

Information Management



# IBM Information On Demand Conference 2009

Optimize Your Business for  
Greater Results with  
InfoSphere Warehouse

Chi-Chung Hui

*Senior I/T Specialist*

*IBM Software Group, Hong Kong*



# Agenda

- InfoSphere Family Overview
- InfoSphere Warehouse
  - Modular Design
  - Performance and Scalability
  - No Copy Analytics and Cubing Services
  - Deep Compression
  - Workload Management





# The IBM InfoSphere Vision



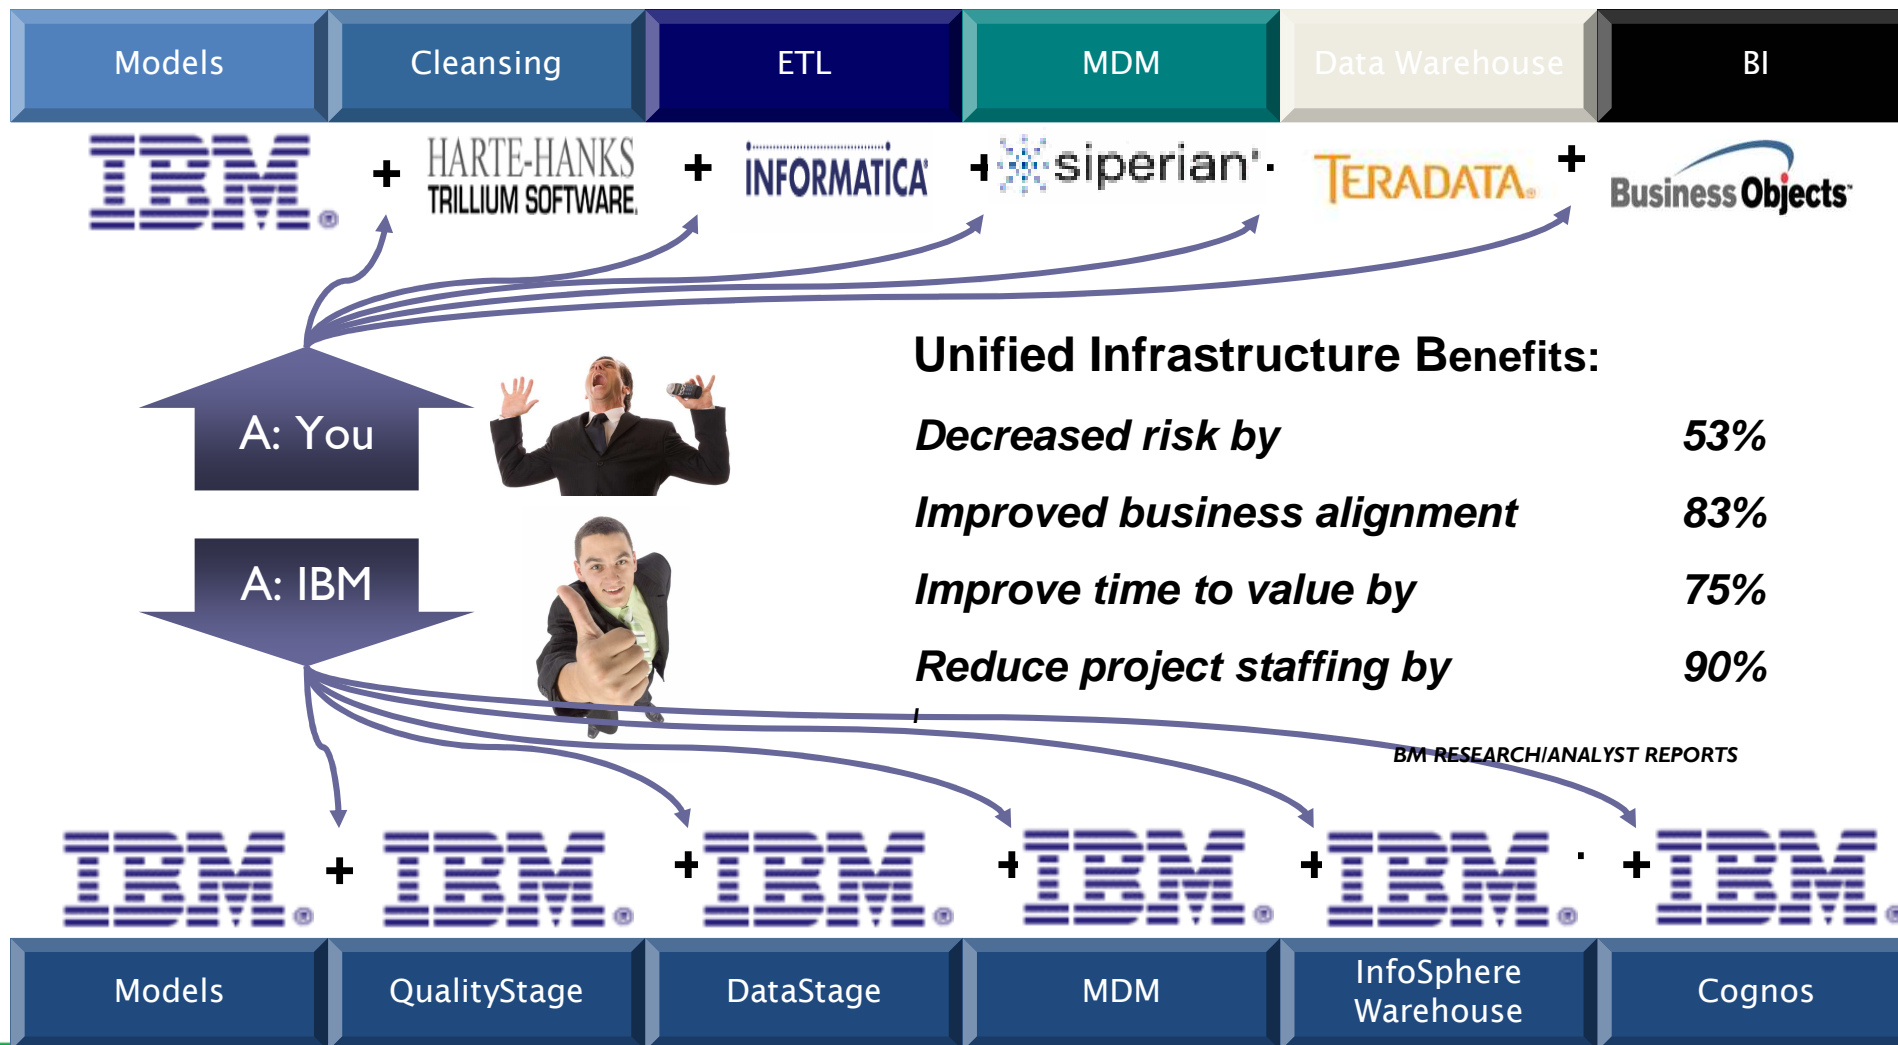
## An Industry Unique Information Management Platform

- Simplify the delivery of Information
- Accelerate client value
- Promote collaboration
- Reduce risk
- IBM Industry Models Integrated
- Project to Enterprise



# IBM Positioning in Business Intelligence

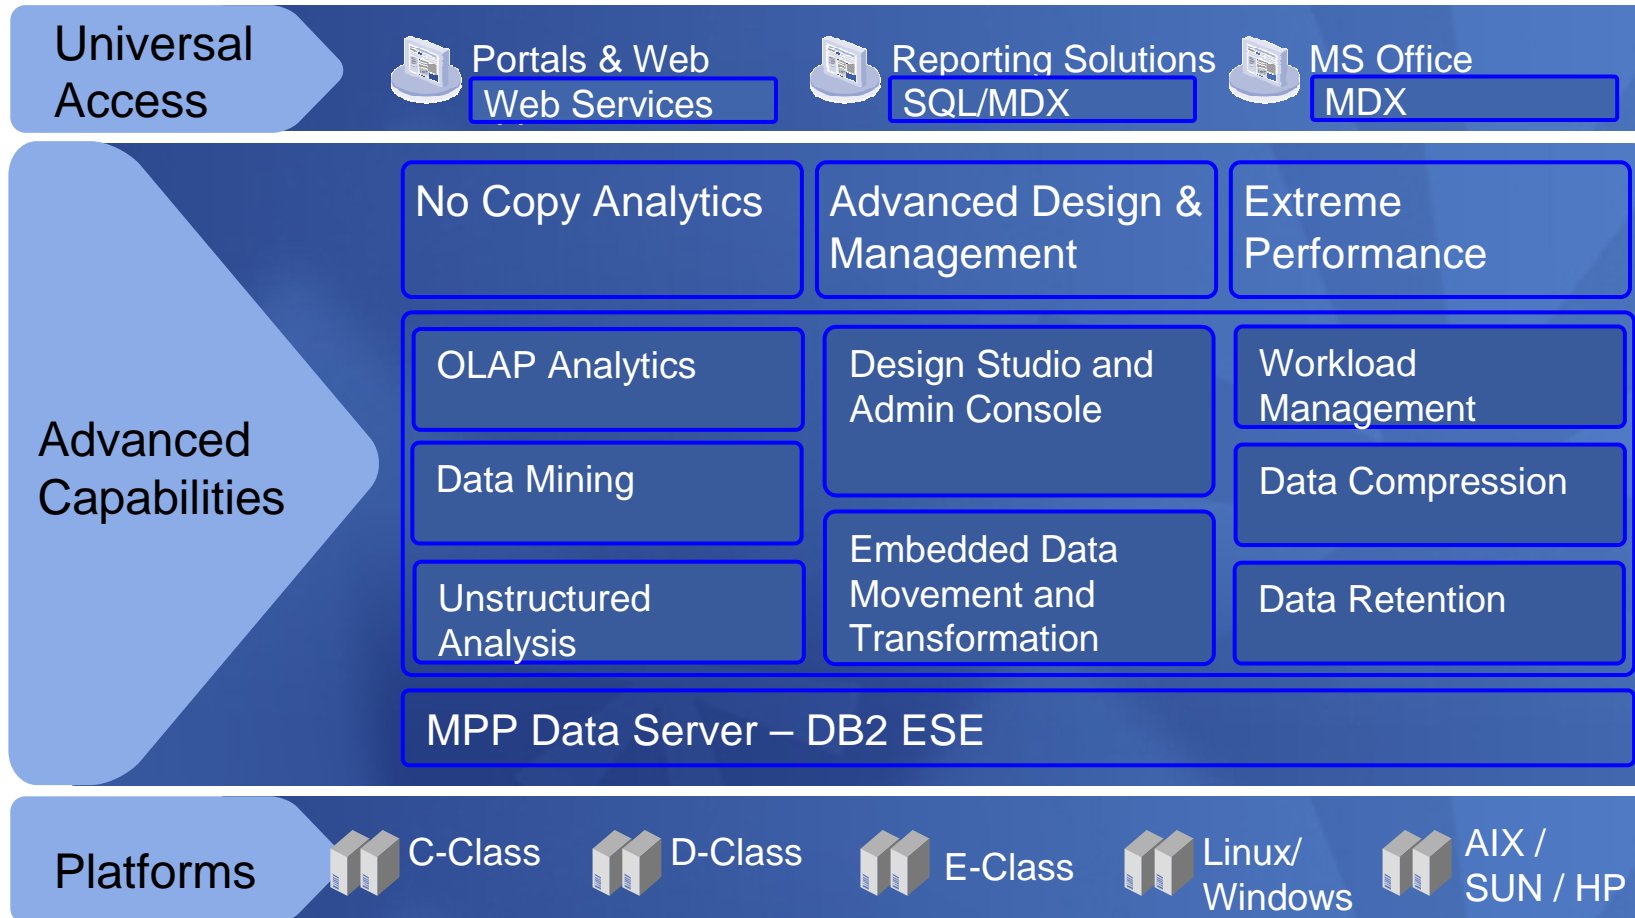
## Single Integrated Infrastructure Of Best Of Breed Technologies



# InfoSphere Warehouse



## Full Featured Data Warehouse Platform



# What is InfoSphere Warehouse?

The InfoSphere Warehouse is...

1. A **proven building block methodology** that focuses on balance as a key design principle for a warehouse. This methodology has been used by IBM for years.
2. IBM Balanced Warehouse solutions are **rigorously tested and tuned** for performance and reliability.
3. Before a Balanced Warehouse is released, **IBM's laboratories** configure and stress test the systems for maximum reliability and security as well as fine tune them to ensure that the highest levels of query performance are met.

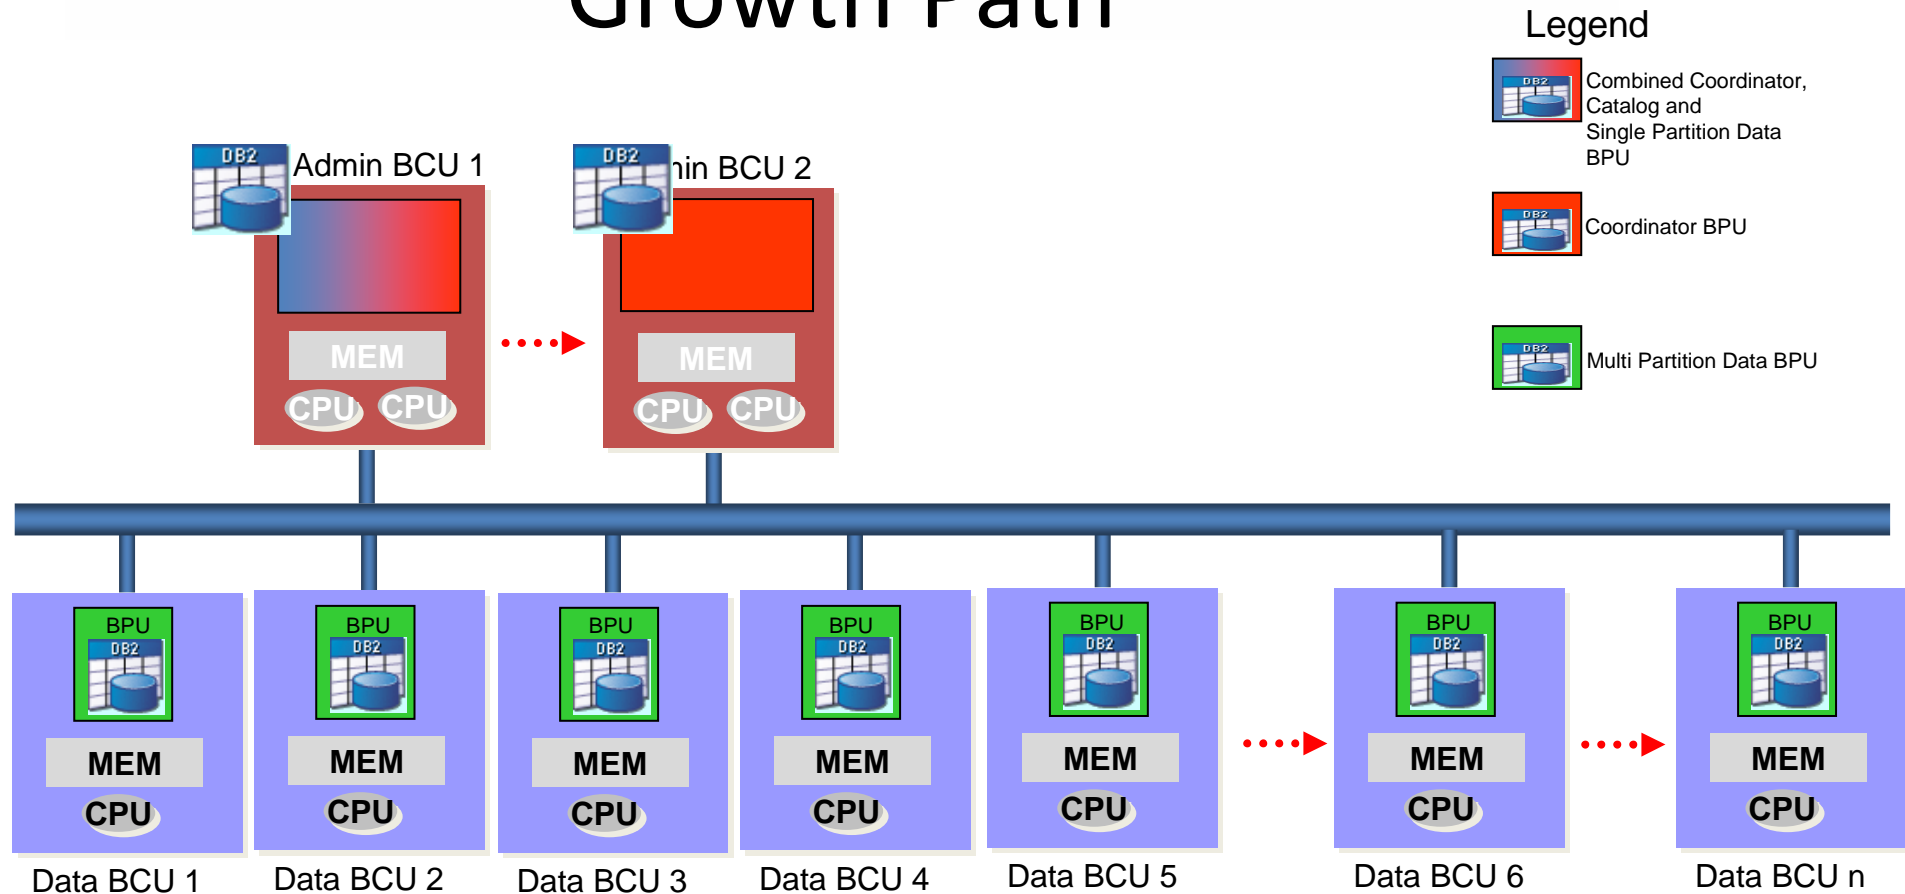
**InfoSphere Warehouse is a solution bundle that takes the “guess work” out of configuring a data warehouse**

**A set of document and best practice exists that help you to build a successful data warehouse!**





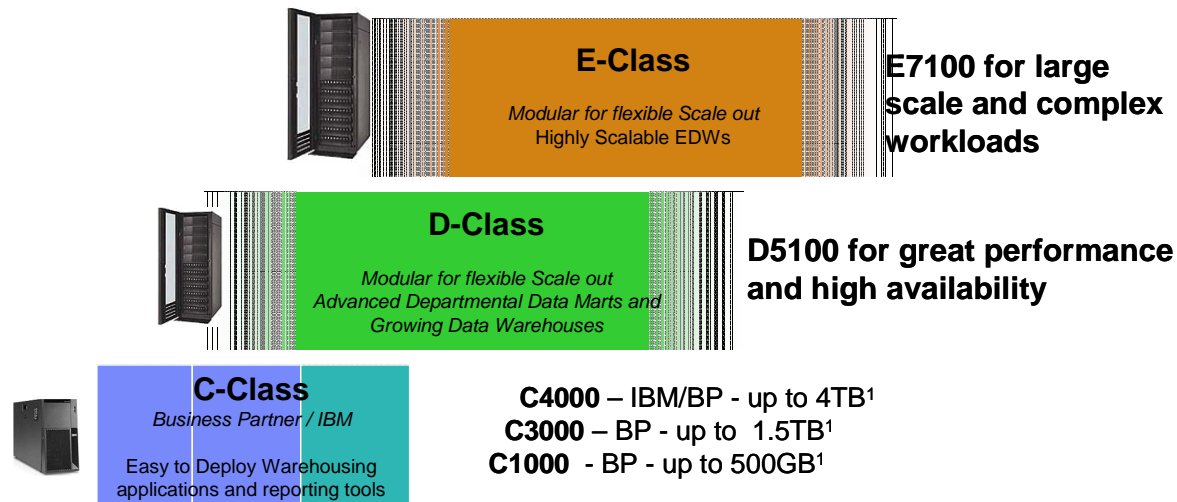
# Incremental Investment: Typical Growth Path



# IBM InfoSphere Balanced Warehouse

*Pre-built, Pre-packaged Warehouse Appliance Offering*

- **Reduces Risk**
  - Standard Pre-tested Configuration
  - Ongoing Supported Package
  - Predictable performance
- **Reduces your Time to Value**
  - Create Table Ready
  - Faster to Deploy
  - Built for Performance
  - Easy to Integrate into any environment
- **Reduces Warehouse Cost of Ownership**
  - Automated Management
  - Single Product Support
  - System Built for Performance
- **Configured for your Business Requirements**
  - Multiple configurations
  - Modularly expandable
  - Available in different sizes




<1      3      10      50      150+


<sup>1</sup> Raw Active User Data




# Features for Data Warehousing Solutions



- **Distribute** data across nodes (**DPF**)
- Range partition the rows in table partitions per node (**Range Partitioning**)
- Cluster the data by multidimensional column values (**MDC**)




- Precalculate summary / aggregate data by group by's into **MQTs**.



- Reduce storage and increase performance with advanced **Compression**.

# No Partitioning

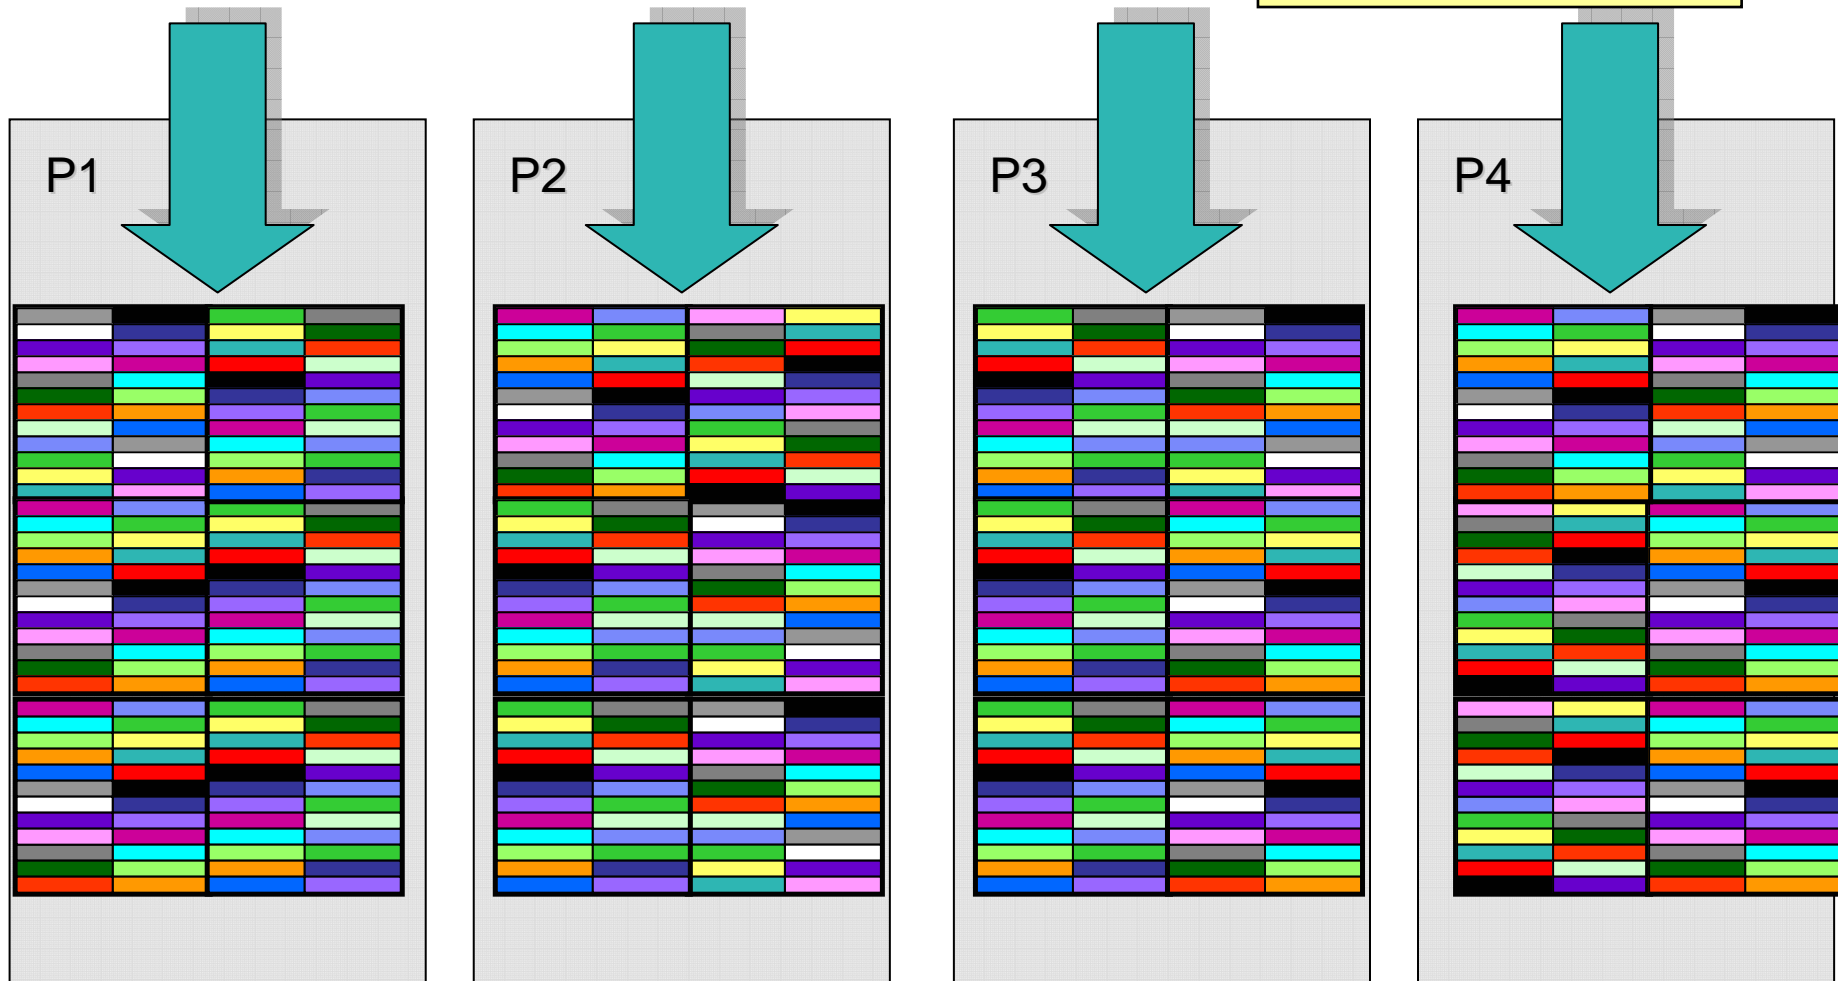
Speedup = 1



# Divide Data and Put in Multiple Servers

## *Divide & Conquer Parallelism*

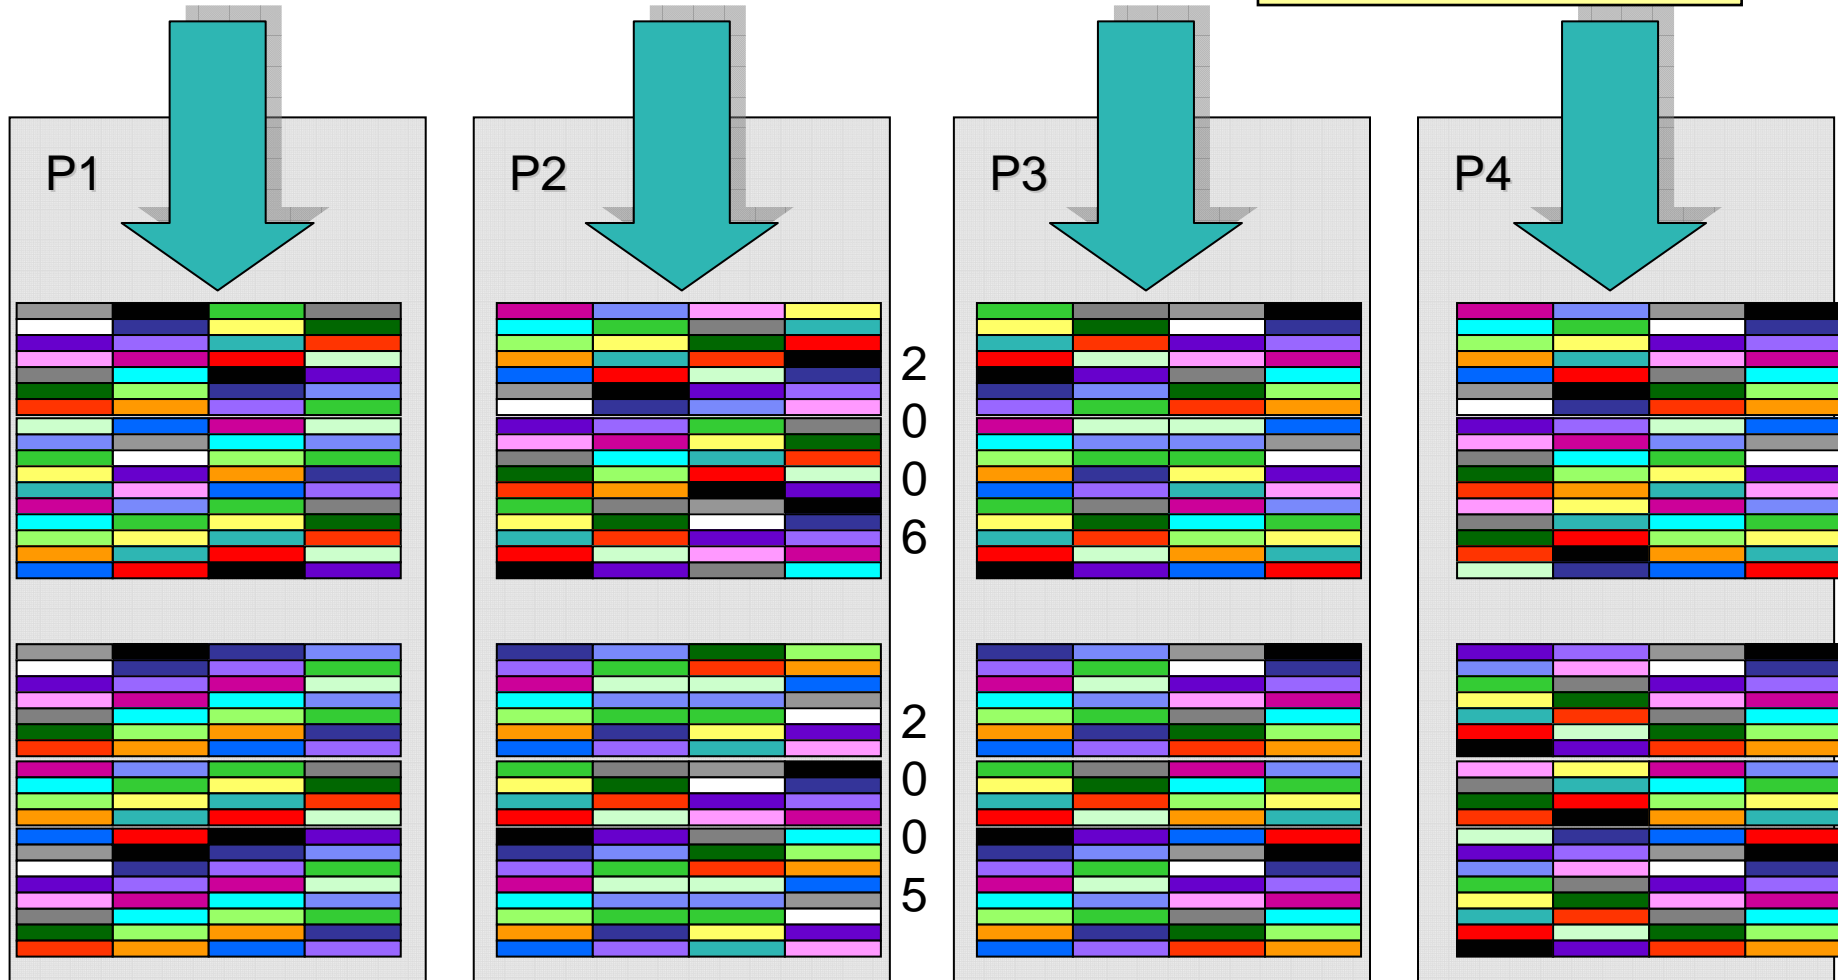
Speedup = 4



# Add Range Partition: Eliminate Unused Partitions

*When analyze 2006 data only...*


**Speedup = 8**



# Further Group Data by Types (Assumed 6 Types)

*When analyze 2006 purple data only.*

**Speedup = 48!**



**Before any Index is Built At All**

# Single Version with *Near Real-time Warehouse*

Optimizing the business with right-time information

Giving customers the right information to enable bigger sales

Customer Self Service

Customer Service/ Care

Online Sales

Marketing/Sales Interactions

Focus is on having the right information at the right time

Financial Trading

Financial Planning

Making the right decision at the right time

Fraud Avoidance

Understanding the choices and planning with accurate information

Regulatory Compliance

Supply Chain Optimization






# InfoSphere Warehouse and No Copy Analytics

- The No Copy approach to delivering business analytics has many benefits
  - Enables low latency, right-time Analytics
  - Increased Flexibility and adaptability to information changes
  - Reduced development and Management Effort
  - Flexibility to add and change Analytical Applications cost effectively
- Leading to outstanding business results across the spectrum, including
  - Business trend detection and analysis
  - Fraud Detection
  - Customer Churn Prediction
  - Event Based Decision Support
- Overall Greater Return On Investment and Lower Cost of Ownership!

# How does No-Copy Analytics work?



**COGNOS**  
AN IBM COMPANY



DataQuant




Microsoft Office




3rd Party





## Universal Access




Data Mining



Unstructured Text  
Annotation



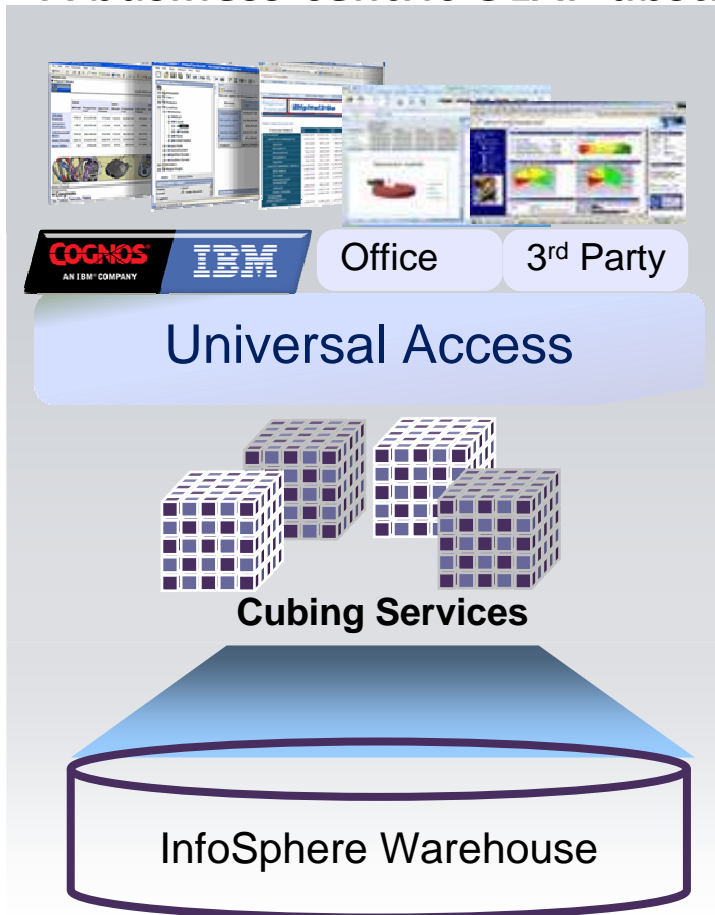
Cubing



- No data extraction
- Data always (only) in the data warehouse
- Full (universal) access to all analytical information
- Low latency response to new/changed data

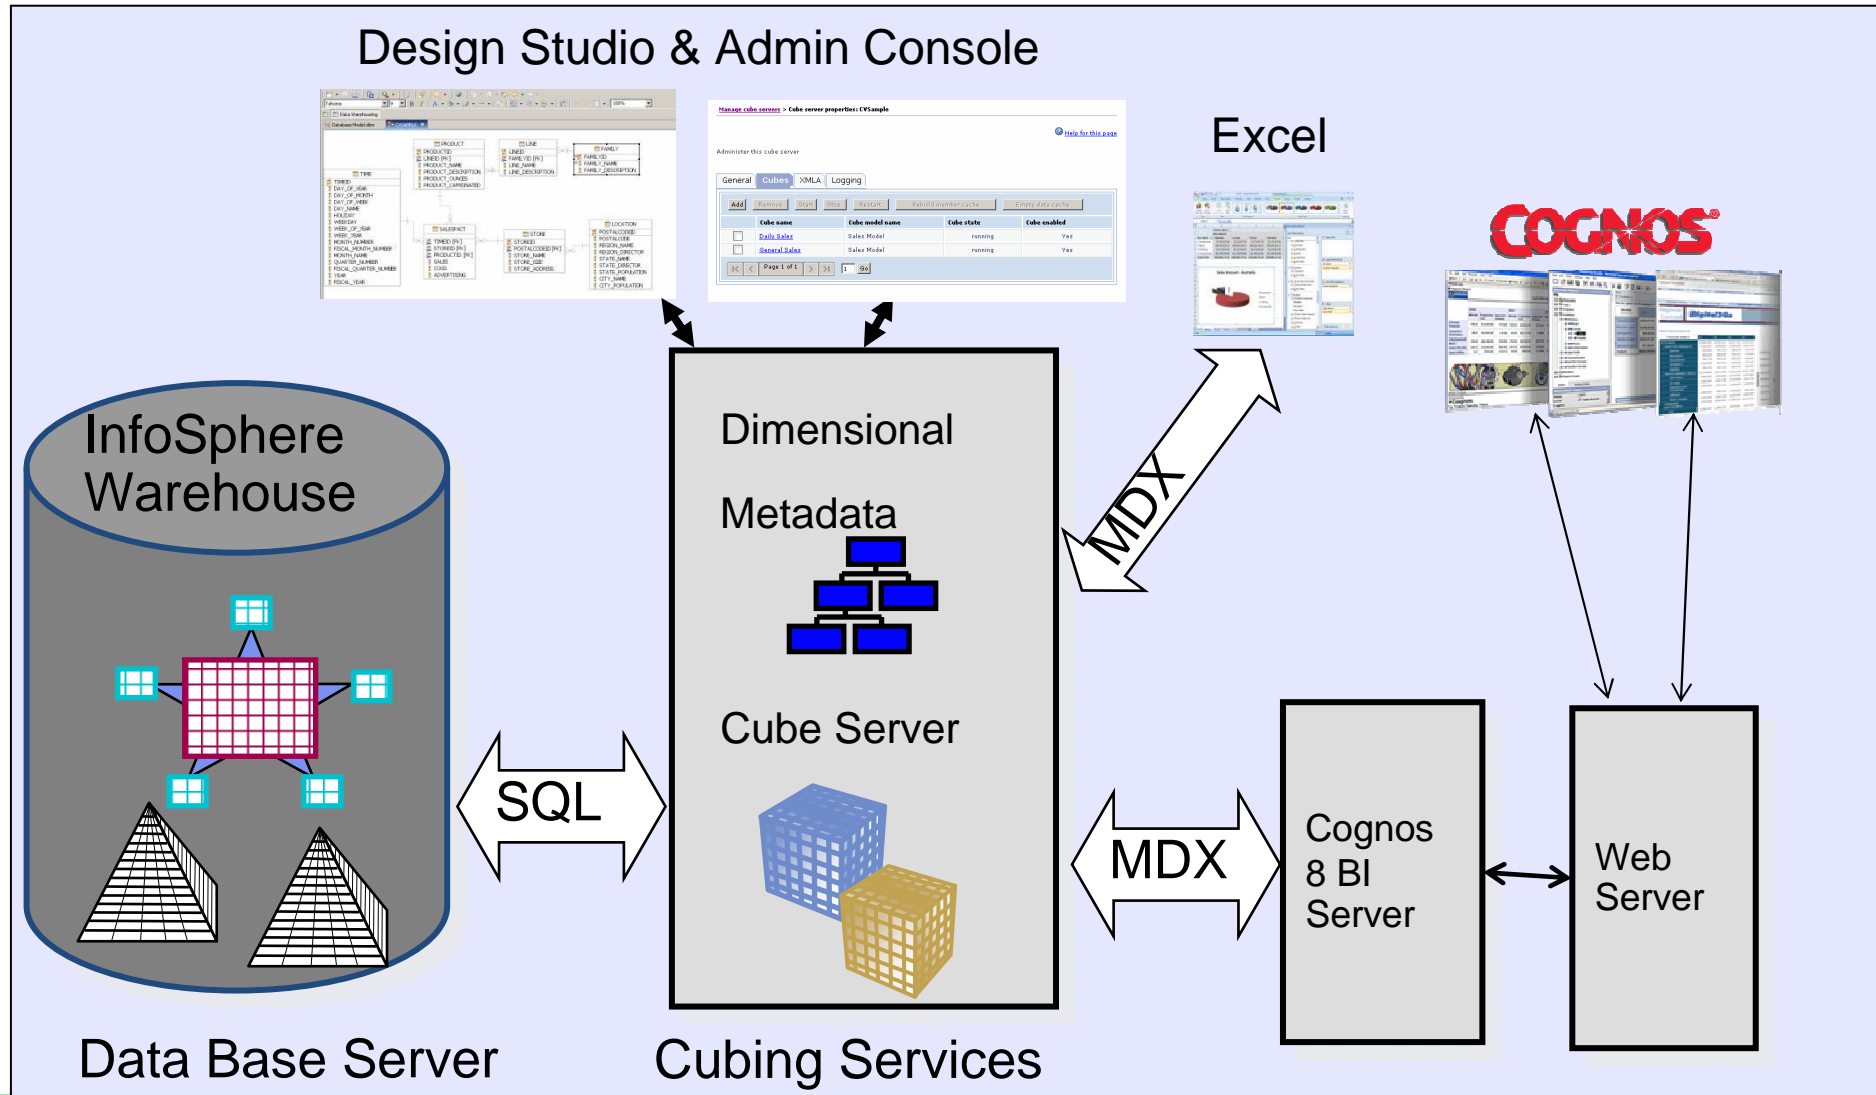
# InfoSphere Warehouse Cubing Services

*A business-centric OLAP abstraction of the data warehouse information*

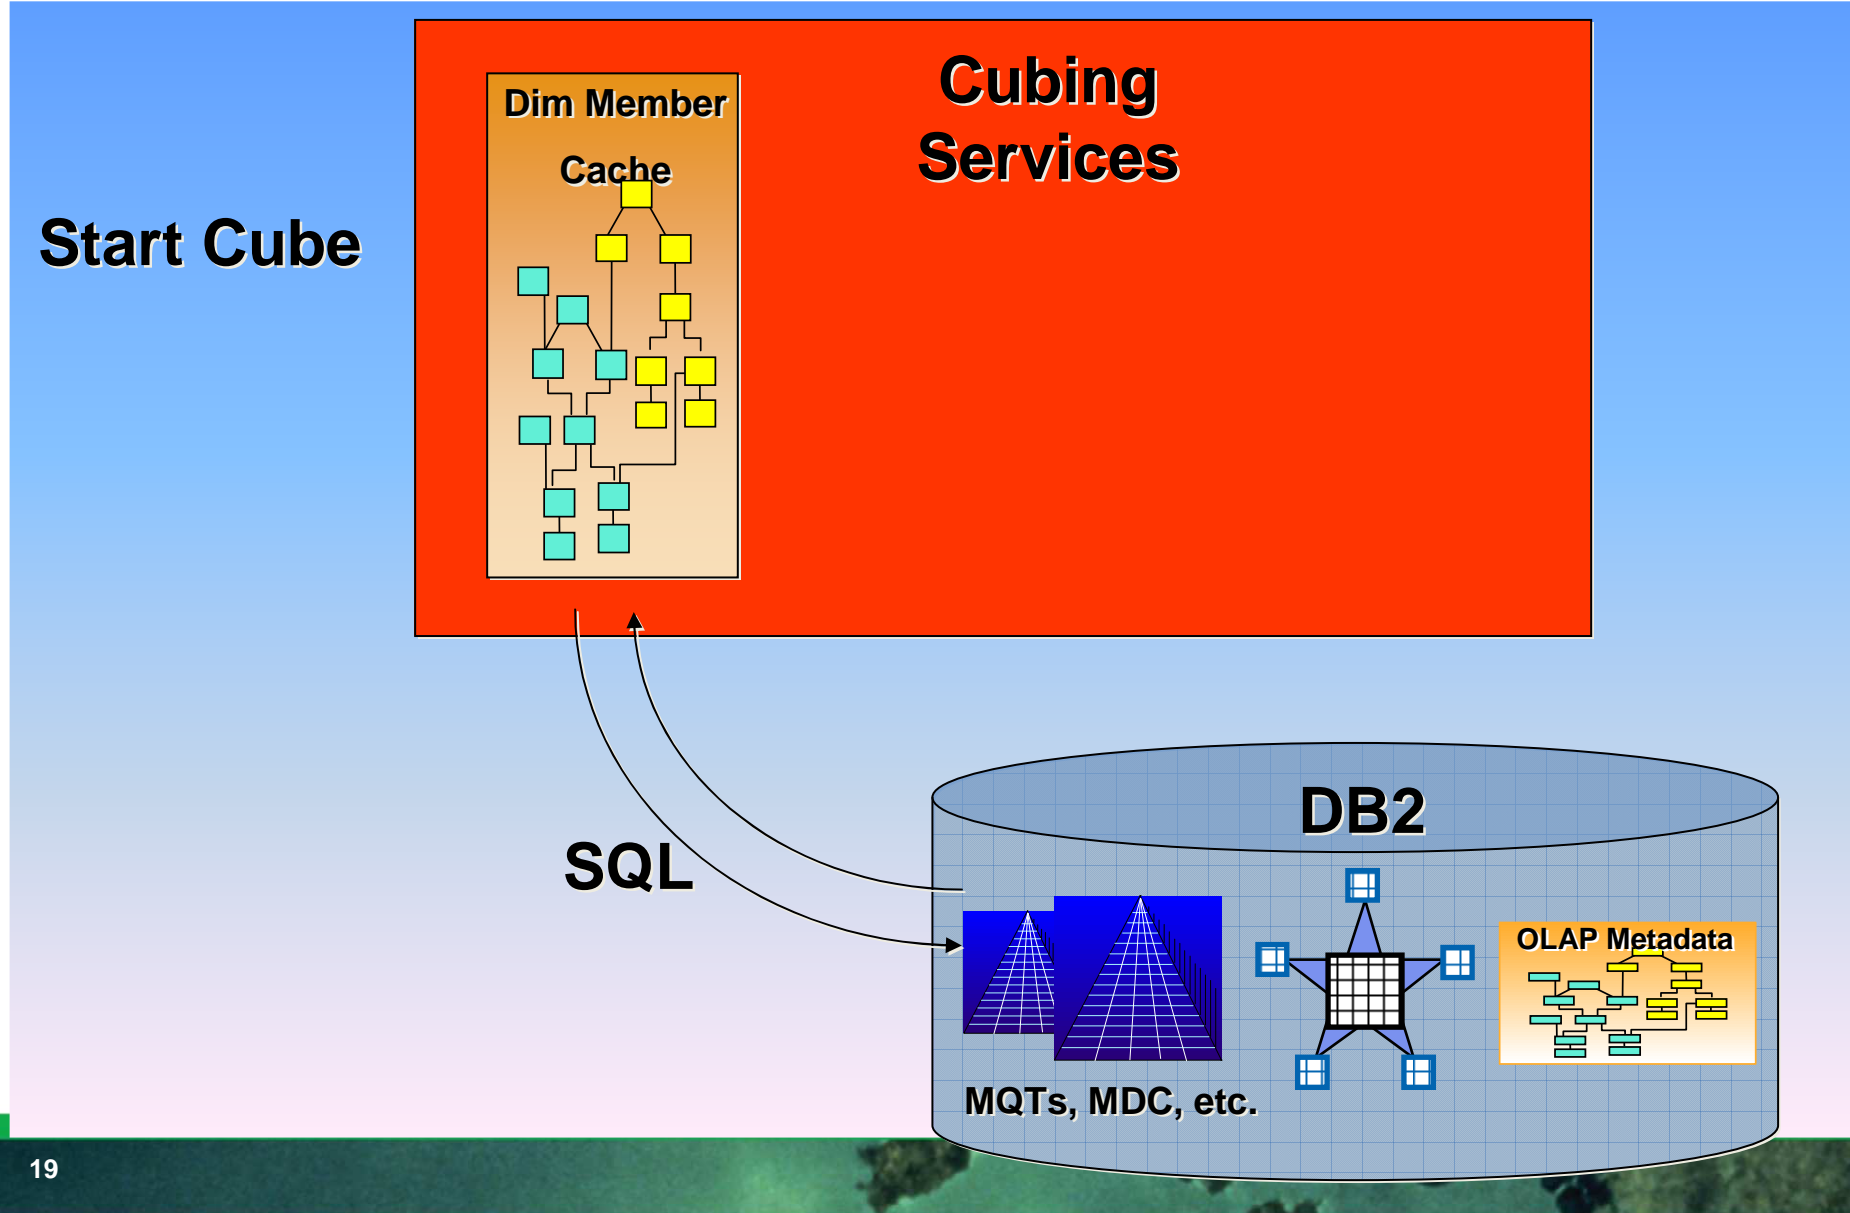


- Benefits
  - Scalable low-latency OLAP
  - Very large dimensions over very large data sets
  - Optimized for the InfoSphere Warehouse
  - Integrated design and maintenance
  - No copy analytics with universal heterogeneous BI tool access
  - Support MDX, ODBO, XMLA
  
- Business Value
  - Enriched reporting and decision support for more members of the organization – greater business impact from information assets
  - Accelerated delivery of high performance business intelligence solutions across a spectrum of business processes and industries.

# Cubing Services Architecture with Cognos

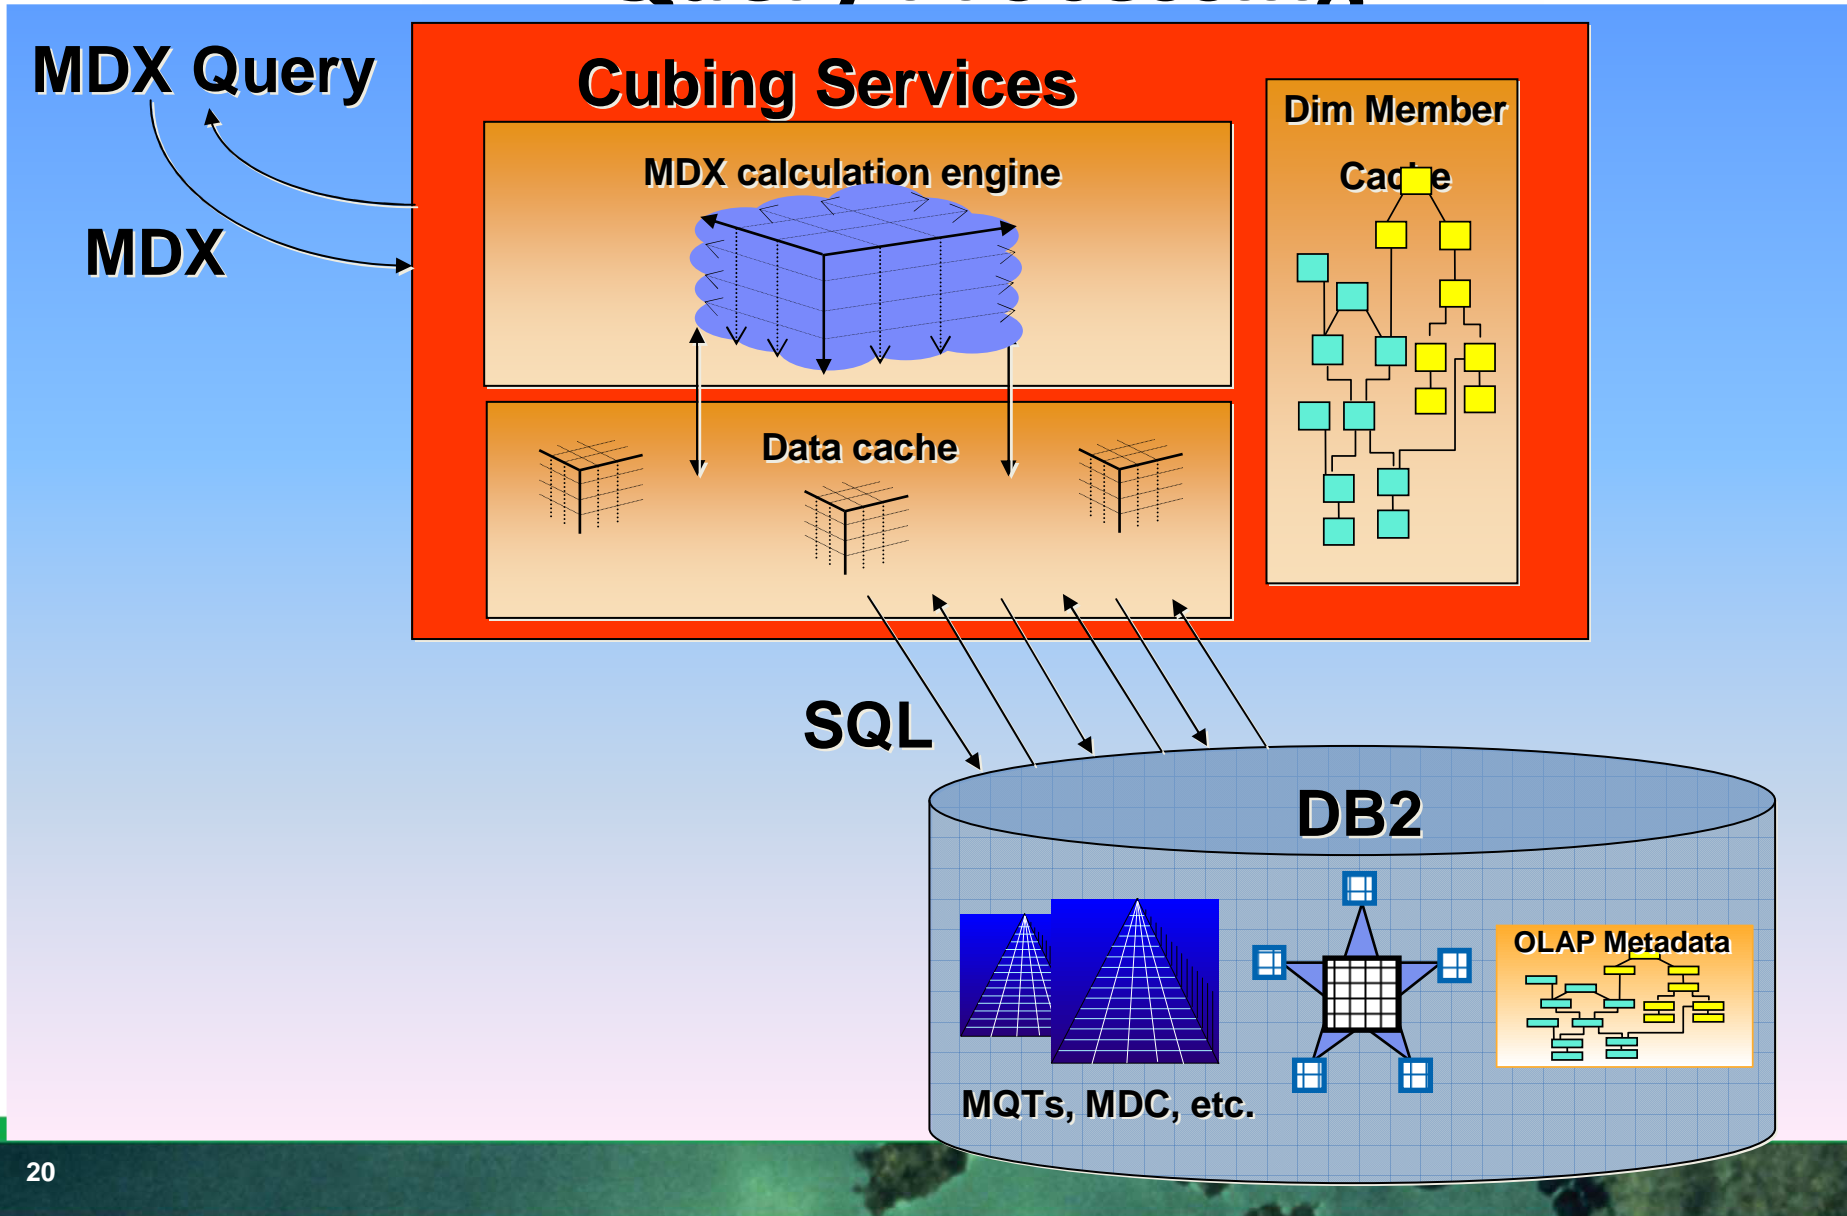


# Cube Server in Action – Startup






# Cube Server in Action – Query Processing





# Cubing Services and BI Tools



**IBM COGNOS 8 Framework Manager**  
Framework Manager allows you to model metadata and publish packages.

**IBM COGNOS 8**

**Projects**  
Create a new project  
Open a project...

**Recent Projects**

Name
IoD Demo Proje
CS Project 1

**My Content**

- Home
- IBM Cognos content

**Administration**

**My Actions**

- Query my data
- Analyze my business
- Create professional reports
- Manage my events




Cubing Services

## Deep Compression Using a Compression Dictionary

- Repeating patterns within the data (and just within each row) is the key to good compression. Text data tends to compress well because of reoccurring strings as well as data with lots of repeating characters, leading or trailing blanks

Name	Dept	Salary	City	State	ZipCode
Fred	500	10000	Plano	TX	24355
John	500	20000	Plano	TX	24355




# Row Compression

Uncompressed Row




Compressed Row




Common sequences of consecutive bytes in row replaced with 12 bit symbol

Data page with uncompressed rows




Data page with compressed rows



- Less storage
- Over **70%** off in storage saving in real cases!
- Much faster, for example, table scan

# Even Faster Performance with Compression

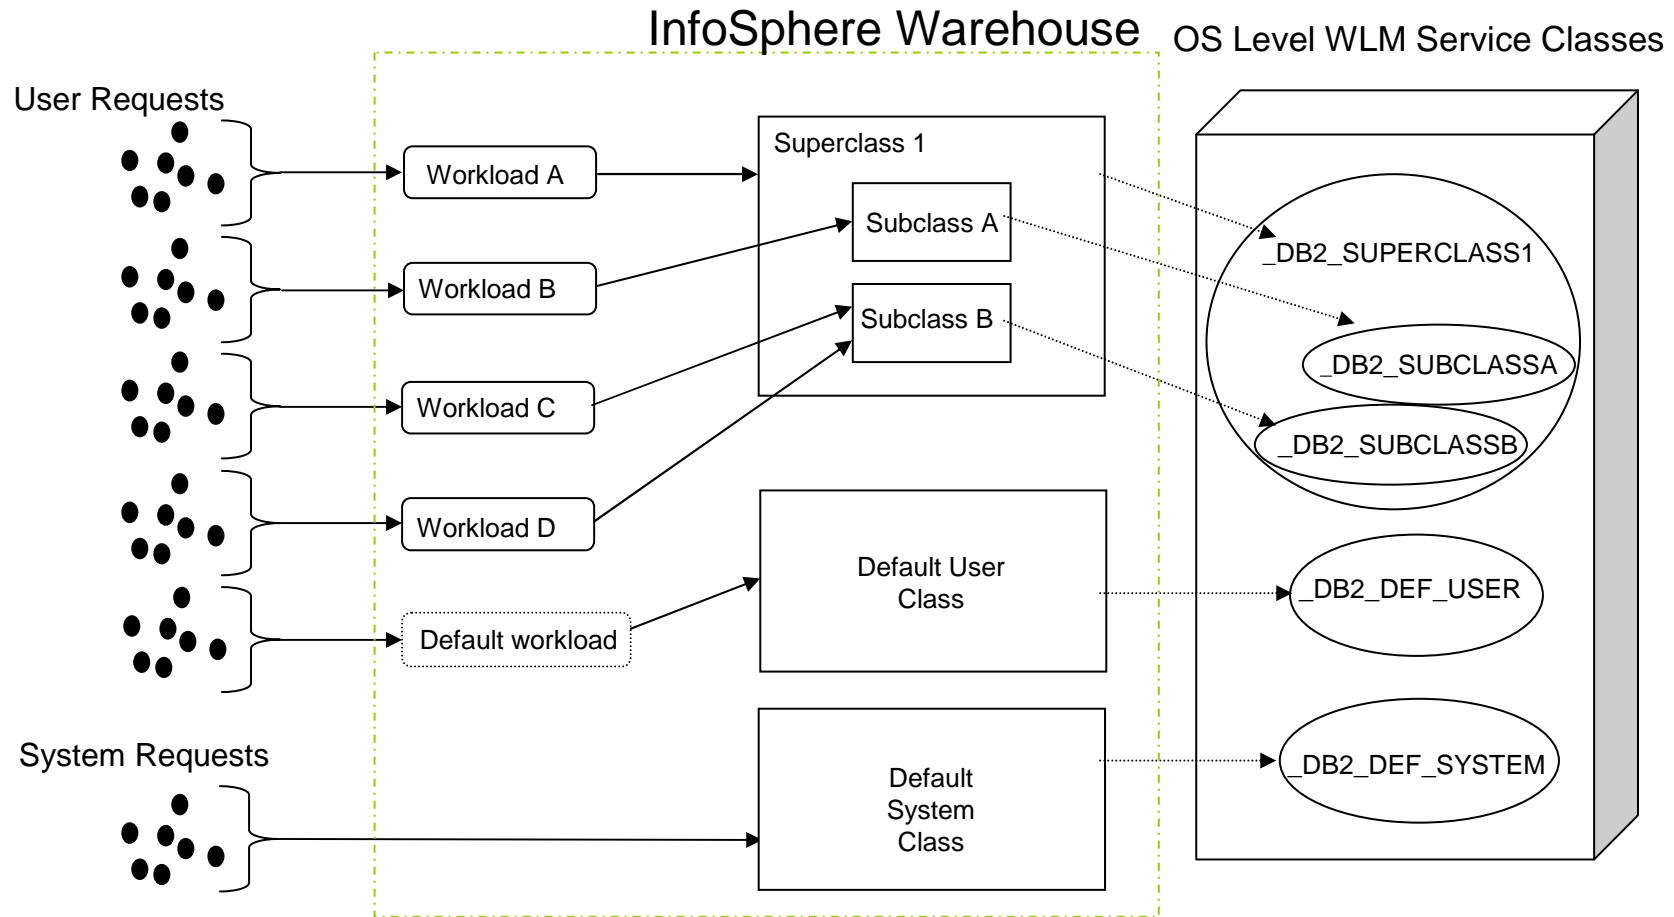
Overall Workload Performance:  
Uncompressed Data versus Compressed Data



"With the new compression technology in DB2 Viper, we realized an **80 Percent improvement in space savings** for our most critical tables in our data warehouse. We were even more pleased with this technology when we found that Viper's compression capability helped us process queries to the database an average of **40 Percent faster** than before. We're looking forward to seeing the same results with our Operational Data Store and OLTP systems." Donny Ledbetter, Senior DBA


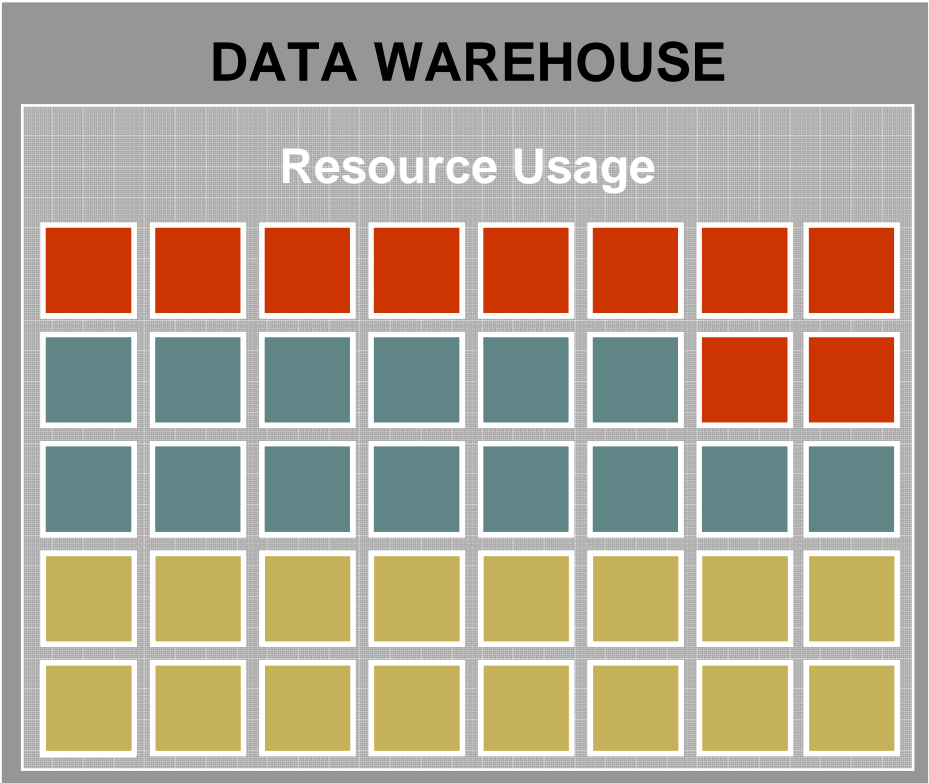
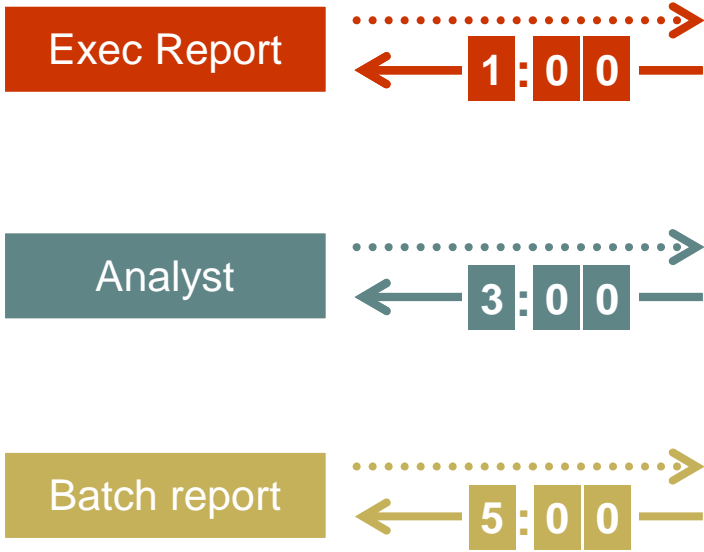
**AutoZone®**

# Workload Management Guarantees Service Level Agreement




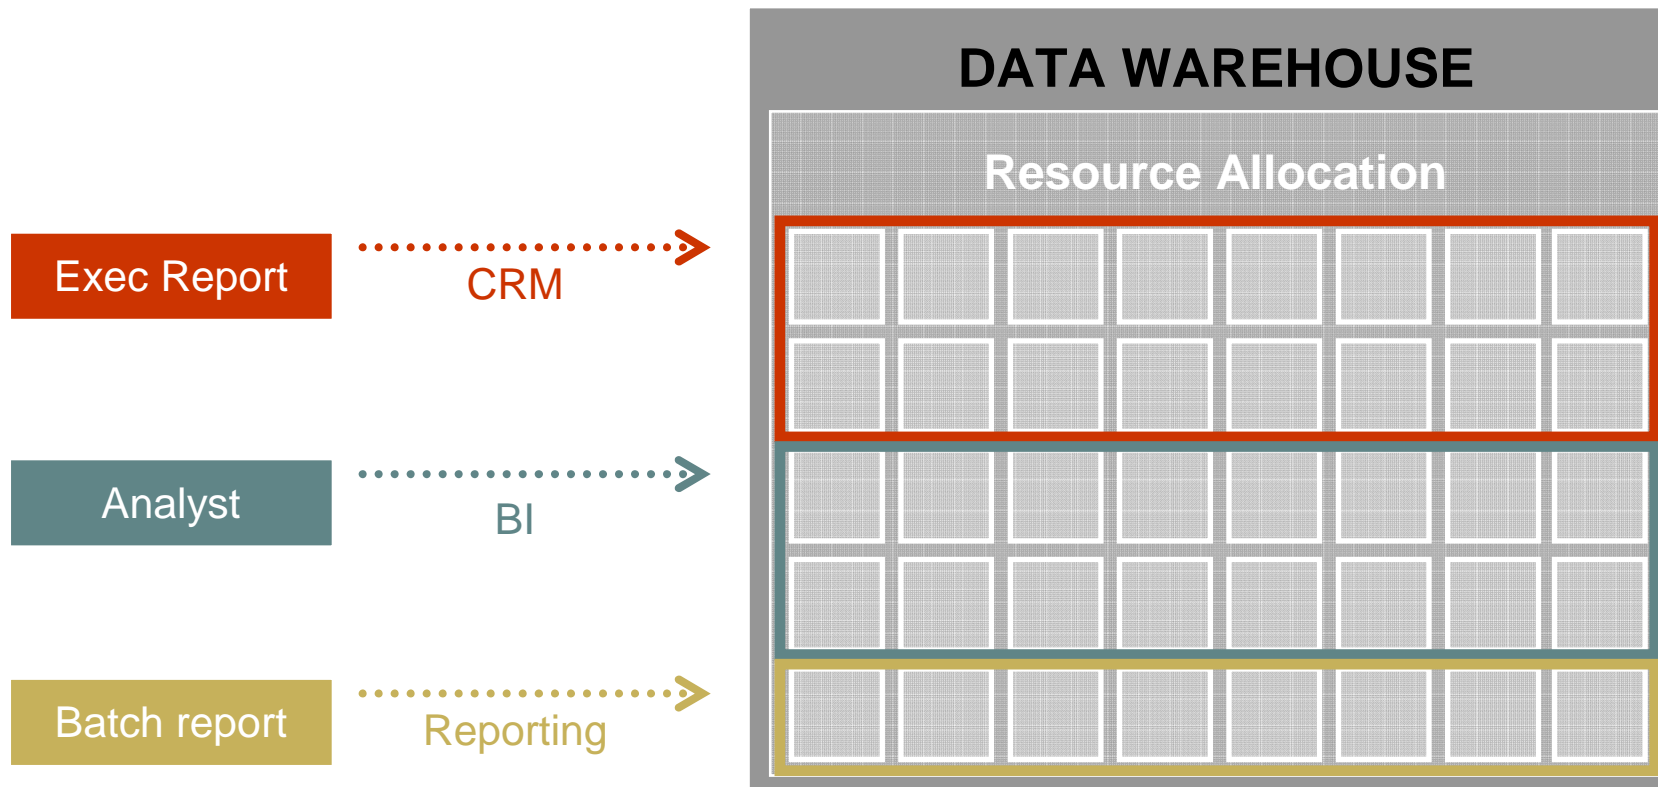
# Why does it matter? Consider a typical warehouse...

Unacceptable  
Response Time!


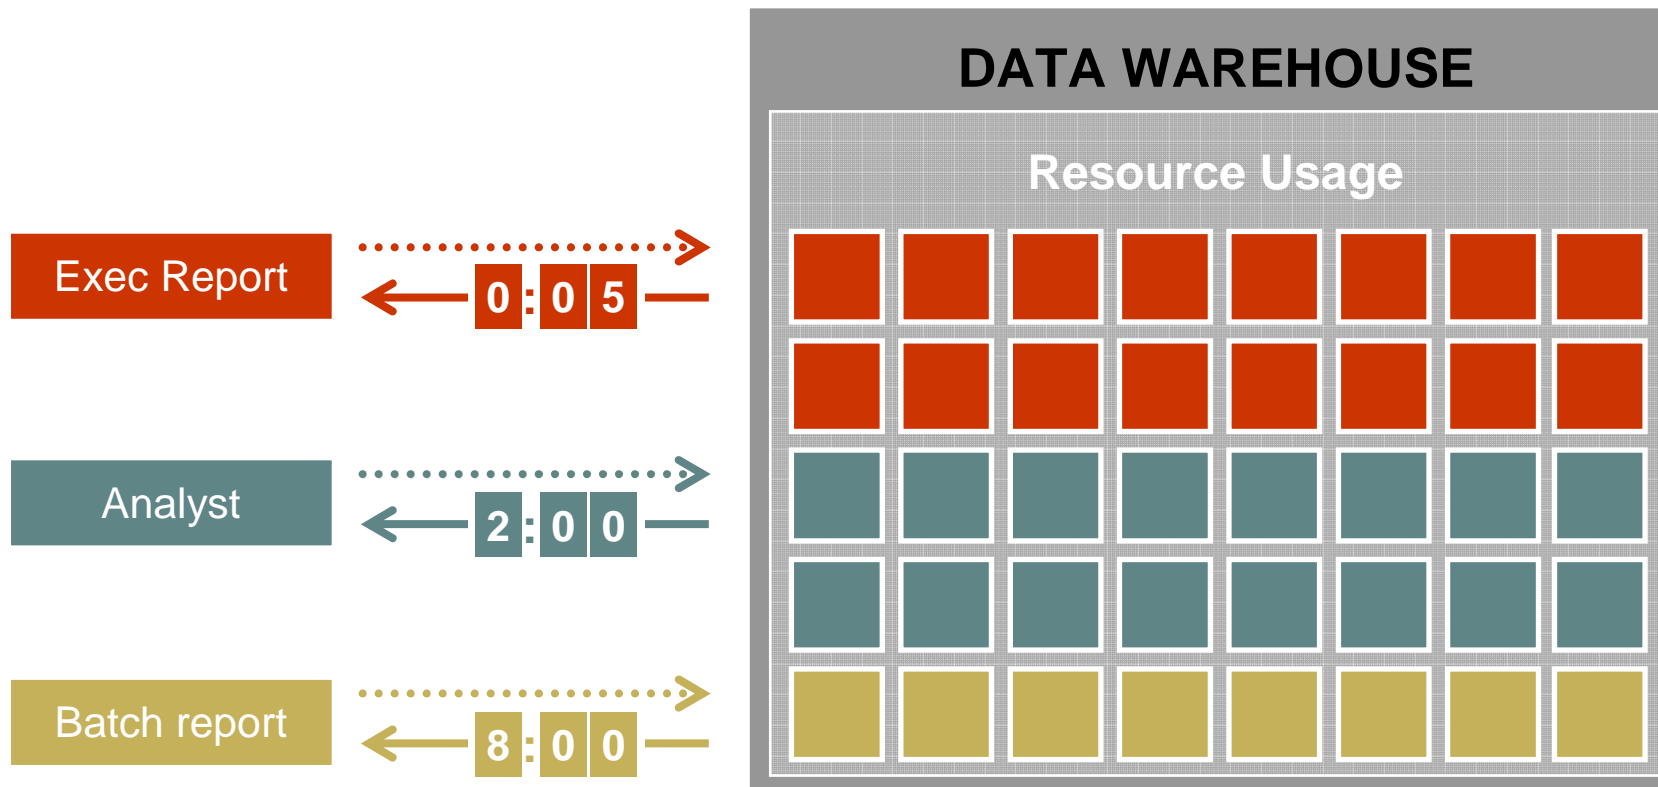




# Extreme workload management warehouse lets you align resources according to business priorities ...




... to ensure operational applications get the response time required—regardless of the number and types of requests



# Unexpected Workload is put on Queue

ID	Query	started	released	completed	wait	exec	total
5238	Q3	17:03:34	17:03:34	17:11:09	00:00	07:35	07:35
5239	Q1	17:03:34	17:03:34	17:06:35	00:00	03:01	03:01
5240	Q1	17:03:34	17:03:34	17:06:43	00:00	03:09	03:09
5241	Q3	17:03:35	17:03:35	17:11:09	00:00	07:34	07:34
5242	Q3	17:03:35	17:03:35	17:11:09	00:00	07:34	07:34
5243	Q1	17:03:35	17:03:35	17:06:21	00:00	02:46	02:46
5244	Q3	17:03:35	17:03:35	17:11:09	00:00	07:34	07:34
5245	Q1	17:03:35	17:03:35	17:06:36	00:00	03:01	03:01
5246	Q1	17:03:36	17:05:30	17:07:15	01:54	01:45	03:39
5247	Q3	17:03:36	17:06:14	17:13:28	02:38	07:14	09:52
5248	Q3	17:03:36	17:06:15	17:13:28	02:39	07:13	09:52
5249	Q3	17:03:36	17:06:21	17:13:13	02:45	06:52	09:37
5250	Q1	17:03:36	17:06:26	17:08:15	02:50	01:49	04:39
5251	Q1	17:03:36	17:06:33	17:08:11	02:57	01:38	04:35
5252	Q3	17:03:36	17:06:33	17:13:26	02:57	06:53	09:50
5253	Q3	17:03:36	17:06:35	17:13:27	02:59	06:52	09:51
5254	Q1	17:03:36	17:06:36	17:08:29	03:00	01:53	04:53
5255	Q3	17:03:36	17:06:43	17:13:22	03:07	06:39	09:46



## Research Showed Delayed Extra Jobs Improves Response Time and Performance (IEEE Concurrency Journal @ 1999)

Improved Strategies for Dynamic Load Balancing - Microsoft Internet Explorer

File Edit View Favorites Tools Help

Back Forward Stop Home Search Favorites RSS Feeds

Address <http://www2.computer.org/portal/web/csd/doi/10.1109/4434.788780> Go Links

### Improved Strategies for Dynamic Load Balancing

IEEE Concurrency  
RSS feed for this publication

DOI Bookmark: <http://doi.ieeeecomputersociety.org/10.1109/4434.788780> July-September 1999 (vol. 7 no. 3) pp. 58-67

**Chi-Chung Hui**  
**Samuel T. Chanson**

#### ABSTRACT

This article investigates load balancing in highly dynamic local area networks connecting computers of different speeds. Most dynamic load-balancing techniques ignore the lag time in load-information update, and dispatch jobs immediately upon arrival irrespective of the overall LAN loading. This can lead to system saturation and thrashing. Moreover, these techniques focus on steady-state system throughput without considering the behavior in transient periods. So, they might not perform satisfactorily in situations that often occur in practice, where jobs arrive in clusters and system loading fluctuates widely. To tackle these problems, the authors propose two strategies: One takes network delay into account to avoid errors in scheduling jobs; the other dynamically delays job execution when the system is fully used. These strategies are general and can augment existing algorithms. Experimental results show that these strategies adapt well to load fluctuation, minimize system loading while producing shorter schedules, and improve job fairness compared to some popular schedulers.

#### ADDITIONAL INFORMATION

**Index Terms:**  
job scheduling, load balancing, local area network, task allocation, workstation cluster

**Citation:**  
Chi-Chung Hui, Samuel T. Chanson, "Improved Strategies for Dynamic Load Balancing," *IEEE Concurrency*, vol. 7, no. 3, pp. 58-67, July-Sept. 1999, doi:10.1109/4434.788780

#### This Article

- PURCHASE ARTICLE: \$0
- PDF
- HTML
- IEEE Xplore Subscribers

#### Share

- Email this Article to a friend

#### Bibliographic References

- ASCII Text
- BibTex
- RefWorks Procite/RefMan/Endnote

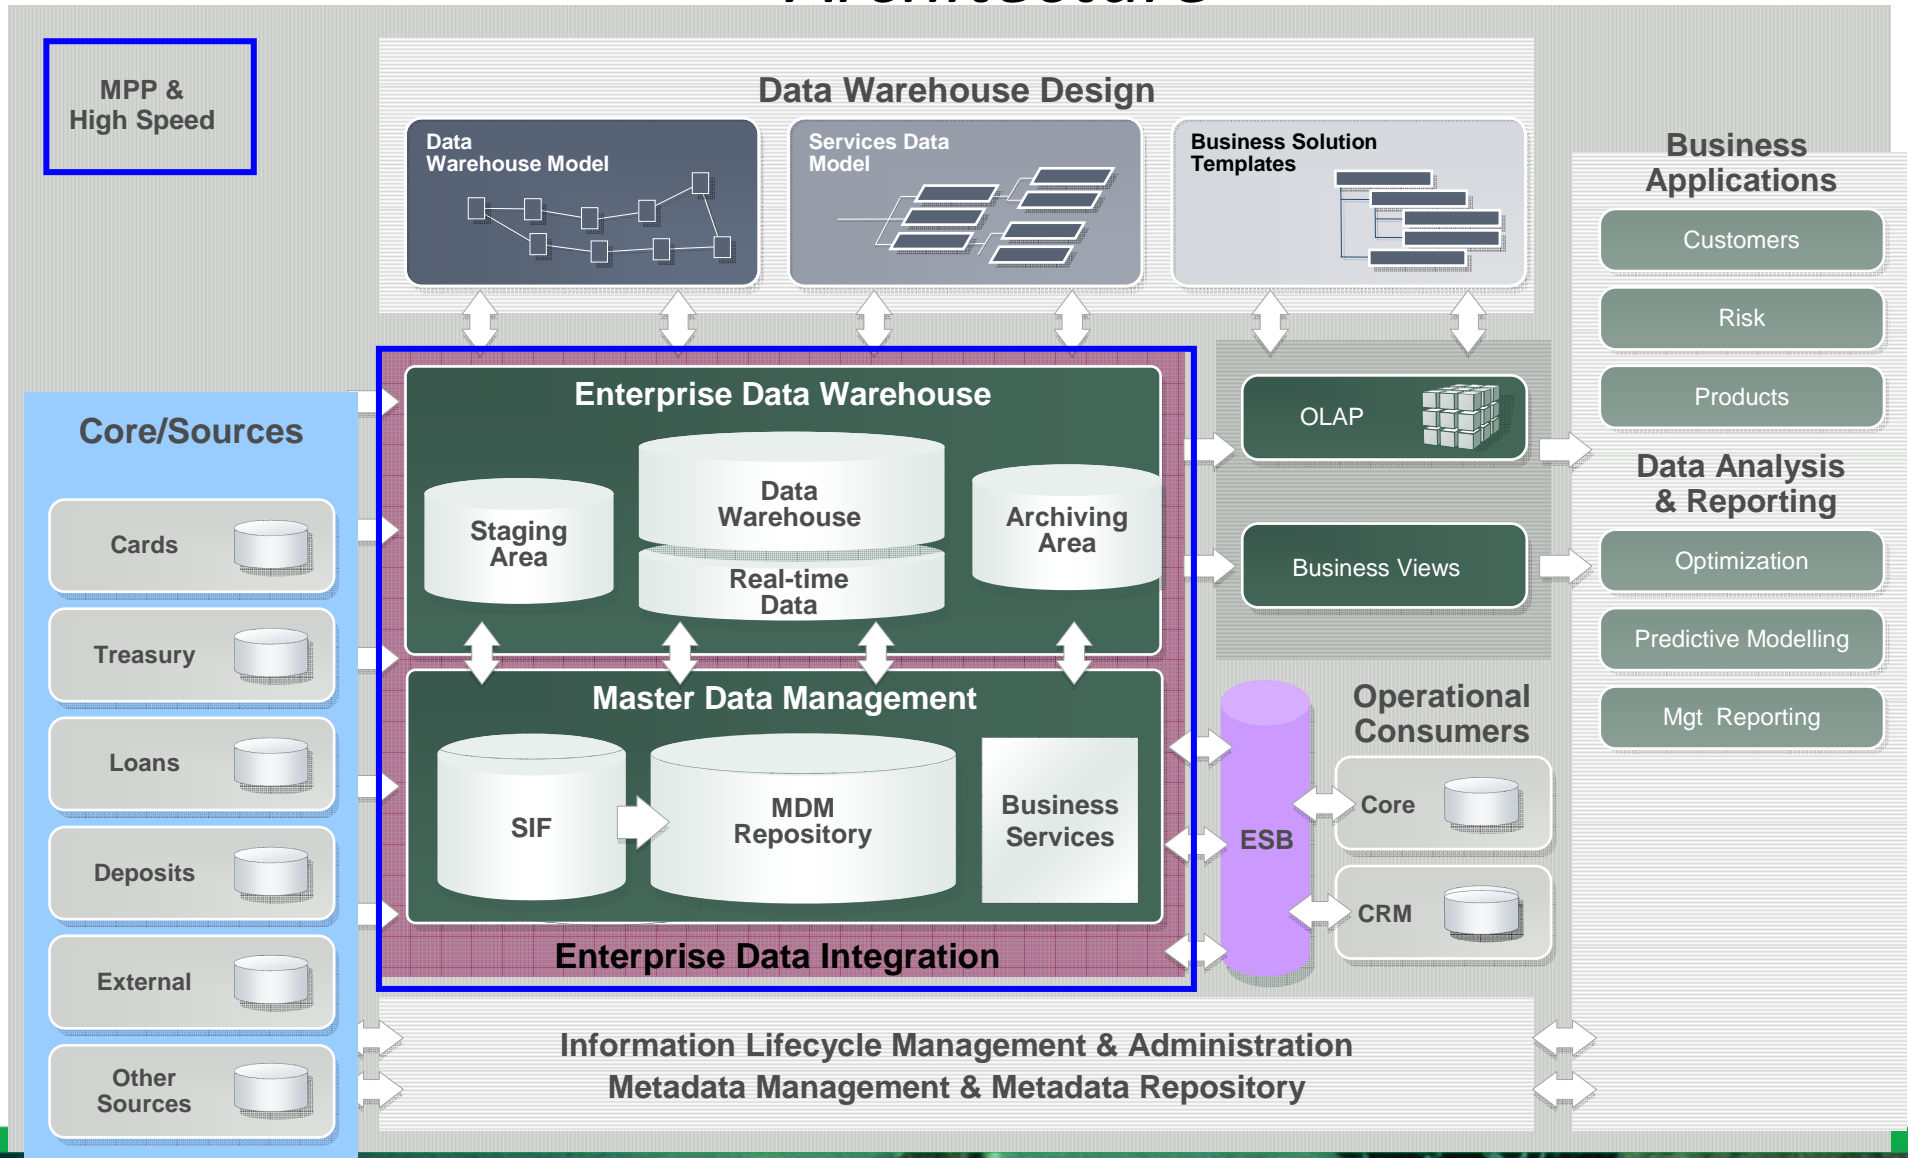
#### Add to:

- Digg
- Spurl
- Simpy
- Del.icio.us
- Furl
- Blink
- Google
- Y!MyWeb

#### Search

- Similar Articles

# Summary: Smart Enterprise Information Management Architecture







Information Management



# IBM Information On Demand Conference 2009

**Thank You!**