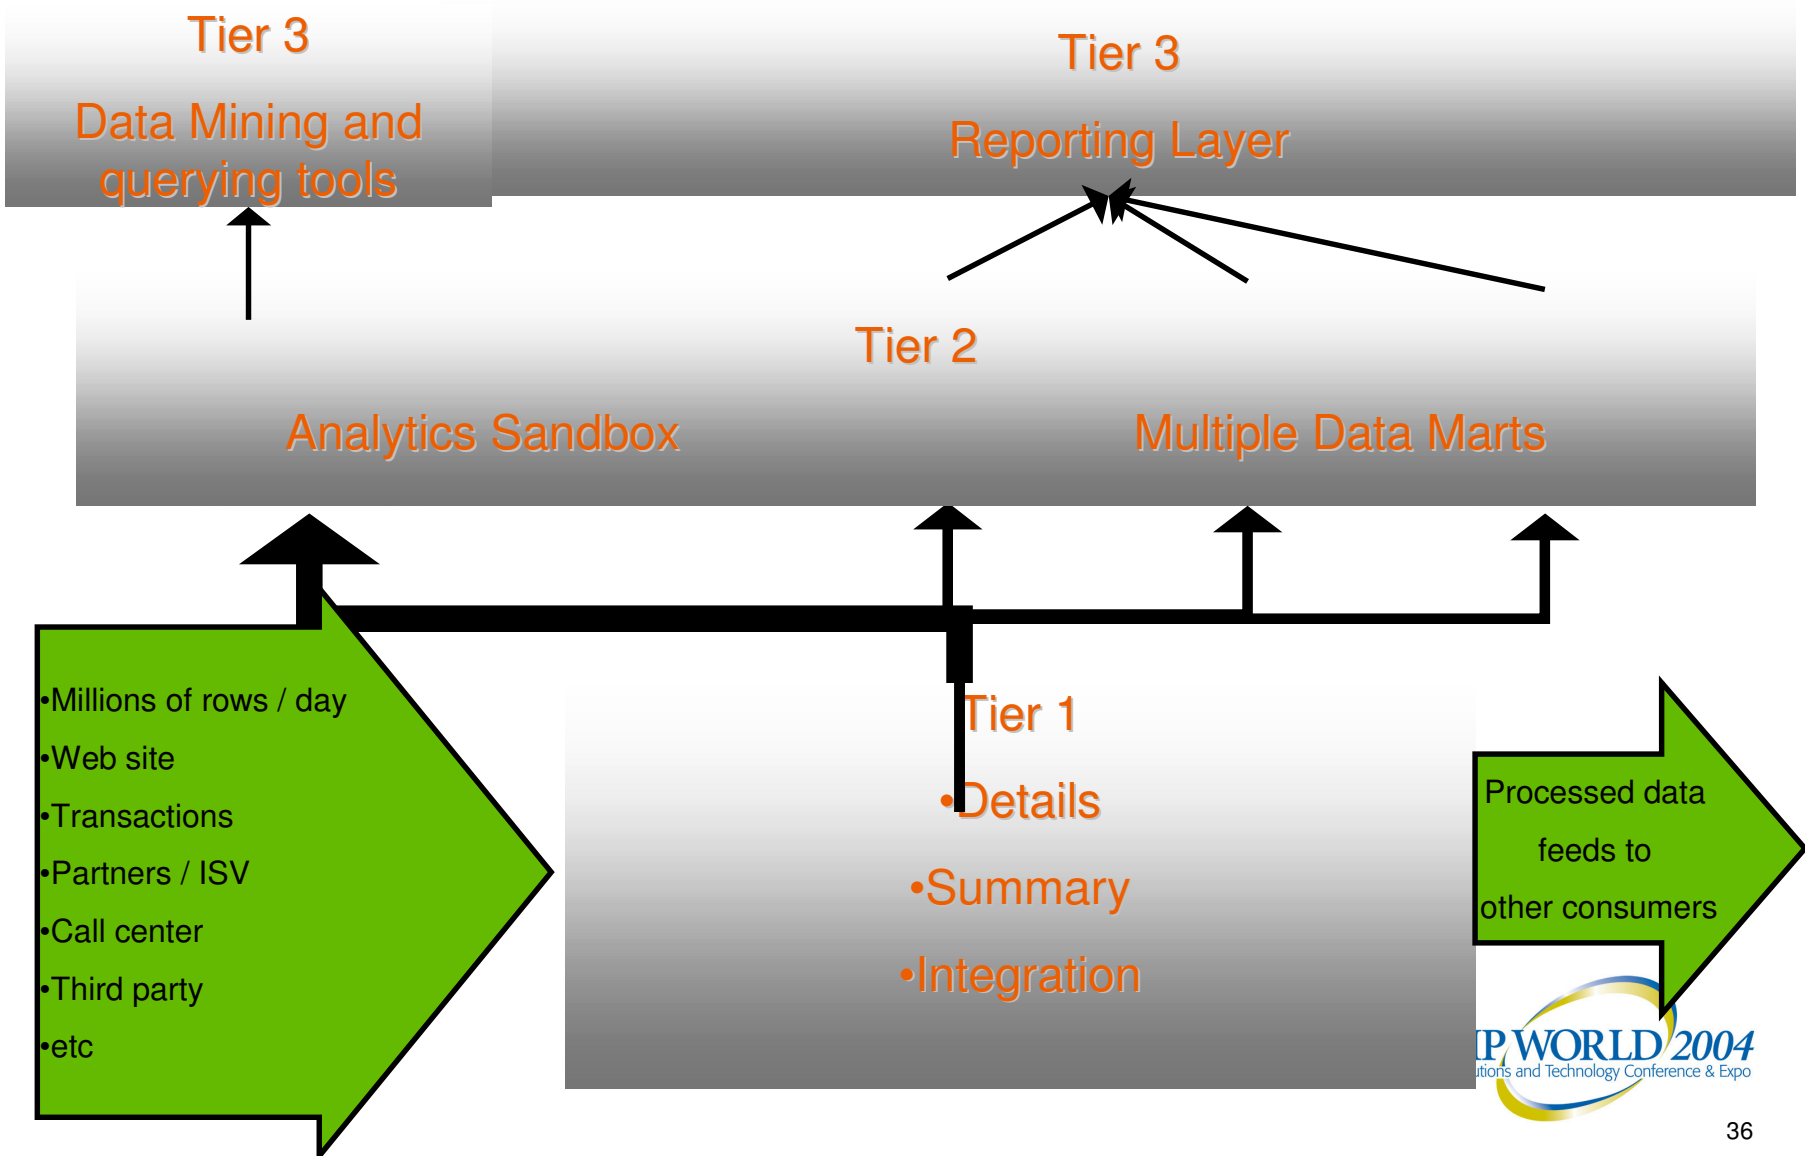

HP Analytics & Research
Hardware & Software Details

- HP Integrity Superdome 32 way, 64 GB
- 4 x EVA 5000
  - Raw = ~30 TB
    RAID5 = ~3 TB
    RAID10 = ~13 TB
- SQL 2000 EE 64 bit
- Windows 2003 Data Center 64 bit

Results from a Test run by Microsoft

Performance Test
Hardware Profile

SQL Server 2000 Relational DB Test Configuration

HP RX5670
Quad 900 MHz Intel® Itanium® 2 processors
16 GB Ram
C: Drive 292 GB Internal HD (4 x 73GB); RAID 1
Windows Server 2003 Enterprise Edition
SQL Server 2000 Enterprise Edition 64-bit

HP VA 7400 - fibre drive chassis storage array
396 GB (11 x 36 GB 10k RPM)
216 GB (12 x 18 GB 15k RPM)
584 GB (8 x 73 GB 10k RPM)
==========
1196 GB Total

Actual storage was divided into two separate logical drives and configured as AutoRaid (psuedo RAID 5)

Data Profile

- Microsoft SQL Accelerator for BI (SSABI) schema
- Tests Conducted on Sales_Fact tables in the DW:

Table name	1 billion row-database	Tbl Size (GB)
Fact_Backlog	382,756,178	~44 GB
Fact_Orders	324,823,696	~38 GB
Fact_Sales	293,907,214	~50 GB
Total	1,001,487,088	~132 GB

http://www.microsoft.com/technet/treeview/default.asp?url=/technet/prodtechnol/sql/maintain/Optimize/SSABIPER.asp

64-bit RDBMS Testing
Complex Query (Cold Cache)

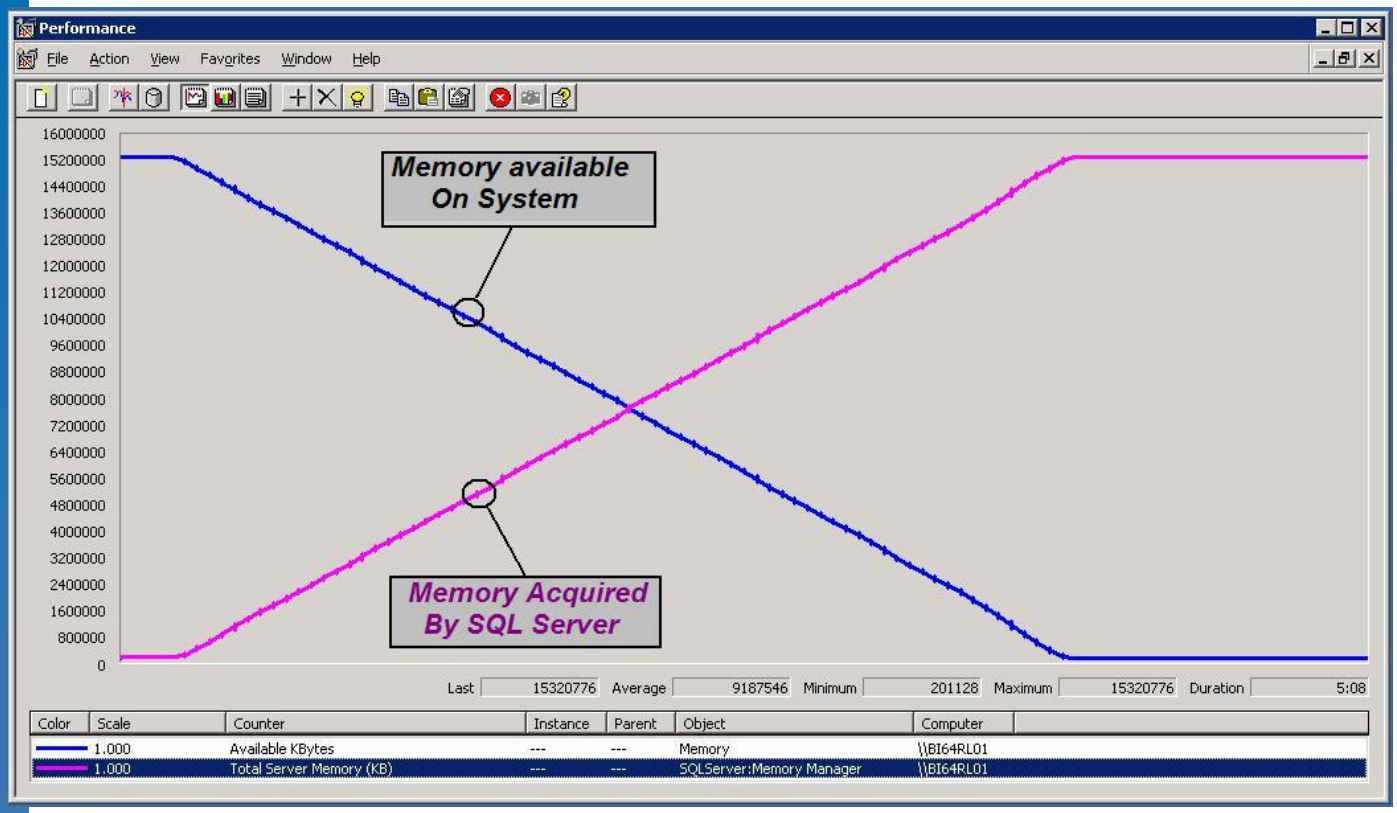

*Graph illustrates use of large memory to resolve a **large multi-step query** as shown in the plan at right*

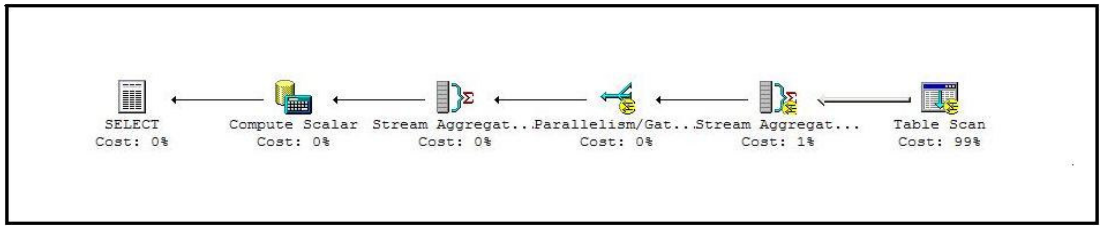

64-bit RDBMS Testing
Cold & Warm Cache

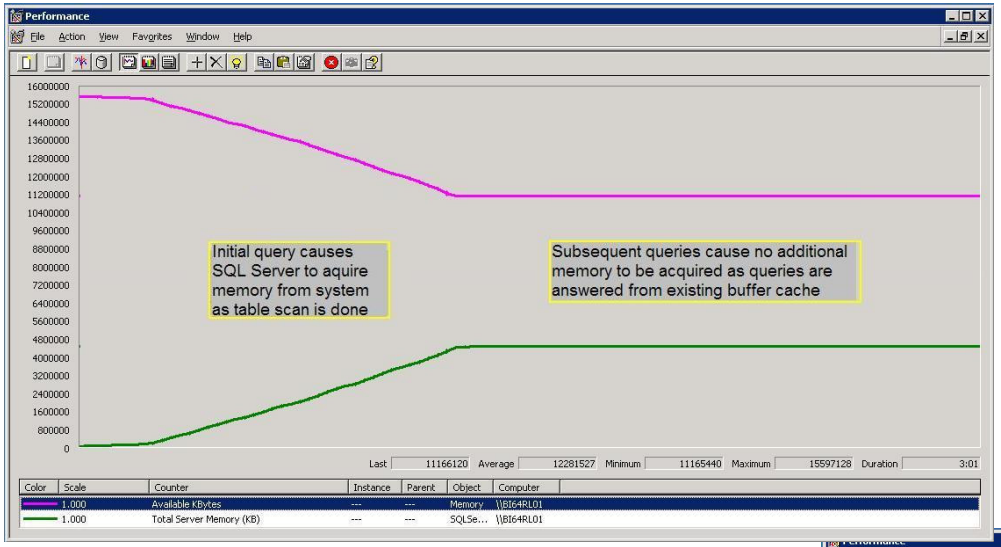

Graph on left illustrates use of large memory to satisfy similar queries from buffer cache.

Graph below illustrates CPU Activity

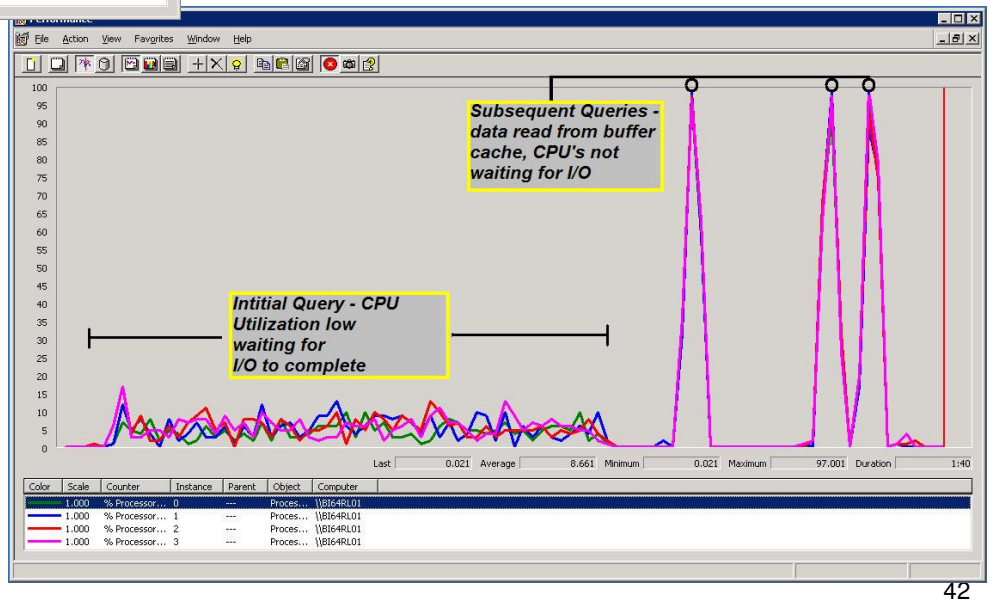

NOTE - initial query slower & CPU utilization is low while waiting on Disk I/O to read table data into RAM.

Subsequent queries run quickly once table is in cache!

Analysis Services (64-bit)

Advantages Of Analysis Services (64-Bit)

- Not constrained by 3GB memory limit
  - Dimension memory + Shadow dimensions + Buffer space (partitions)
  - Cache fills remaining space
- Huge dimensions supported in memory
  - MOLAP dimensions are extremely large, but provide best query performance
  - Replicas for dimension security
- Large Process Buffers for processing
  - Temp files can be avoided
- Large Data Cache
  - Less evictions = more warm cache queries.
- Parallel processing of partitions
- Simultaneous processing and querying

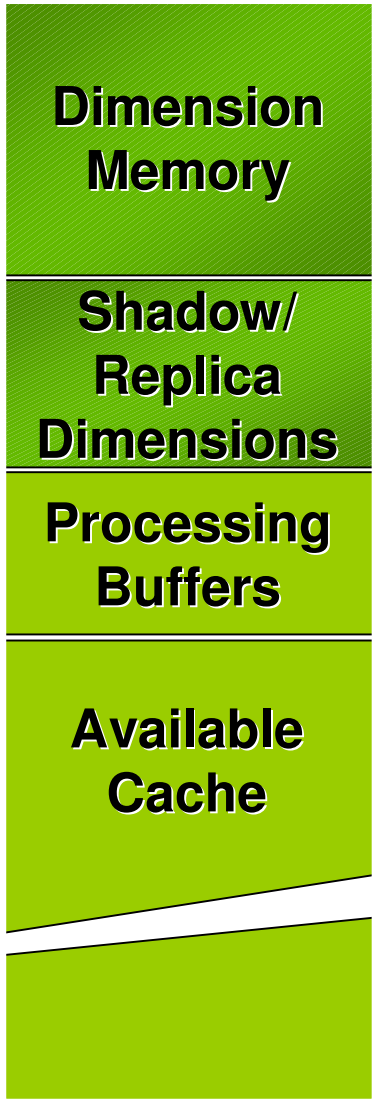

When to consider 64-bit Analysis Services?

- Total count of dimension members is in millions.
  - Don't forget about shadow dimensions.
- Extensive use of dimension security.
- Processing of partitions requires temp file usage.
  - Parallel processing, distinct count, large and/or many aggregations defined.
  - Monitor with perfmon or filemon.exe
- Significant rate of aggregation cache evictions and users could benefit from cache.
- Availability of appropriate 64-bit OLEDB providers.

Analysis Services
Migration Tips/Issues

- 32-bit server is required for administration.
  - Gigabit, same subnet, same segment.
- SQL Server instance is required.
  - Query Log and Repository stored in SQL Server.
- UDFs need to be compiled to 64-bit.
  - Windows SDK contains compiler.
- Excel functions are not available.
- VBA Functions are available.
- Copy and Paste database and reprocess on 64-bit server.
- MemoryASMB Registry Entry.
- VLDM is disabled.

Analysis Services
Migration Tips/Issues

- Analysis Services server and 32-bit administration box must be on the same domain.
- Backup and Restore via msmdmarch.exe not recommended for routine maintenance.
  - Use file system backups and backup the OLAP repository.
- You must configure the 32-bit client exactly the same as the 64-bit server.

Analysis Services
Data Access Summary

- Analysis Services requires OLEDB
- Provider availability
  - SQL Server 2000 OLEDB provider included.
  - Oracle and IBM DB 2 are in beta
    - Successfully tested the Oracle provider, but we didn't test all data types in dims and partitions.
  - There is no Microsoft OLEDB for ODBC provider.
    - You cannot use any ODBC drivers.
  - Contact the database vendor for 64-bit providers .

Test Background
Analysis Services

- SSABI Sales and Marketing Schema
- Customer.Direct dimension contains 50 million members
- Storage mode is MOLAP
- Server is HP 4 x 1.0 Ghz
- Server RAM is 32.0 GB
- Process Buffer is set to 1.5 GB
- 5 Years of fact data generated in 3 cubes.
- Total record count is 1.8 Billion facts seasonally adjusted
  - 4.6 million to 23.5 million records per fact table

Test Scenarios
Analysis Services

- Very Large dimensions on 64-bit.
  - How many members?
  - How long does it take to process?
  - What's query performance like against it?
- Simultaneous query and processing.
- What did we learn from this experience?

Dimension Memory Utilization

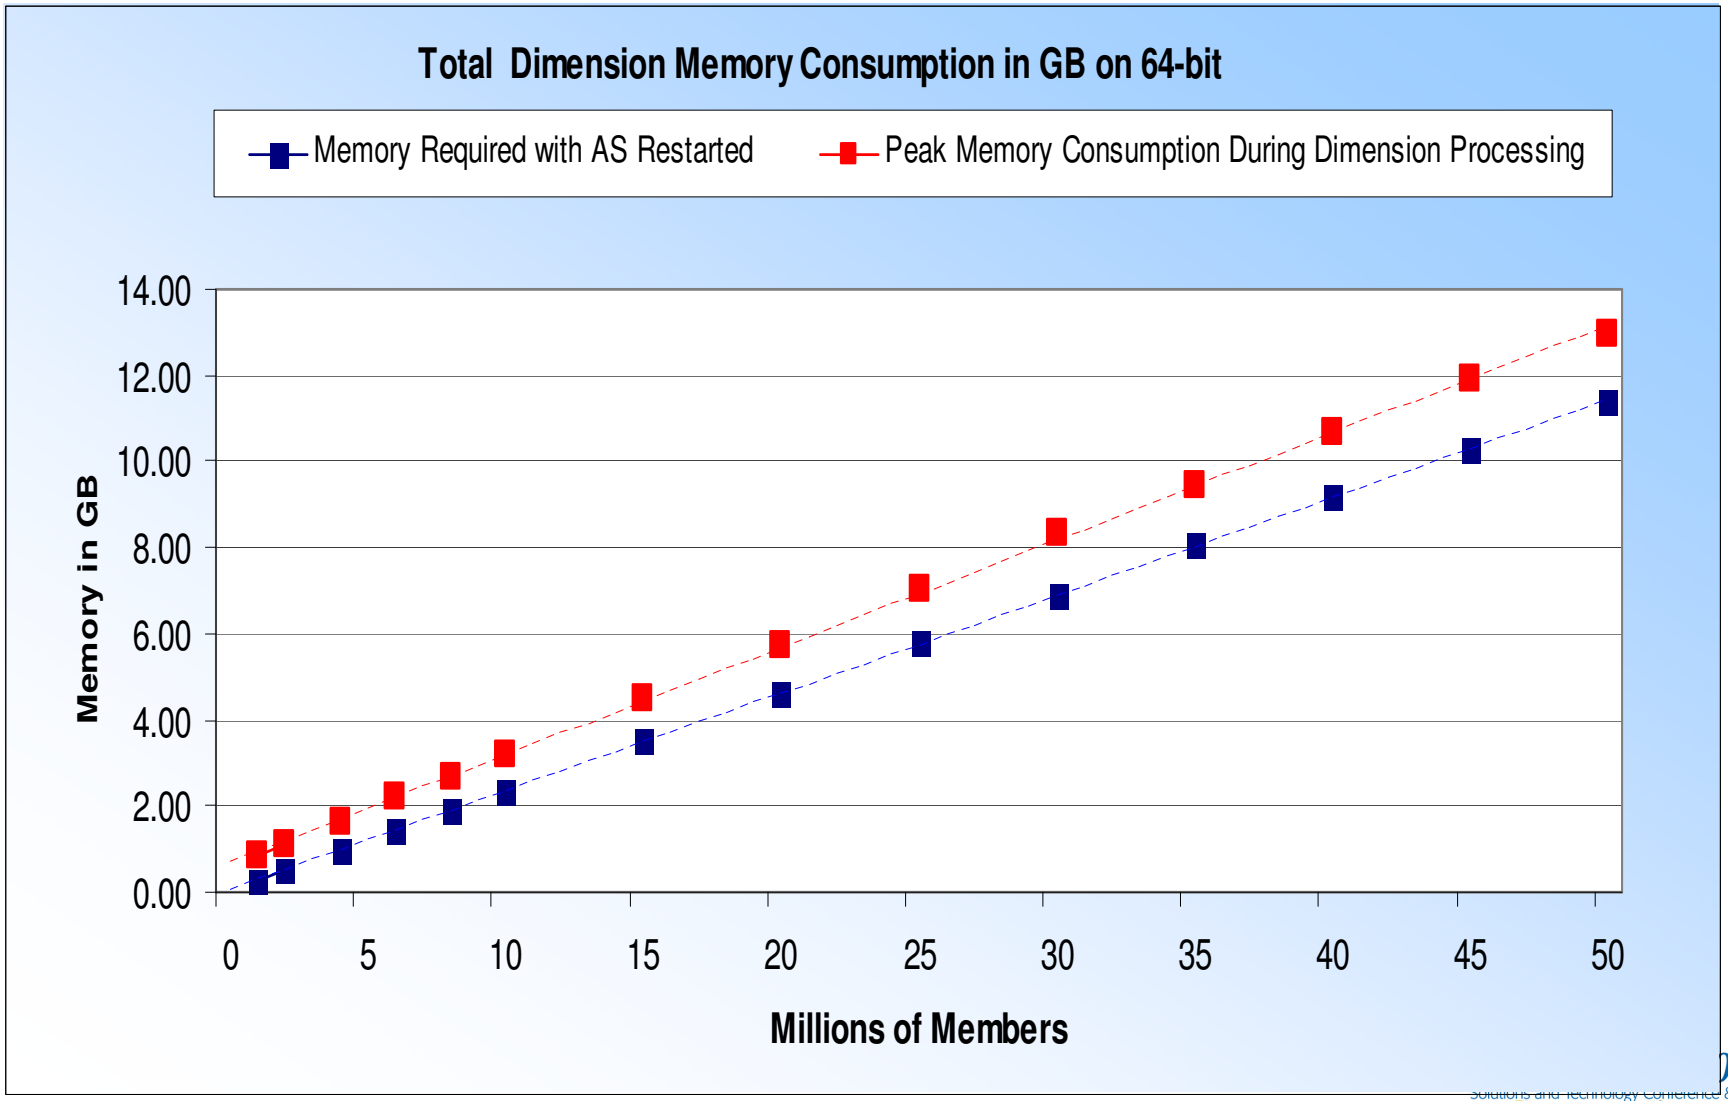

Dimension Processing

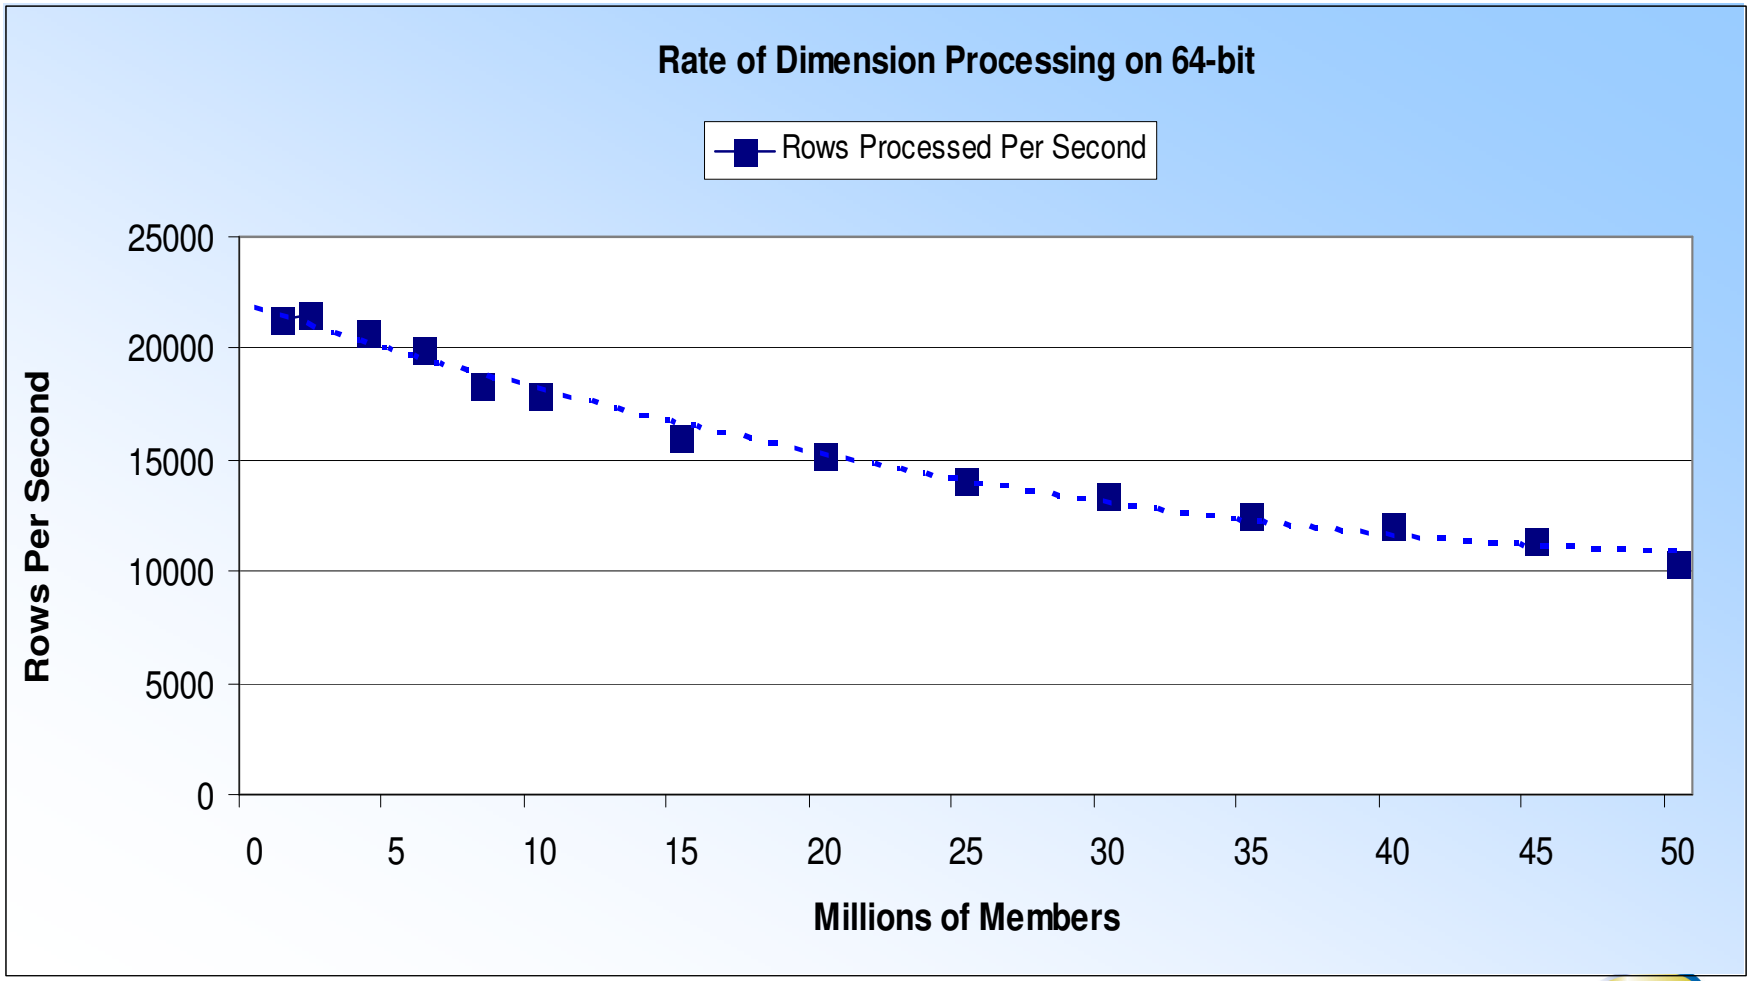

Roughly 80 minutes to process a 50 million member dimension

Analysis Services
Dimension Memory Tips

- How much memory do I need?
- Dimension memory utilization is linear based on the number of members for a given set of dimension characteristics.
- Steps for calculating dimension memory
  - 1. Process a subset of the dimension.
  - 2. Calculate the slope of the line.
  - 3. Compute amount of memory needed based on how many members you expect.
- Caveats
  - Shadow dimensions, dimension security

Query Tests
Defining Cold and Warm Cache

- Cold Cache
  - All queries answered by the server accessing the file system.
- Warm Cache
  - All queries answered by retrieving records from the aggregation cache.
- Customer Reality
  - Typical customer query workload will consist of ***BOTH*** cold and warm cache queries.

Query Tests
Query Workload Characteristics

- Guaranteeing Cold Cache Queries
  - 12,000 unique queries generated.
  - No pyramid operations allowed.
    - No result set can be derived from any another result set
  - Limit members per query to keep server from downloading sibling result sets.
- Problem: It isn't an accurate reflection of a typical customer workload.
  - How do you define a 'realistic' query load?
    - It's going to be different for every customer!
- Use perfmon Analysis Services:Query counters

Query Performance
No Aggregations

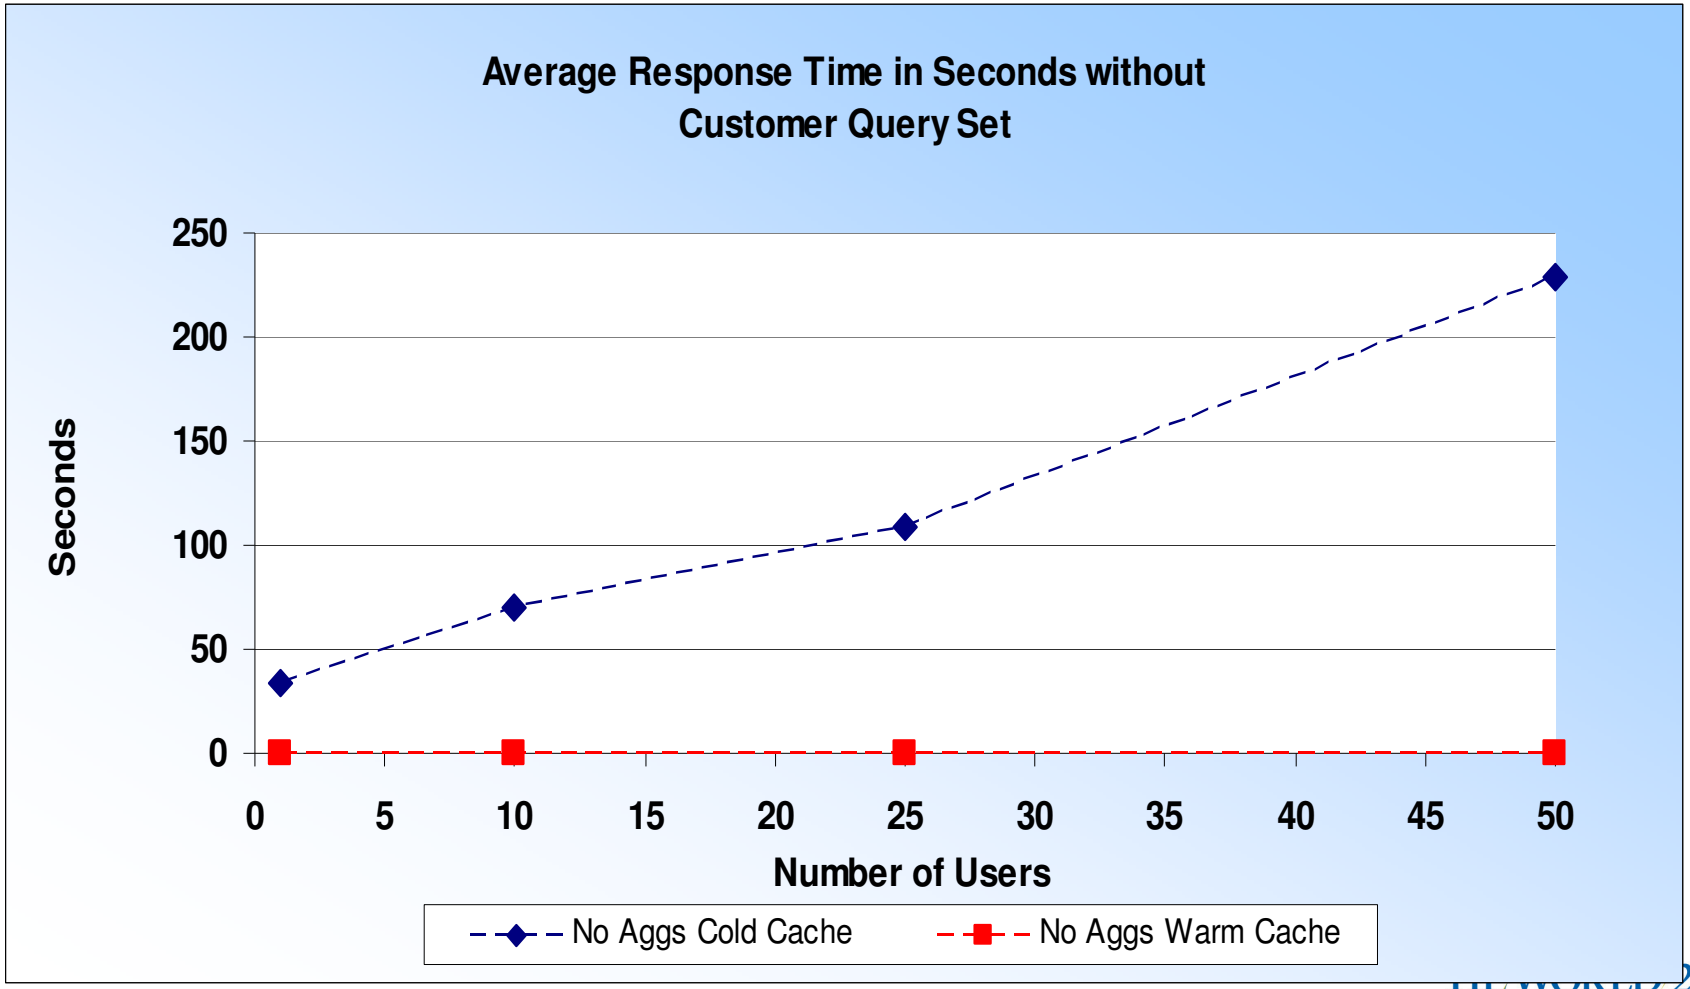

Query Performance
UBO plus Targeted Aggregations

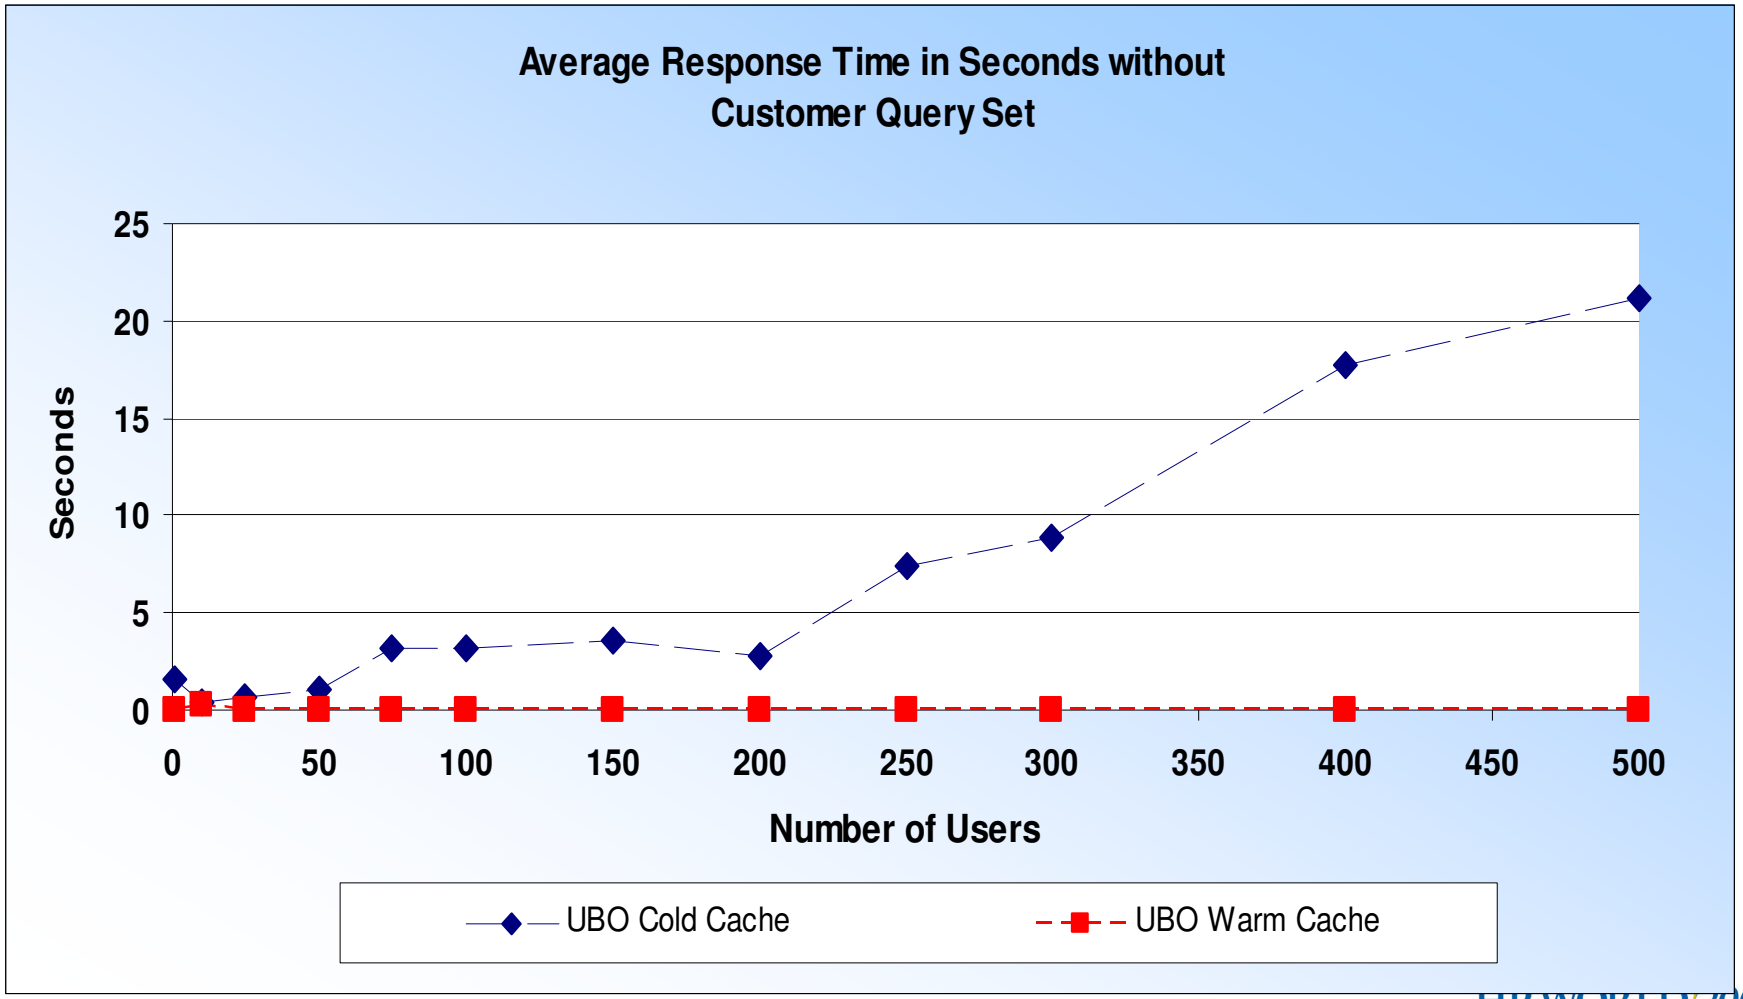

Query Performance
Querying the Customer Level

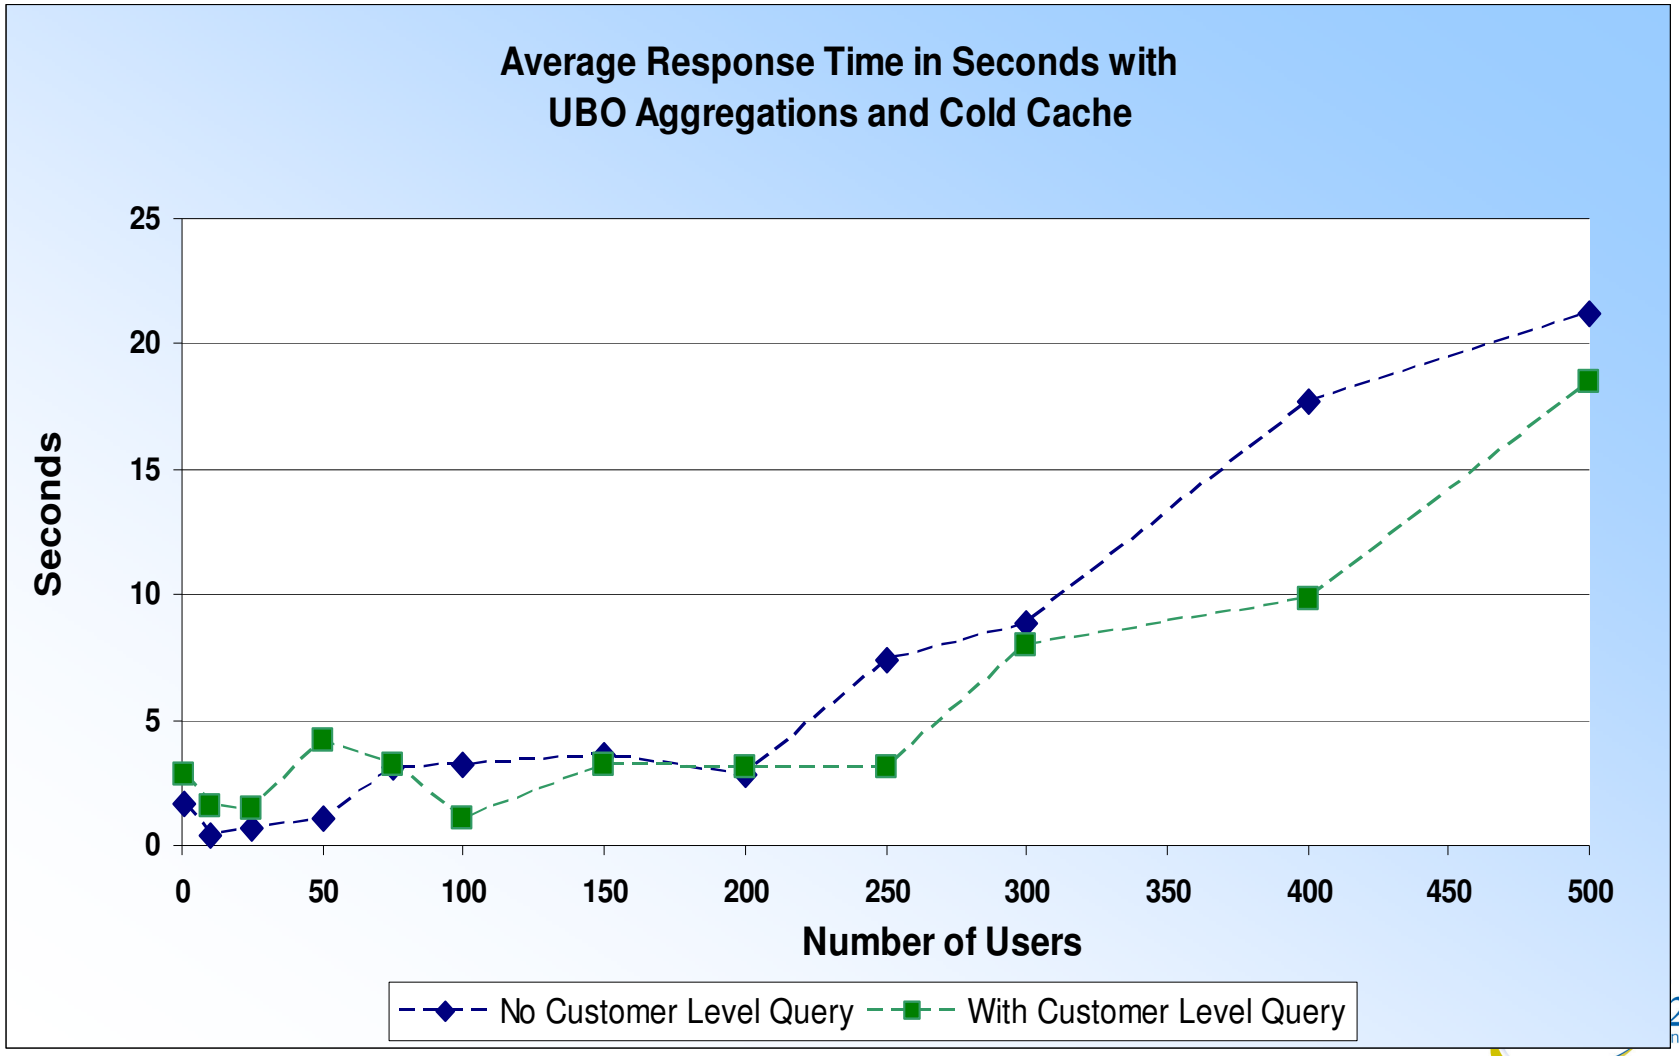

Query Performance
Simultaneous Processing and Querying

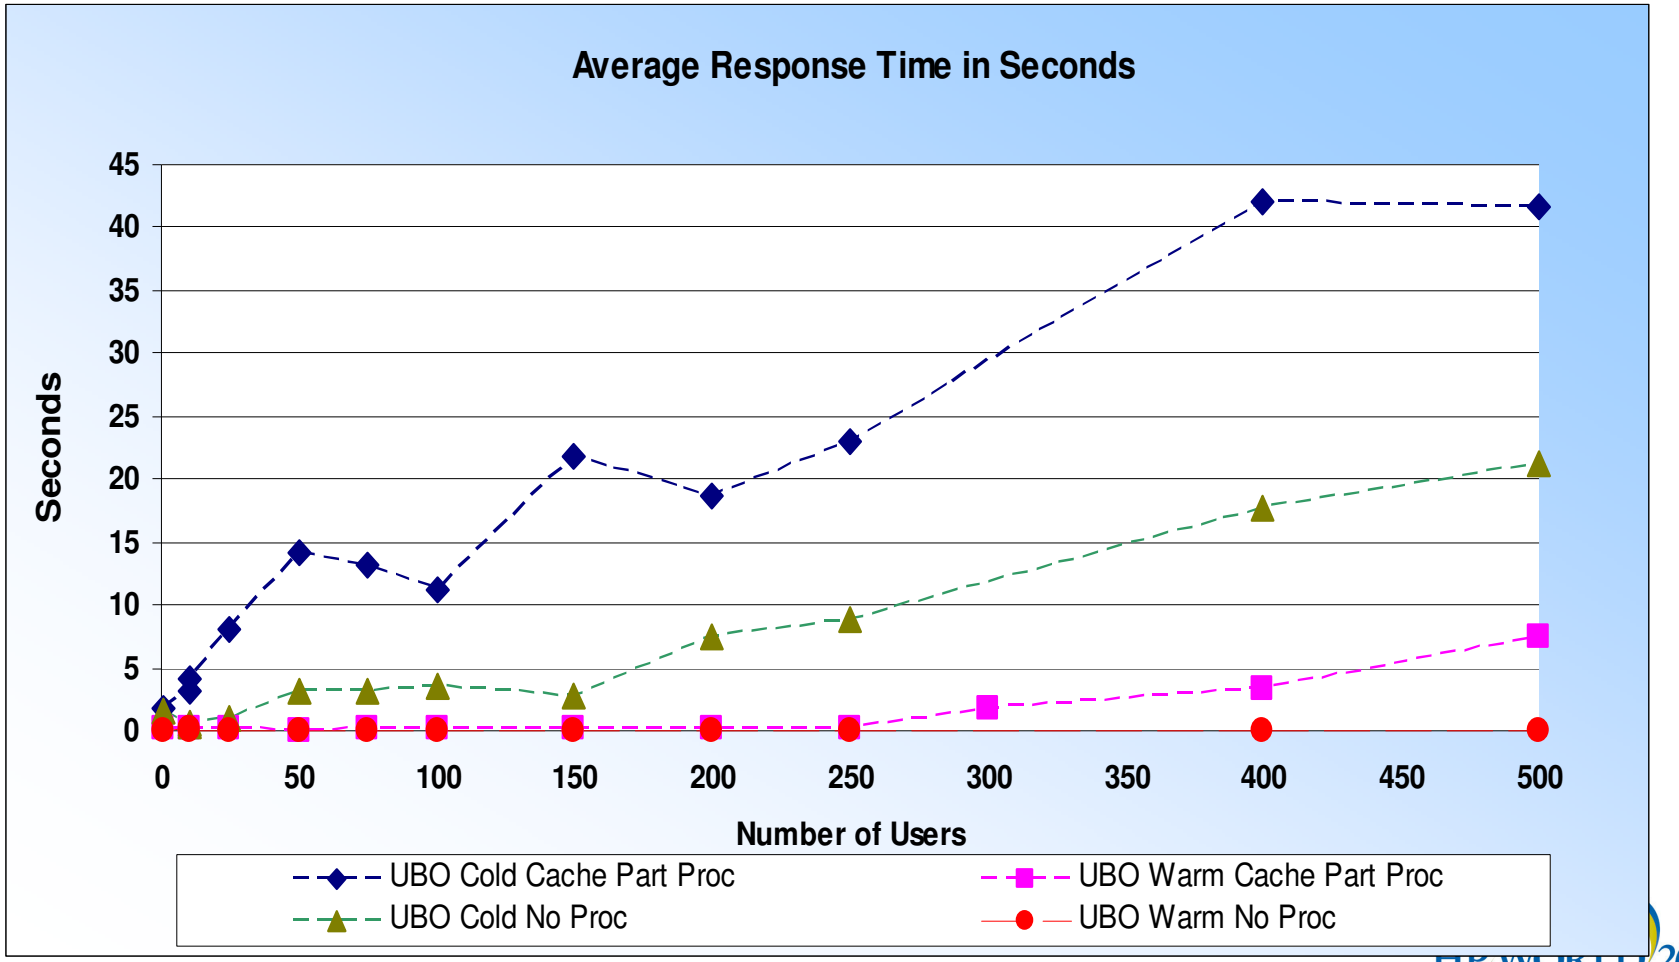

Analysis Services
Simultaneous Processing and Querying

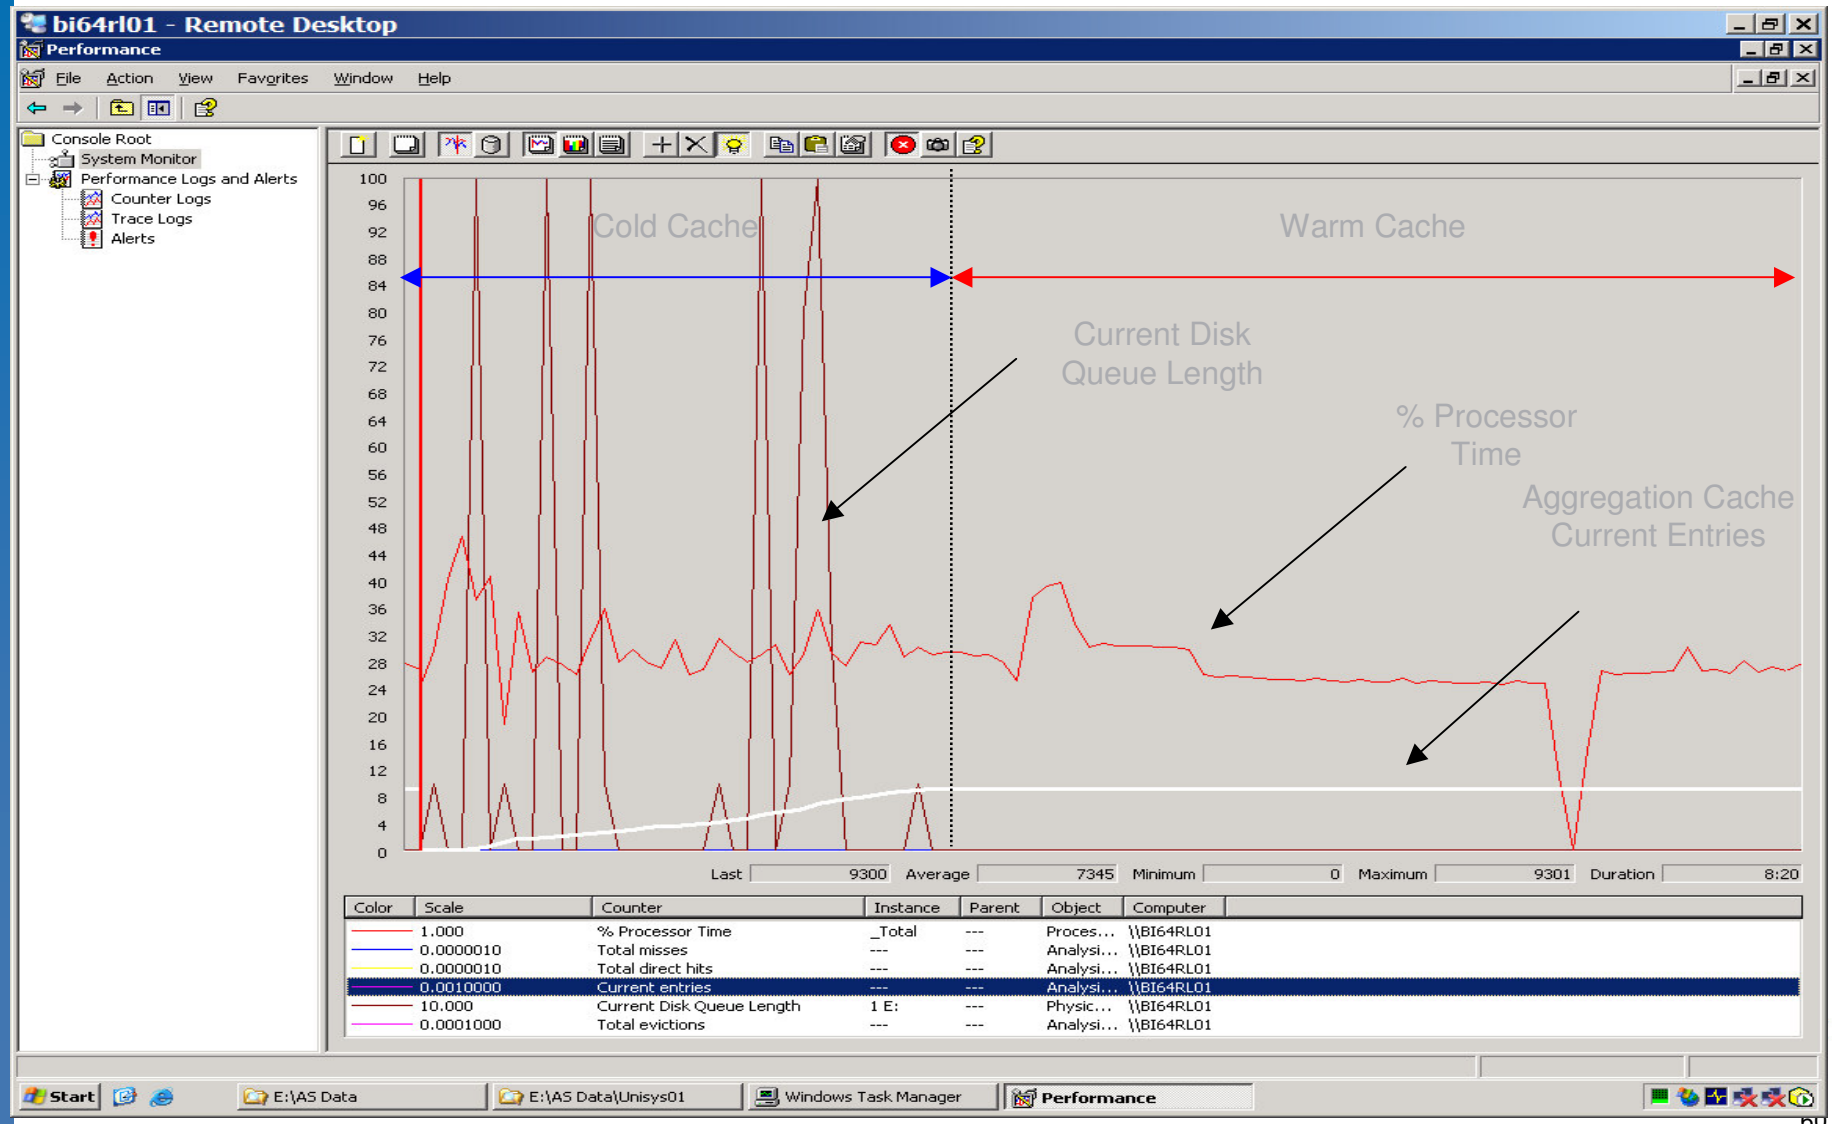

Conclusions

Lessons Learned

- 64-bit is not *directly* about performance, it's about capabilities.
- Client memory consumption with large dimensions
  - Particularly important for running PTS in the middle tier → Memory as high as 1 GB per client when querying Customer dimension.
  - Create as many intermediate levels as possible.
  - Combine cubes into virtual cubes to avoid downloading members *n* times.
- Analysis Services serially loads dimension members at startup.
  - Our server would take about an hour and a half to start up.

32 bit versus 64 bit

A. Understand 64 bit product features
   - DTS and management tool set is still 32 bit
   - What is your tool set?
     - ETL
     - OLAP
     - Modeling
     - Scheduling
     - Monitoring
     - Reporting

B. New Application or migration from other products
   - Can you estimate (peak) load?

C. Existing application from 32 bit SQL
   1. No performance issues, leave it alone
   2. Has DBA tuned application (if allowed)
   3. Determine bottle neck from performance counters
   4. Bottleneck = LAN, solve it via LAN
   5. Bottleneck = disk, solve it via disks/SAN
   6. Bottleneck = RAM, can we add RAM or go to 64 bit (>3GB working set)
   7. Bottleneck = CPU, licensing mode, add CPU, go to 64 bit

D. MSFT allows moving licenses from 32 bit to 64 bit

SQL Server 2000 64-bit

- Consider upgrading to 64-bit when:
  - Using very large DB's with many concurrent users
  - query plans indicate use in-memory Hash Joins, Hash Aggregates and Sorts
  - applications that involve multiple threads of work, where performance would be enhanced by true parallelism on the processor
  - Queries that can use parallel query execution with more memory
  - Analysis Services with cubes and a large number of dimensions
  - Server consolidation – consider multiple instances of SQL Server
  - applications that use complex and compute-intensive calculations
  - Expected growth in database size and
  - Expected growth in number of users

HP WORLD 2004

Solutions and Technology Conference & Expo

Co-produced by:

interex

shared knowledge • shared power

encompass

AN HP USER GROUP

RECOMMENDED TRAINING VENUE FOR THE

HP Certified Professional

SQL Server 2000 64-bit

- Design Goals –
  - improve database performance for large complex queries through large memory addressability and
  - reduced disk swapping as compared to 32-bit server with AWE enabled
  - Code-compatible T-SQL
  - No changes in on-disk database files

SQL Server 2000 64-bit

- Analysis Services (AS)
  - Can effectively use larger memory addressability for higher performance
  - benefits from parallelism in partition processing
  - cube queries & processing faster
  - very large dimensions can be supported in memory.

Itanium Processor Family

- Itanium (IPF) Architecture
  - Twice as fast versus 32-bit? Madison (3rd generation IPF) 30-to-50 percent greater performance than McKinley
  - More than just large memory addressability. 128 floating-point and 128 integer registers.
  - EPIC – Explicitly Parallel Instruction Computing
    - Compiler enhancements
      - Predication – execute both sides of a branch in parallel, only the "real" side actually completes execution
      - Speculation – run code before branch direction is computed, check results later, and clean up as needed
    - More instructions per cycle
  - On-chip Cache
    - 3MB L3
    - 6MB L3 in the next Processor release
  - Improved Bus architecture
    - Faster and wider
    - 128-bit at 6.4Gb/sec

Migration Considerations
SQL Server 2000 (64-bit)

- Minimal Migration hurdles
  - Scripting languages should “just work” on SQL Server 2000 (64-bit)
  - Extended Stored Procedures, however, need to be recompiled
    - 64-bit compilers ship with Platform SDK
- Same on-disk format
  - Detach database from 32-bit server
  - Transfer the file to the target system
  - Attach database to 64-bit server
- Inter-operability with other SQL Server installations

Product Positioning (RDBMS)

- A highly scalable database platform for memory intensive, performance-critical business applications
- Improved performance for large complex queries through
  - Large memory addressing
  - Nearly unlimited virtual memory
  - And reduced disk thrashing due to large buffer pools that save on I/O
- Viable alternative to expensive Unix solutions

64 bit Significant Changes

- Windows Installer (Darwin) Setup (Merge Modules)
  - Single Feature Tree
  - Easily integrated into ISV application setup
  - Improved reliability
  - Analysis Services integrated into the database installation procedure
- Leverages Windows Server 2003 64-bit components
  - MDAC 2.8 stack
  - Distributed Transaction Coordinator
  - HTML Help
  - MMC
- New in 64-bit Edition
  - Analysis Services now uses SQL Server to host its repository and query log
  - Supports 8-node Failover Clusters

64 bit Unsupported Features

- Upgrades from SQL Server 6.5 and 7.0
- Remote Installation
- Running in WOW
- Others
  - DBLib
  - ESQL
  - Any Access-based provider/driver (JET)
  - For developers, no VB runtime and any components that require it

Sample bar chart

Subtitle (or chart title) here

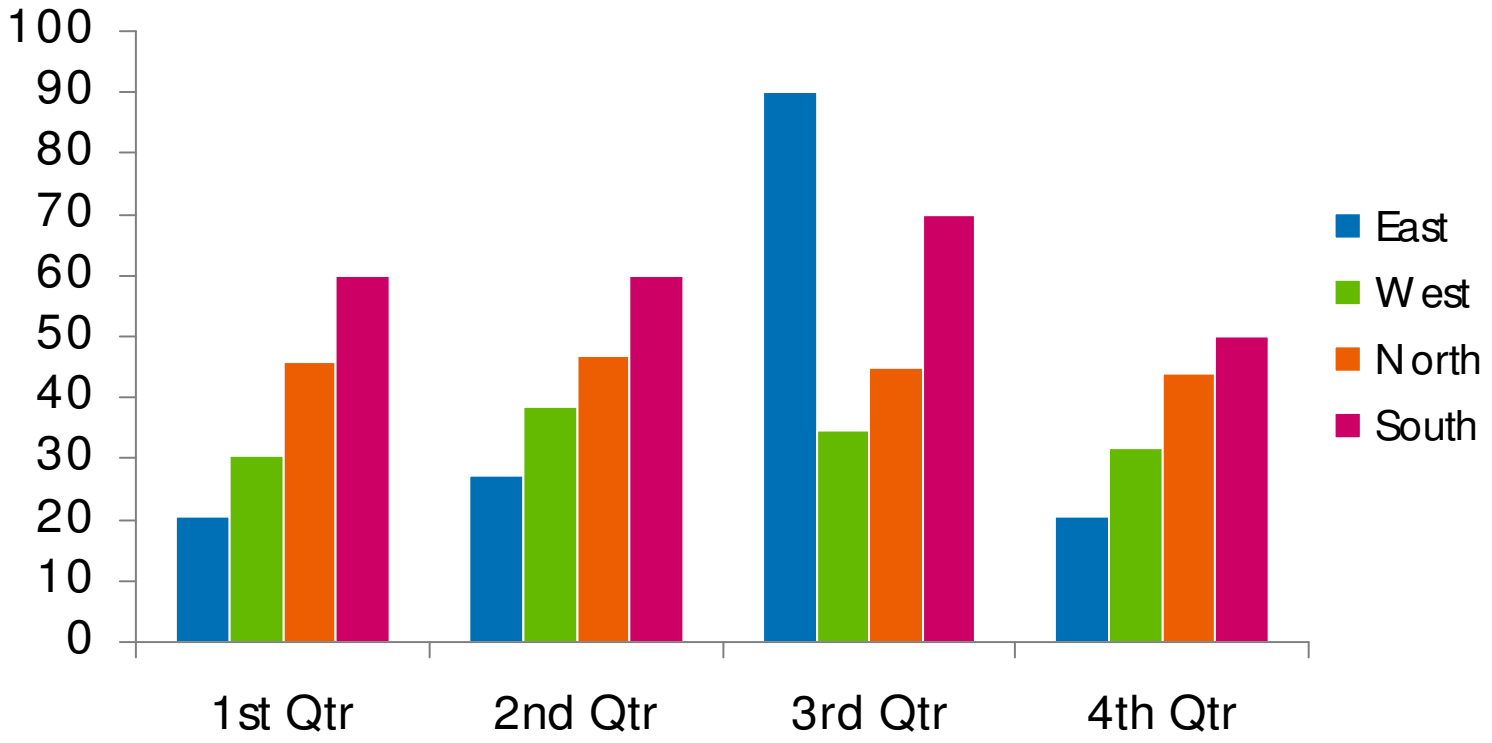

Source: Use this size and position for source information

This is a standard quote slide. Type your quote inside the quotation marks.”

Transition slide headline text goes here

Grid

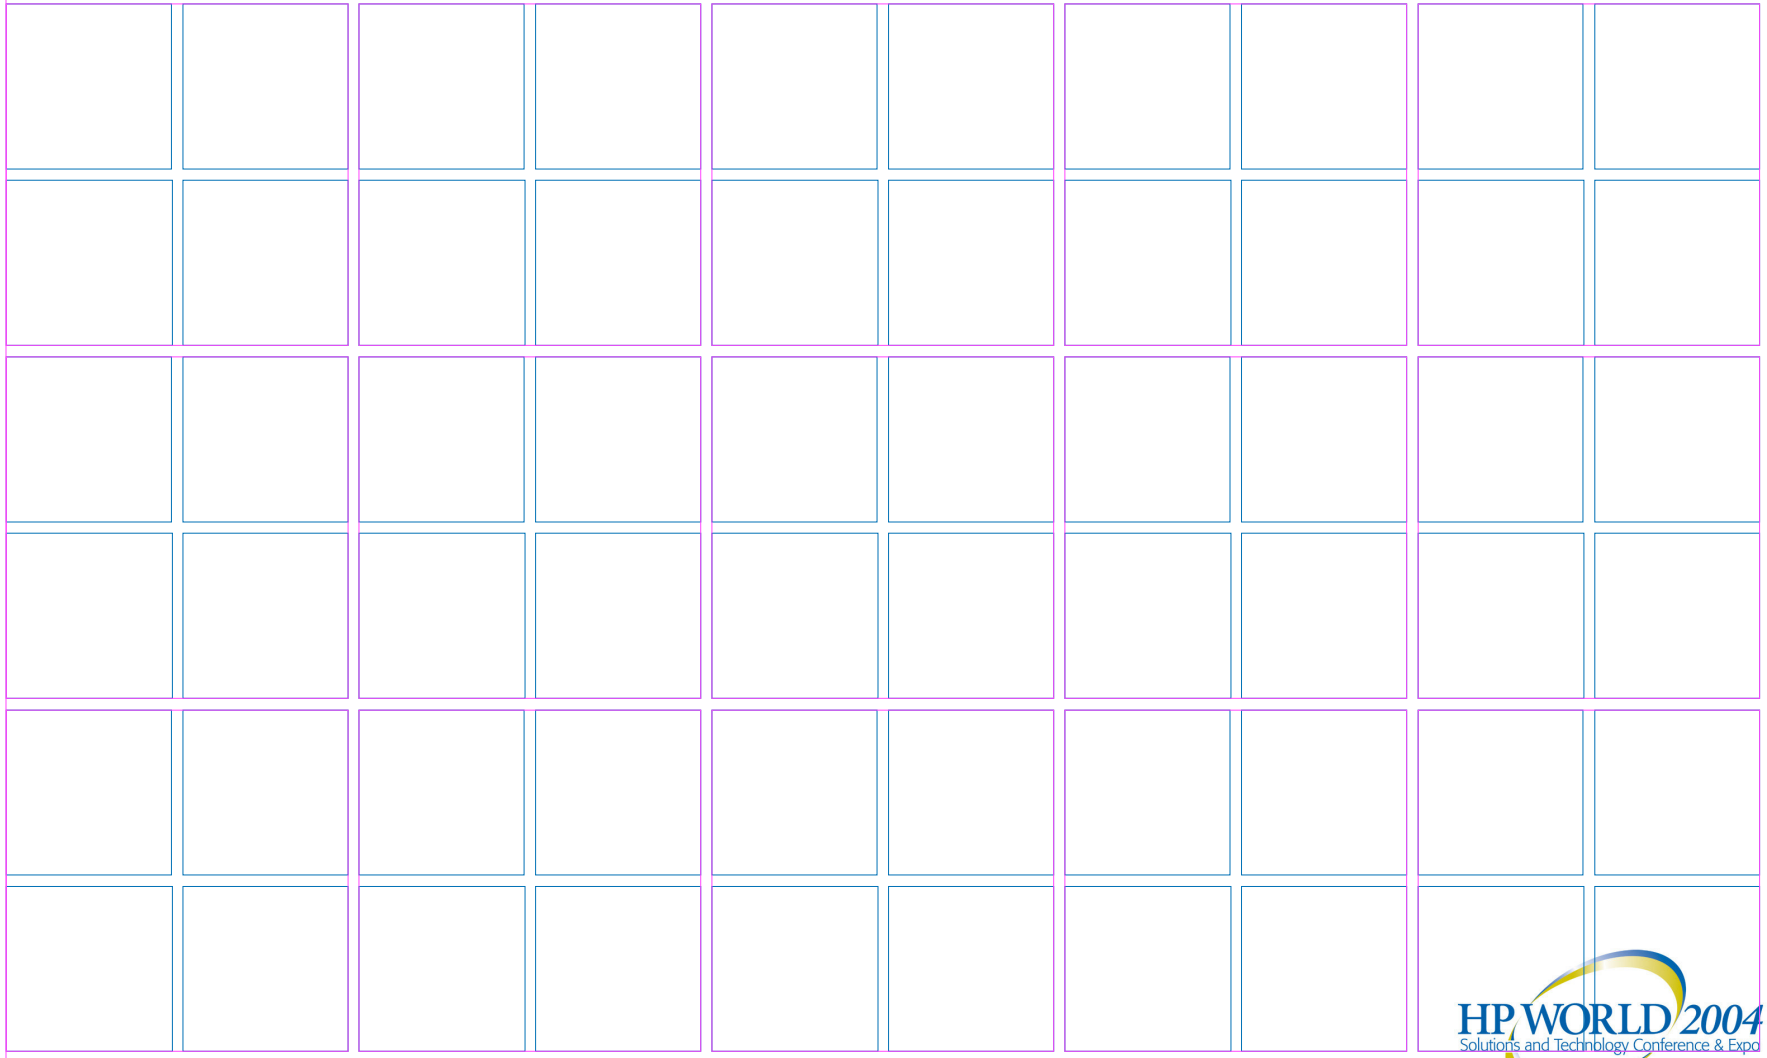

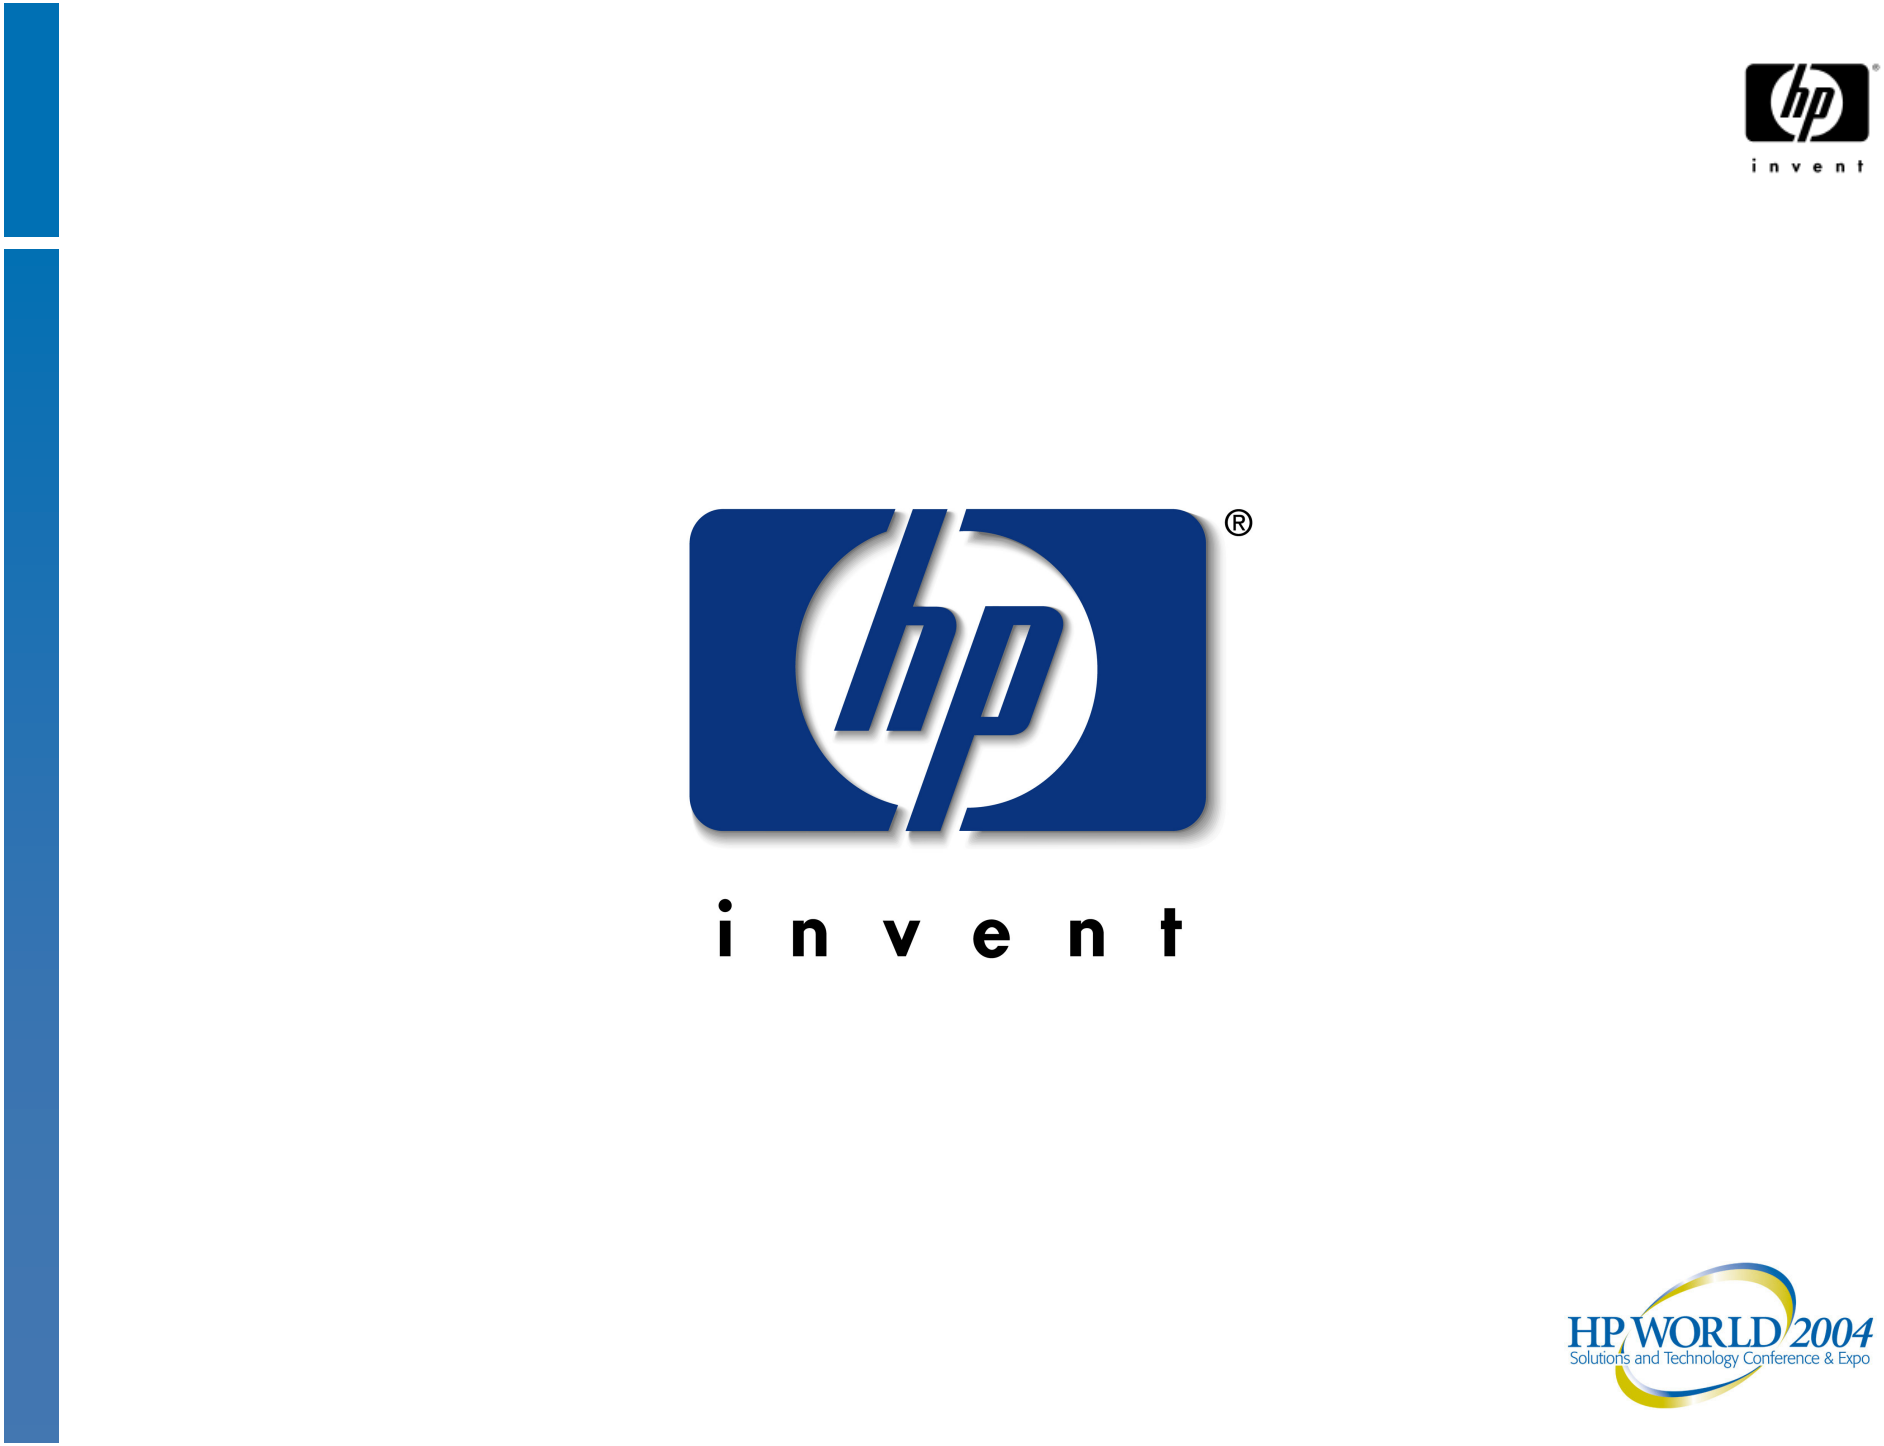
invent
invent
HP WORLD 2004
Solutions and Technology Conference & Expo

How Business Intelligence Helps

- BI serves two main purposes:
  - It monitors the financial and operational health of the organization
    - Reports, alerts, alarms, analysis tools, key performance indicators (KPIs) and dashboards
  - It regulates the operation of the organization
    - Two-way integration with operational systems, information feedback analysis

BUSINESS OBJECTIVES AND ISSUES

- What are the objectives of your organization? What are you trying to accomplish? What are your top priority business goals?
- What are your success metrics? How do you know you're doing well? How often do you measure key success factors?
- What functions and departments within the organization are most crucial to ensuring that these key success factors are achieved? What role do they play? How do they work together to ensure success?
- What are the key business issues you face today? What prevents you from meeting your business objectives? What's the impact on the organization?
- How do you identify problems/exceptions or know you're headed for trouble?
- What do you see as opportunities for additional profit that are not being addressed today?
- Where do you stand compared to your competition in the use of information technology?
- Are you able to respond quickly to market conditions and assure productivity of your staff?
- Describe your products (or other key business dimension such as customer, vendor, etc.). How do you distinguish between products? Natural way you categorize products? How would you narrow a list of thousands of products?
- How often do these categorizations change? What should happen with your business analysis following a change?

ANALYSES REQUIREMENTS

- What role does data analysis play in decisions that you and other managers make to run the business?
- What key information is required to make or support the decisions you make in the process of achieving your goals and overcoming obstacles? How do you get this information today?
- Is there other information which is not available to you today that you believe would have significant impact on helping meet your goals?
- Which reports do you currently use? What data on the report is important? How do you use the information? If the report were dynamic, what would the report do differently?
- What analytic capabilities would you like to have?
- What opportunities exist to dramatically improve your business based on improved access to information? What's the financial impact?
- What type of routine analysis do you currently perform? What data is used? How do you currently get the data? What do you do with the information once you get it?
- What analysis would you like to perform? Are there potential improvements to your current method/process?
- What type of on-the-fly analysis do you typically perform? Who requests ad hoc analysis? What do they do with the analysis? Do you have time to ask the follow-up questions?
- Are there specific bottlenecks to getting at information?
- How much historical information is required?

USER SUPPORT / ANALYSES AND DATA REQUIREMENTS

- What is the current process used to disseminate information?
- What tools are used to access/analyze information today? Who uses them?
- Are you asked to perform routine analyses? Do you create standardized reports?
- Describe typical ad hoc requests. How long does it take to fulfill these requests?
- Who are the most frequent requesters of analysis and/or data?
- What is the technical and analytical sophistication of the users?
- Describe the current user support mechanism - centralized vs. experts located in user departments?
- What is the biggest bottleneck/issues with current data access process?

DATA AVAILABILITY AND QUALITY

- Which source systems are used for frequently-requested information?
  - How do production systems relate to each other? Which systems feed others?
  - What is the granularity?
  - How often is the data updated? Availability following update?
  - How much history is available?
  - What is an estimated size of this data (preliminary # of rows)?
- What are the known data gotchas in current source systems?
  - Which fields are not populated (e.g., not required and/or validated at input)?
  - Are there dual-purpose fields depending on context?
  - What is the availability of decodes? Are they buried in reporting programs?
- What master files do you have? Describe the maintenance of these master files.
  - Do you currently have common source files?
  - Who maintains the source files?
  - How are keys maintained? Are keys reassigned?
  - What is the cardinality (# distinct values)?
- Frequency of hierarchy changes within key business dimensions (product, vendor, facility…) How are changes captured?