

April 18th - 20th, 2011



IBM Information Management Technology Showcase:

A New Era for Business Analytics and Data Warehousing

Mark Thomas

Warehouse Sales Lead, ANZ

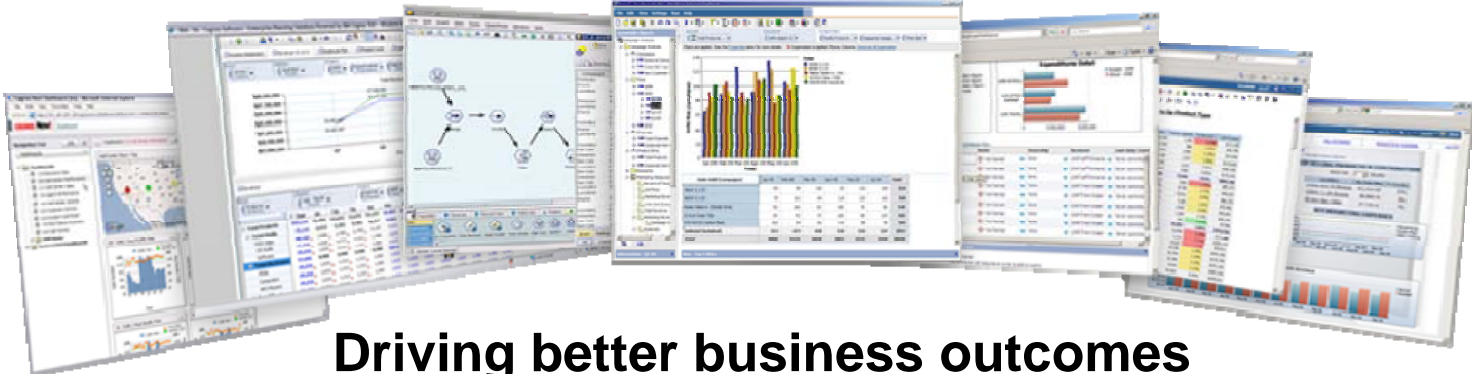


V7.2

Agenda

- Introduction: Challenges that organizations are facing
- Business Analytics - driving better business outcomes
- Data Warehousing – the foundation for business analytics
- Accelerating Business Insight & Analytics
- Demonstration
- Questions

IBM Business Analytics



Driving better business outcomes

IBM Data Warehouse Solutions



The optimized foundation for business analytics

All Organizations are Facing an Information Challenge

Trust 1 in 3

Business leaders frequently make decisions based on information they don't trust, or don't have.

Access 1 in 2

Business leaders say they don't have access to the information they need to do their jobs.

Vision 83%

Of CIOs cited "Business intelligence and analytics" as part of their visionary plans to enhance competitiveness.


Pre-integrated 35%

Of businesses will look to replace their current warehouse with a pre-integrated warehouse solution in the next 3 years.

Data Warehousing Requirements are Changing



For a More Instrumented, Interconnected, Intelligent World



10x growth in data
Digital data is projected to grow tenfold from 2007 to 2011.

80% data unstructured
from all these devices that requiring effort to understand and analyze.

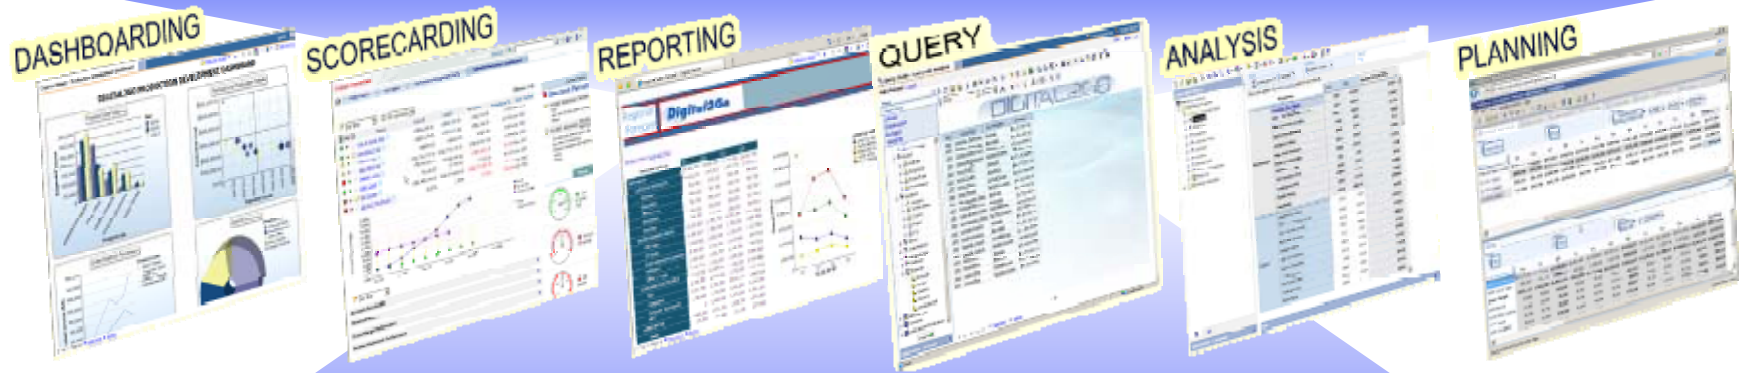
1 trillion devices
connected to the Internet by 2011.

6 terabytes of information
is exchanged over the internet every second.

... wipe out the competition!

70% of companies in the global 1,000 will have to change their data centers to meet requirements.


Information Demands Across the Organization are Increasing



- What is happening?**
- Why are we on/off track ?**
- What is likely to happen?**
- What should we do next ?**

Analytics-Driven Organizations Outperform

Driving Better Financial Performance



Only Recognized Best Of Breed Portfolio Across Entire Information Management Warehouse Value Chain



Simplicity, Flexibility, Choice

IBM Data Warehouse & Analytics Solutions

**IBM
Netezza**



**IBM
Smart Analytics System**



**IBM
Warehouse Software**



**Custom
Solutions**

Warehouse Accelerators

Information Management Portfolio
(Information Server, MDM, Streams, etc)

Simplicity

The right mix of simplicity and flexibility

Flexibility

Are you overpaying for Oracle Database? Hint: you're overpaying for Oracle Database.

The first thing to consider when thinking about DB2® for your business: It's as low as 1/3 the cost of Oracle Database. Then consider DB2 on Power Systems™ with 3x the performance per core of Oracle Database on SPARC. In TPC-C and SAP SD benchmarks. Overall, an ironclad case for IBM. There's more where that came from, too.

ibm.com/facts

COST based on publicly avail U.S. info on 2/10/2011 for IBM DB2 Advanced Enterprise Edition + Oracle software w/comparable capabilities. IBM: 100 Processor Value Units. Oracle: assumes 1.0 processor multiplier. Both incl. Y1 maint/support. PERFORMANCE: www.tpc.org (<http://www.tpc.org>) as of 01/26/11 [IBM Power 780 (3 x 64 C)(24 Ch/192 C/768 Th); 10,366,254 tpmC; \$1.38/tpmC; avail 10/13/10 v. Oracle SPARC SuperCluster w/T3-4 Servers (27 x 64 C)(108 Ch/1728 C/13824 Th); 30,249,688 tpmC; \$1.01/tpmC; avail 6/1/11]. TPC-C is a trademark of Transaction Performance Processing Council. www.sap.com/solutions/benchmark/ (<http://www.sap.com/solutions/benchmark/>) as of 01/26/11 [IBM Power 795 (32 P/256 C/1024 Th); 126063 users/2-tier SAP ERP 6.0 pack4/AIX 7.1 + DB2 9.7; cert 2010046 v. Oracle SPARC Enterprise Server M9000 (64 P/256 C/512 Th); 39100 users/2-tier SAP ERP 6.0/Solaris 10, Oracle 10g; cert 2008042]. SAP is registered trademark of SAP AG in Germany and in several other countries.


Netezza. Up and running in 24 hours, not 24 days.

Get set up in hours instead of days and start counting returns in minutes instead of hours. All with IBM's Netezza data warehouse appliance for high performance analytics. It gives you analytics reports at supersonic speeds. At a fraction of the cost of Oracle Exadata. Get real, actionable business results fast.

ibm.com/facts

COST comparison based on publicly available information as of 2/10/2011 for an Oracle Exadata X2-2 HP Full Rack and a full rack of Netezza TwinFin. The cost to acquire Netezza can be as low as 1/6 of Exadata if a client is acquiring new Oracle database licenses and as low as 1/2 if using existing Oracle database licenses.

IBM offers the widest and deepest portfolio of data warehouse solutions



For times when ultimate flexibility is required:

IBM offers warehousing and analytics software individually for build-your-own solutions



- Choice of platform
- Choice of operating system
- Dynamic growth
- Complex workloads
- Variety of applications
- Mixed and hybrid usage
- Variety of IT and DBA skills



*About 2,500 users and 200,000 reports per month:
We would not have been able to achieve our ambitious goals
in business intelligence without InfoSphere Warehouse*

- Ralf Bruhnke, Controlling and Project Manager for Karstadt



Smart Analytics System

The modular system for business analytics

- Integrated Cognos Business Intelligence
 - Integrated InfoSphere Warehouse
 - In-database cubing and mining
 - Choice of platform and OS
-
- **Scale 'On Demand'**
 - **Modular application interfaces**
 - **Built for complex and mixed workloads**
 - **Autonomic tuning**



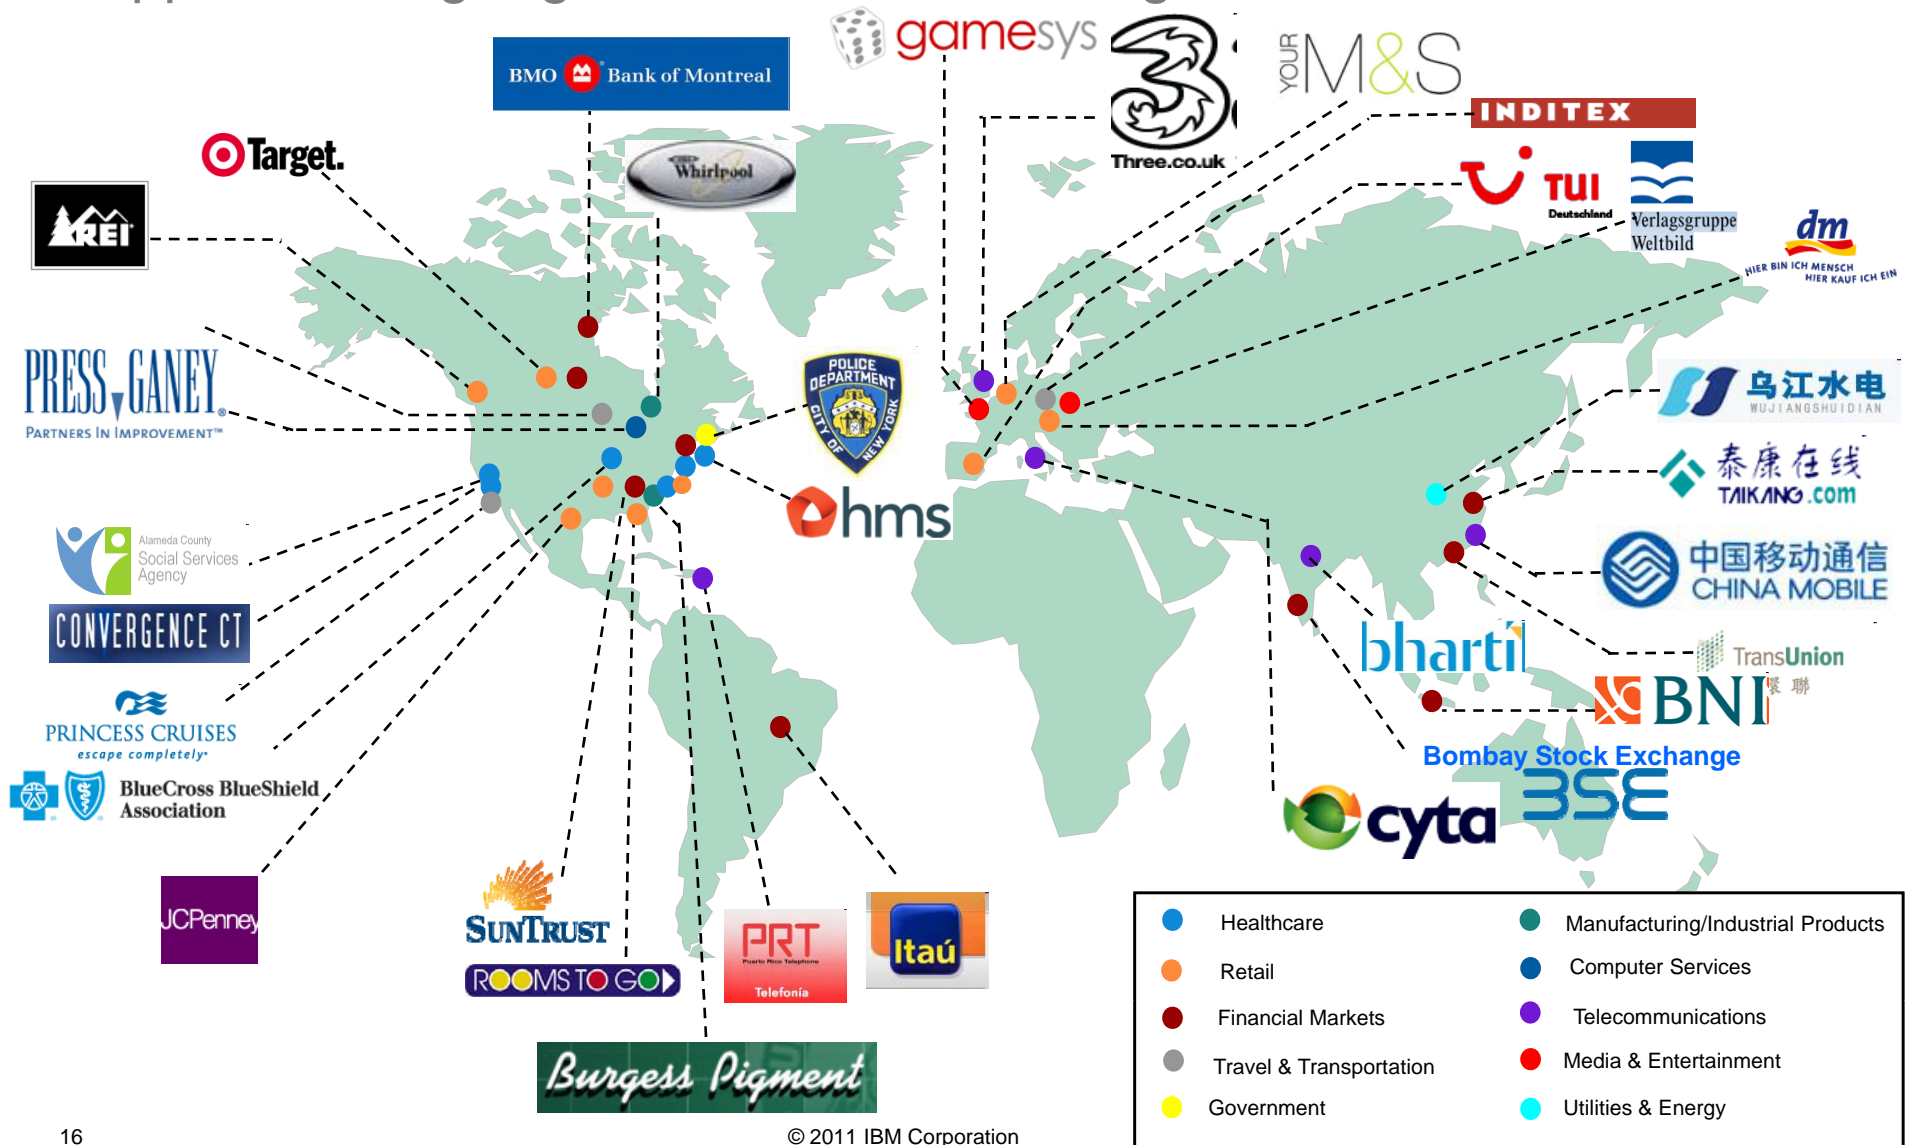
IBM Netezza

The true data warehousing appliance

- Purpose-built analytics engine
 - Integrated database, server and storage
 - Standard interfaces
 - Low total cost of ownership
- **Speed: 10-100x faster than traditional system**
 - **Simplicity: Minimal administration and tuning**
 - **Scalability: Peta-scale user data capacity**
 - **Smart: High-performance advanced analytics**

And the Smart Analytics System & InfoSphere Warehouse

Support leading organizations across the globe



The IBM Netezza appliance empowers all types of industries:

Digital Media



Financial Services



Government



Health & Life Sciences



Retail / Consumer Products



Telecom



Other



In Summary, IBM Continues To Invest: Your Partner in Business Analytics and Data Warehousing Solutions

- More than **\$14B** in Acquisitions Since 2005
- More than **10,000** Technical Professionals
- More than **7,500** Dedicated Consultants
- **Largest** Math Department in Private Industry
- More than **27,000** Business Partner Certifications



2005

2011

Thank You

Mark Thomas
mathomas@au1.ibm.com

Name

Date



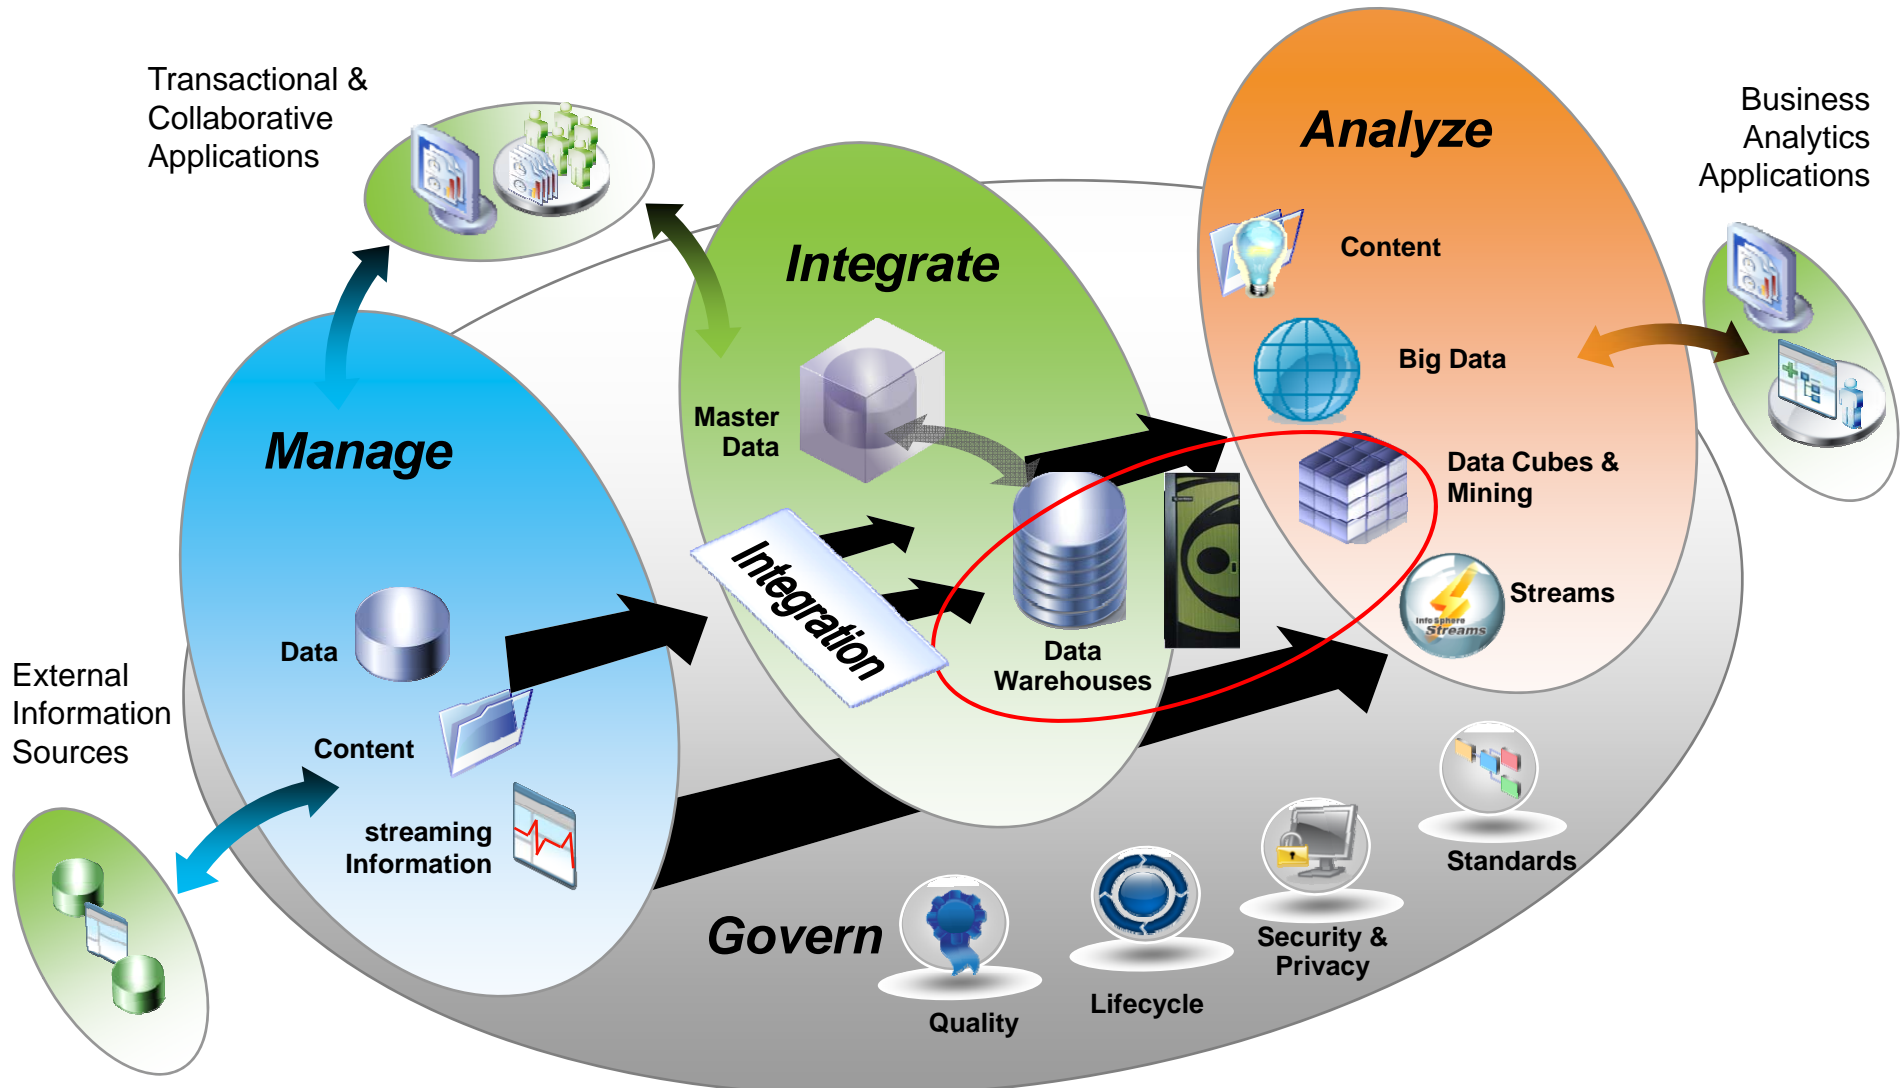
InfoSphere Warehouse – Accelerating Business Insight



V7.2

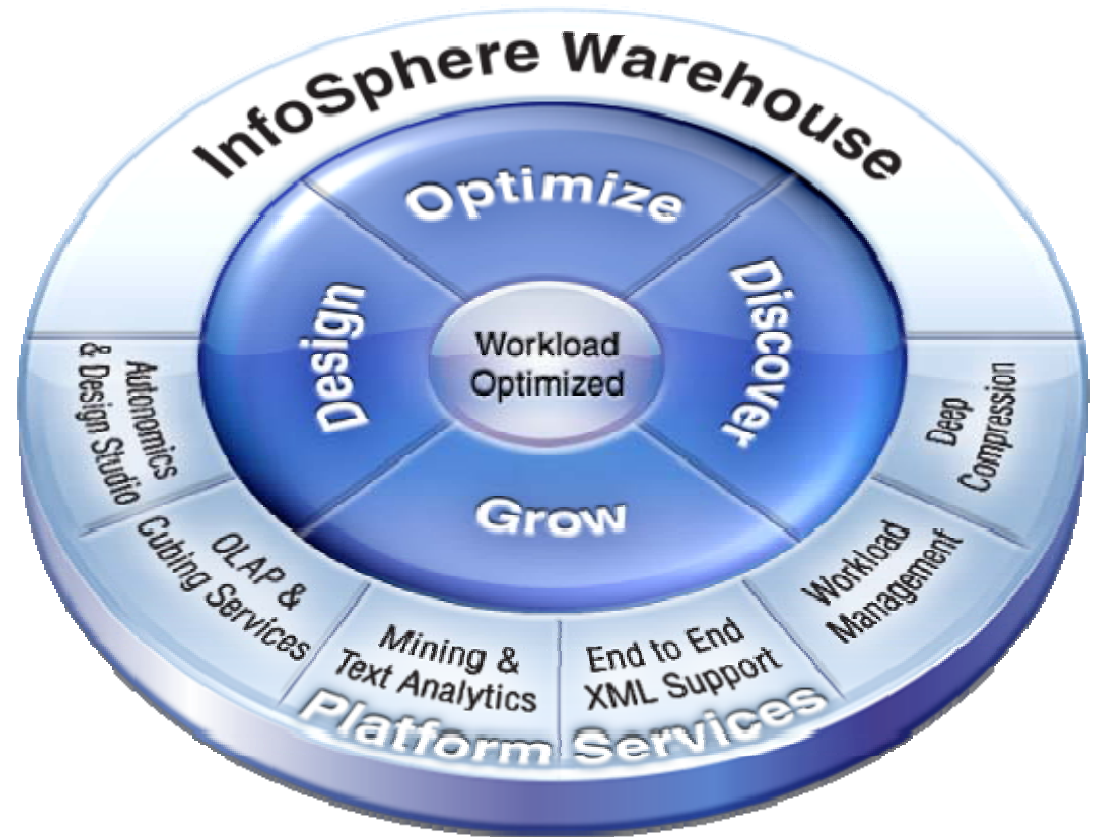
Where Does the Data Warehouse Fit in the IT Environment

IBM Offers a complete information ecosystem



IBM InfoSphere Warehouse

- Provides a comprehensive data warehouse platform
- Built on the powerful DB2 with features for warehousing
- Built-in in-the-database analytics and realtime data mining.



InfoSphere Warehouse

Performance



- Massively Parallel Processing*
- Compression*
- High-Speed Query Performance*
- Mixed Workload Mgmt*

Sophisticated Analytics



- OLAP*
- In-Database Mining*
- Embedded Business Intelligence*
- Mashup Platform*

Flexibility



- | | | |
|--------------------------|-----------------------------|------------------------|
| <i>Autonomics</i> | <i>Native XML Support</i> | <i>Tools</i> |
| <i>High Availability</i> | <i>External Data Access</i> | <i>Business Models</i> |
| <i>Security</i> | | |

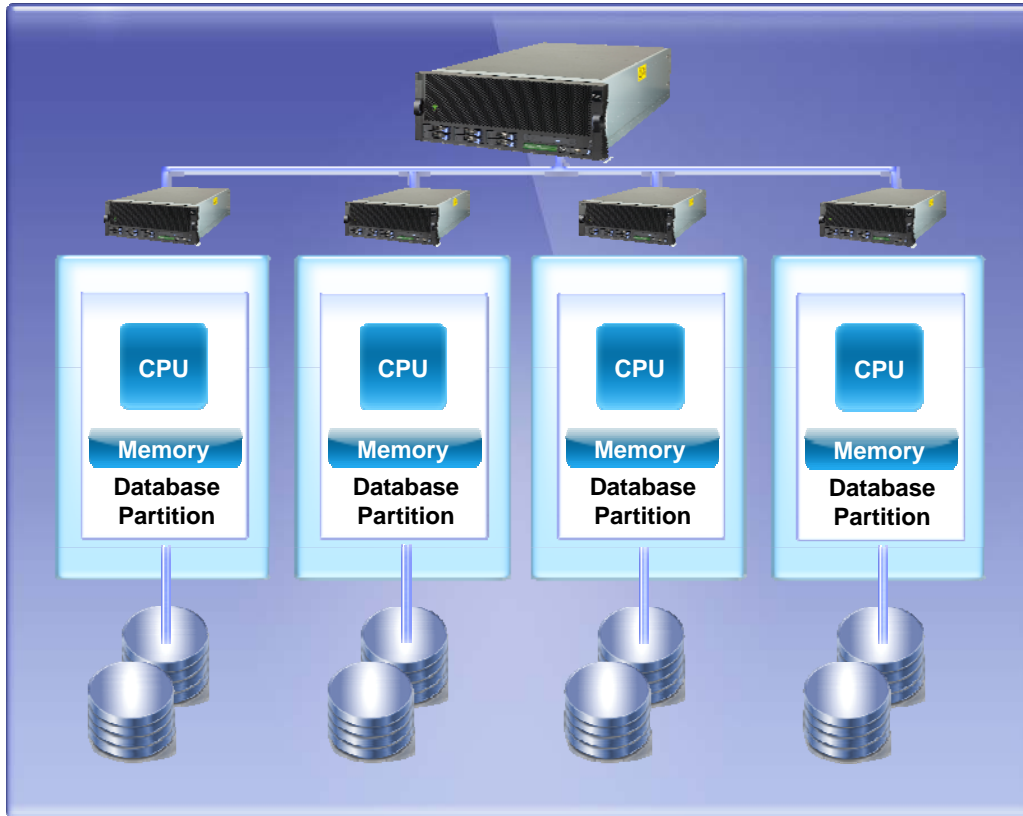
IBM InfoSphere Warehouse – Built on DB2

- Shared Nothing Architecture
- Advanced Cost Based Parallel Query Optimizer
- Flexible Partitioning
- Multi-dimensional Clustering (MDC)
- Materialised Query Tables (MQT)
- Powerful compression
- Workload management



Massively Parallel Processing

The foundation for high performance warehousing




- Database is divided into multiple partitions
- Partitions run on different nodes
- Each partition has dedicated resources (CPU, I/O, and memory.)
- Parallel processing occurs on all partitions and is coordinated by the DBMS
- Single database transparent to user

A warehouse, that works like a supercomputer

Deep Compression

Save time and money



Reduce Storage Requirements

- Database Compression: Reduce database storage requirements by up to 80%
- Backup Compression: Can shrink backup and disaster recovery storage in half.

Improved Performance

- Workloads run against compressed data can result in up to 40% performance improvements.

Save Time

- Save time with utilities processing (especially for database back-up and recovery)


Save Money

- Less storage, means less hardware, smaller server footprint, fewer administrative requirements and reduced power consumption.

Even more methods for reducing query response time

Materialised Query Tables (MQTs)

- Pre-computed query results used to improve the performance when the query is repeated by the same, or different users.



Cost Based Query Optimisation


- Ensures that the operations and methods used to fulfill a query minimize resource consumption and reduce response time.
- Key to scalable performance

Performance Management

- Includes IBM Optim Performance Manager which enables IT staff to identify, diagnose, solve and prevent performance problems in InfoSphere Warehouse.

InfoSphere Balanced Warehouse TPC-H Benchmark

■ Brute Force Parallelism



E7100 Balanced Warehouse delivers outstanding data warehouse performance

Load rate of 6TB / hour



TPC Benchmark, TPC-H, QphH, are trademarks of the Transaction Processing Performance Council. For further TPC-related information, please see <http://www.tpc.org>.

InfoSphere Warehouse 9.5 on IBM System p6 570, (128 core POWER6 4.7GHz), 343551 QphH@10000GB, 32.89 US \$ per QphH@10000GB available: April 15, 2008
Oracle 11g Enterprise Ed w/ Partitioning on HP Integrity Superdome-DC Itanium 2, HP-UX 11i v3 64 bit (128 core Intel Itanium 2 1.6 GHz), 208457 QphH@10000GB, 27.97 US \$ per QphH@10000GB available: September 10, 2008

Results as of Sept 7, 2009

Sophisticated Analytics

For Smarter Decision-Making



OLAP


In-Database Mining

Embedded Business Intelligence


Mashup Platform

Cubing Services

Enables OLAP applications to access large data volumes stored inside the warehouse



In-Database Mining Brings Analytics to the Data



Better Decisions . . .

- Real-time scoring
- Lower cost of deployment
- Reduce data movement & increase process efficiency
- Faster reaction time
- Leverage database security for data privacy

Without Leaving the Warehouse

InfoSphere Warehouse Intelligent Miner

Advanced Analytics with SPSS

In-Database Mining with SAS

InfoSphere Warehouse Intelligent Miner

Improve & Automate Decisions

DISCOVERY MINING
 Find patterns in data to improve decision-making



PREDICTIVE MINING
 Use known results to create predictive models



Flexibility

To Adapt to Your Business Needs



Autonomics	External Data Access
Security	Tools
High Availability	Business Models
Native XML Support	

Autonomics for Performance Tuning

Self-configuring, self-healing, self-optimising and self-protecting

- **Self Tuning Memory Manager:** Monitors memory usage and reallocates memory to optimise workload performance
- **Configuration Advisor:** Provides recommendations for database, and database manager configuration parameters.
- **Design Advisor:** Helps tune the physical design of tables and indexes.

Autonomics for Maintenance

Self-configuring, self-healing, self-optimizing and self-protecting

- **Automatic Storage:** Efficiently manages the use of space.
- **Health Monitor:** Tracks database and OS health indicators. Uses dashboard to present status, and provide recommendations for resolving alerts.
- **Monitored and automated:**
 - Statistics collection
 - Database and index re-organisations
 - Database backups



Security to protect your data assets

Auditing:

- Discover unwanted, unknown, and unacceptable access to the data, and keep history records of the activities in the warehouse.

Roles:

- Assign roles to a user or group.
- Role privileges are available for use in any and all activities undertaken by the connection user.
- Simplifies administration and management of privileges in the warehouse.

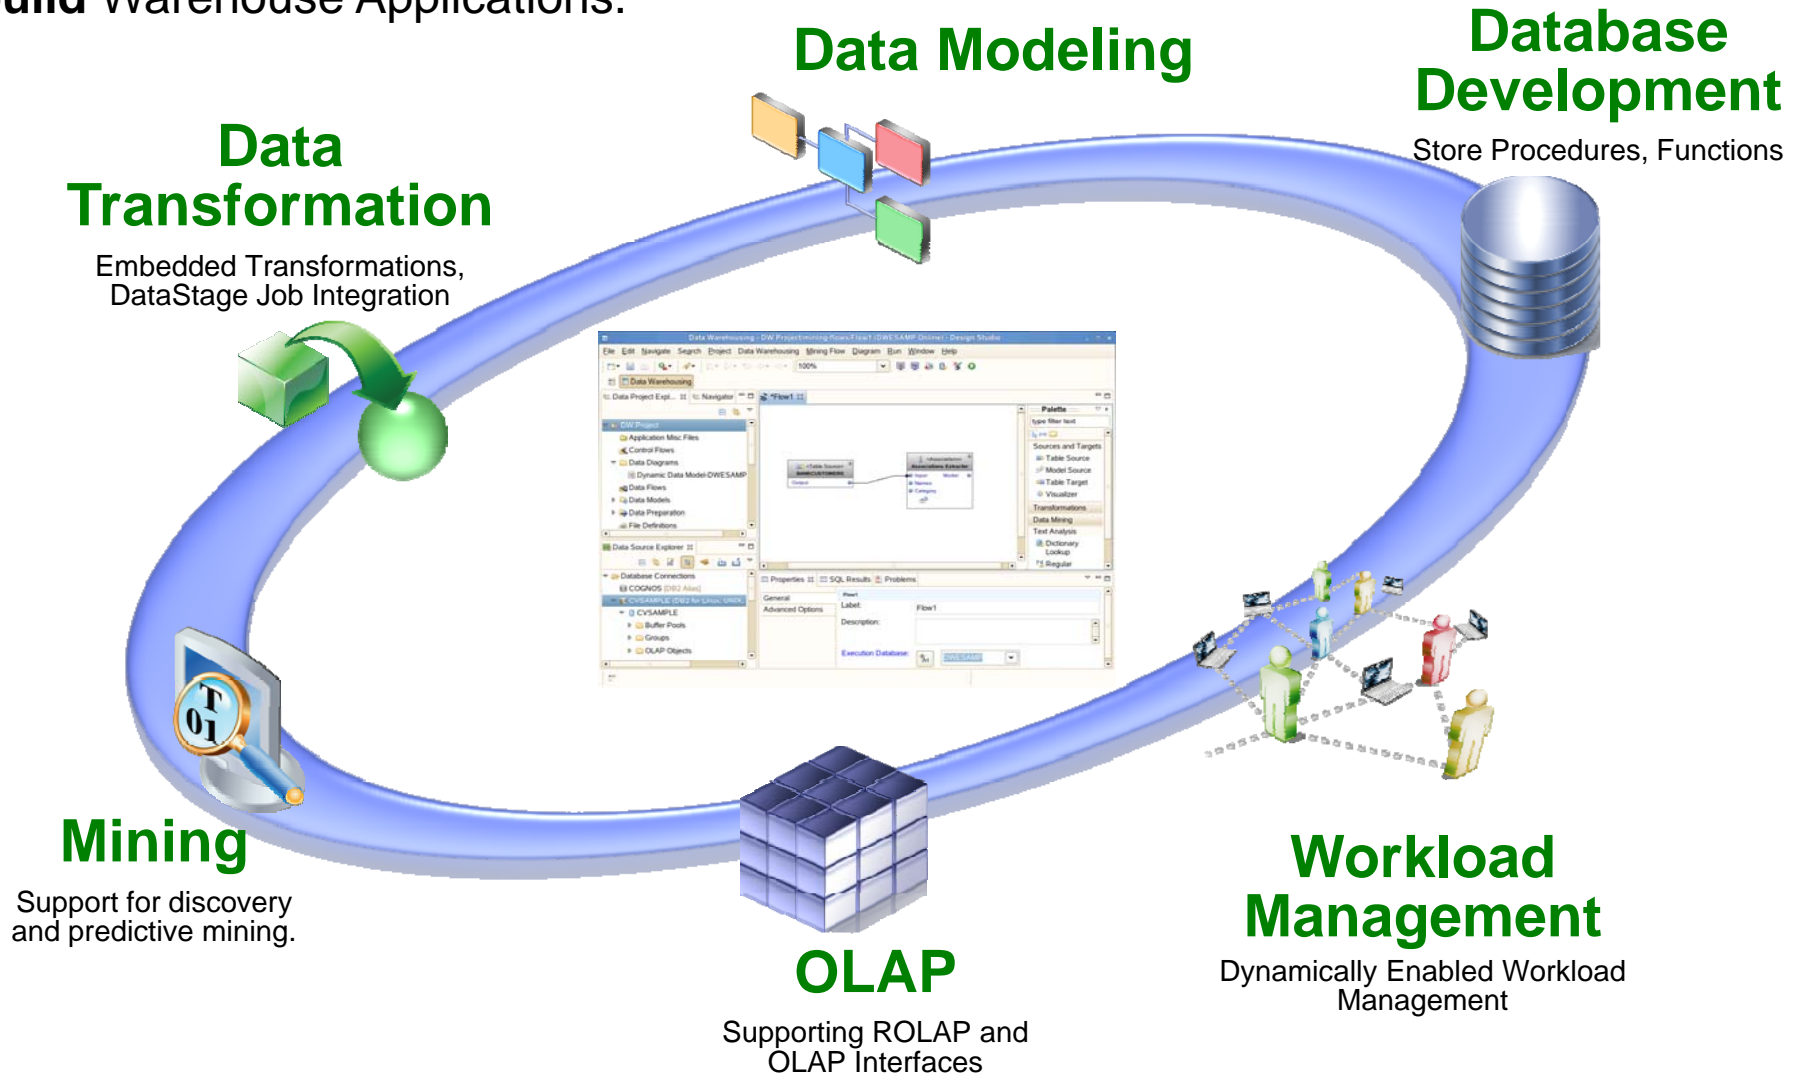


Label-Based Access Control (LBAC):

- Increases the control you have over who can access your data.
- Lets you decide exactly who has write access and who has read access to data at the individual row and column level.

Tools – Design Studio

Eclipse-based tool provides all the capabilities required to **design** and **build** Warehouse Applications.





Industry Models

- Proven approach and methodology
- Embedded industry expertise
- More than 500 global clients




Realising the Value

- Reduces analysis and design of requirements by more than 40%
- Accelerates stakeholder approval by over 50%
- Supports on average 85% of an organisation's data requirements
- Cuts time compared to custom-built data warehousing projects

Banking, Financial Markets, Health Plan, Insurance, Retail, and Telecommunications

InfoSphere Warehouse Model Packs


Pre-built warehouses for advanced business analytics



**Customer
Insight**



**Market &
Campaign Insight**



**Supply Chain
Insight**

Information Server for Data Warehouse Offering v8.5



What we have created:

- Information Server for Data Warehousing
- Provides all the key data integration and data cleansing capabilities
- Supports select IBM Smart Analytics System and IBM Netezza models

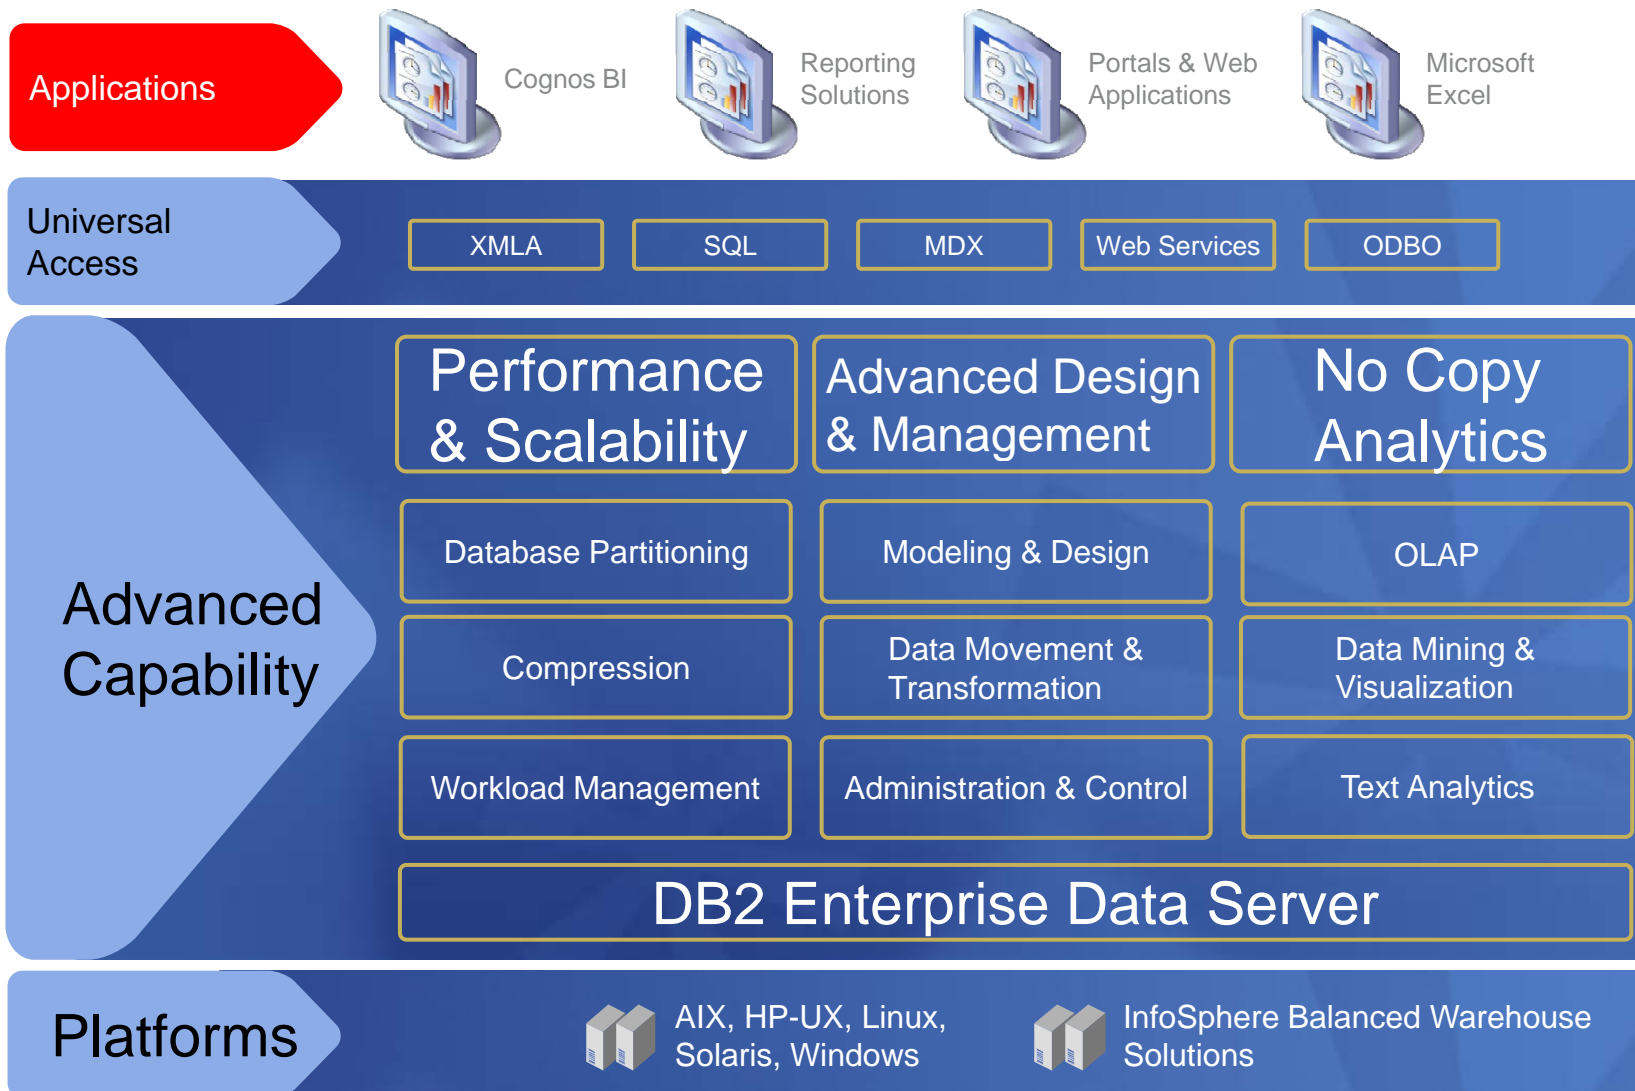


Why?

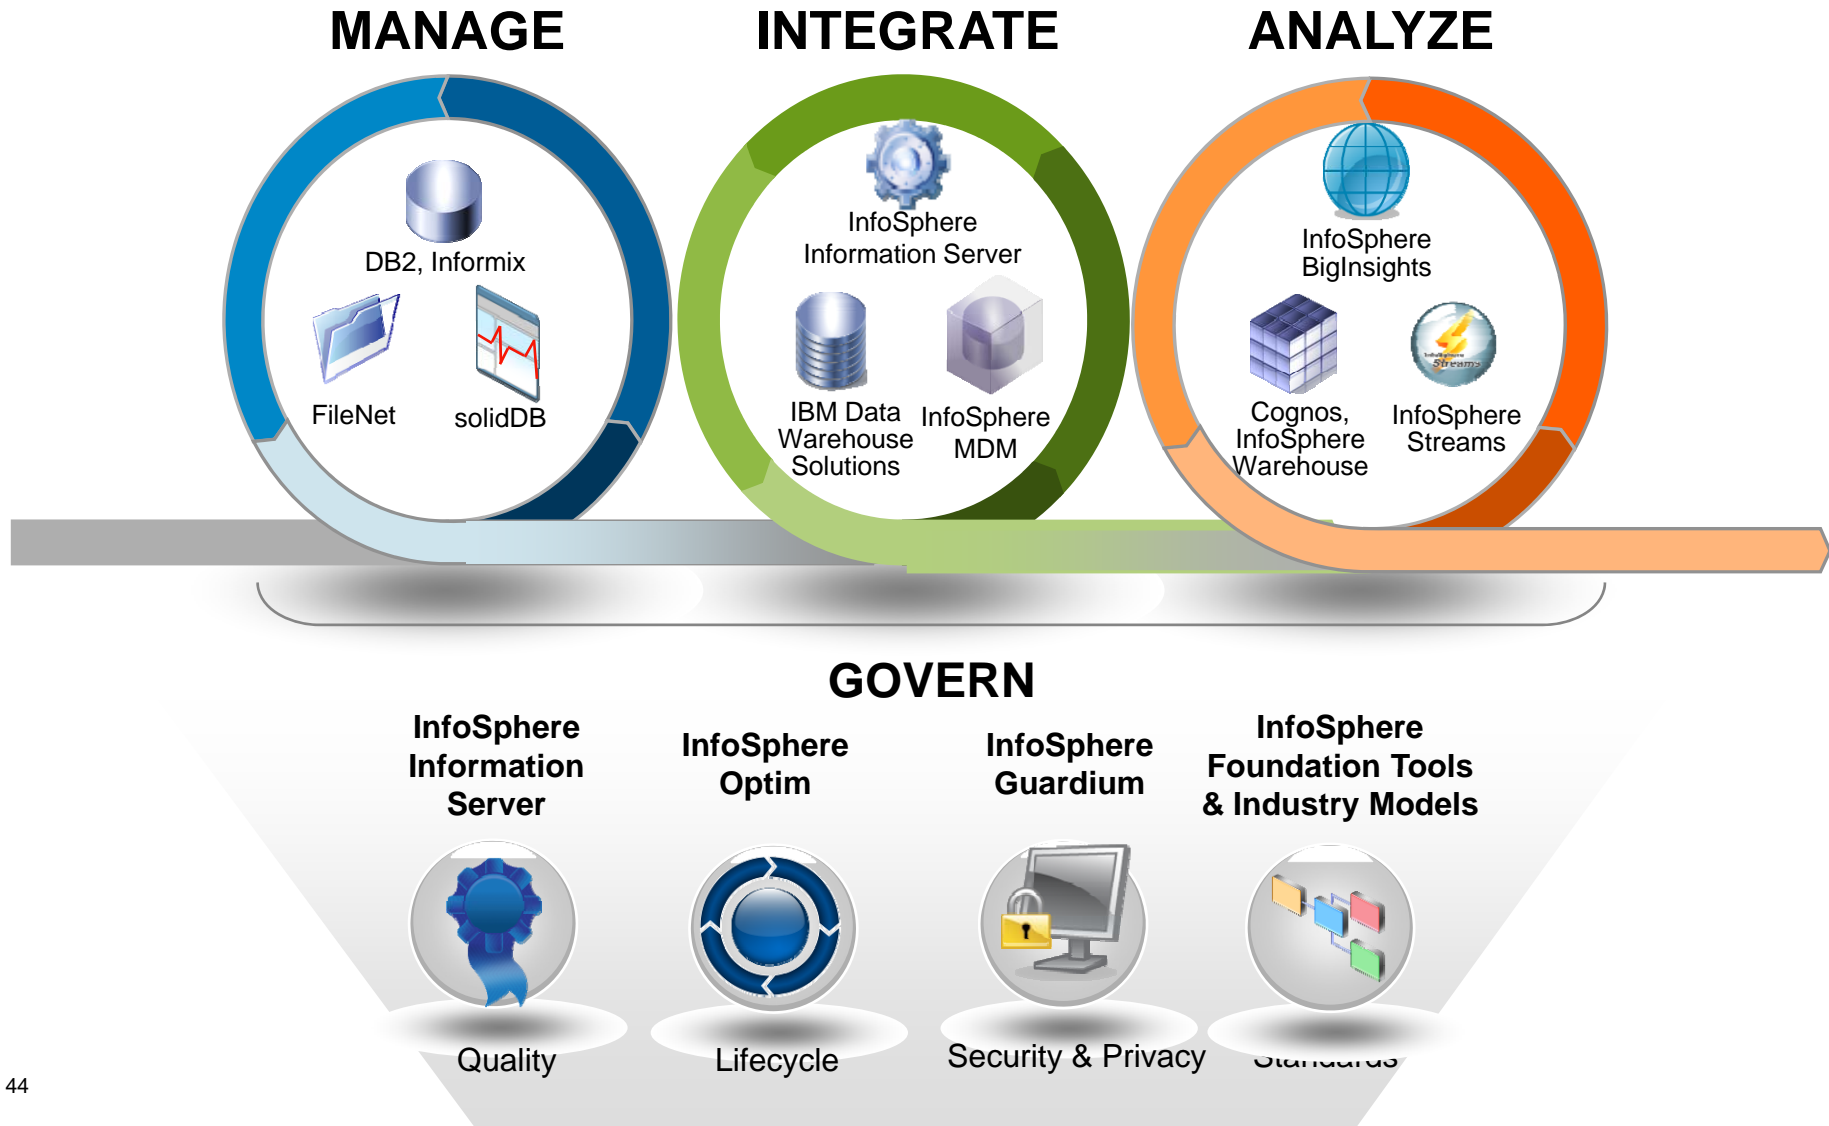
- Help accelerate your warehouse deployment by providing seamlessly integrated information integration capabilities as part of a common infrastructure



InfoSphere Warehouse is More than Just DB2



IBM provides integrated, modular and heterogeneous solutions for governance across your information supply chain



Summary

InfoSphere Warehouse provides a road to actionable insight.

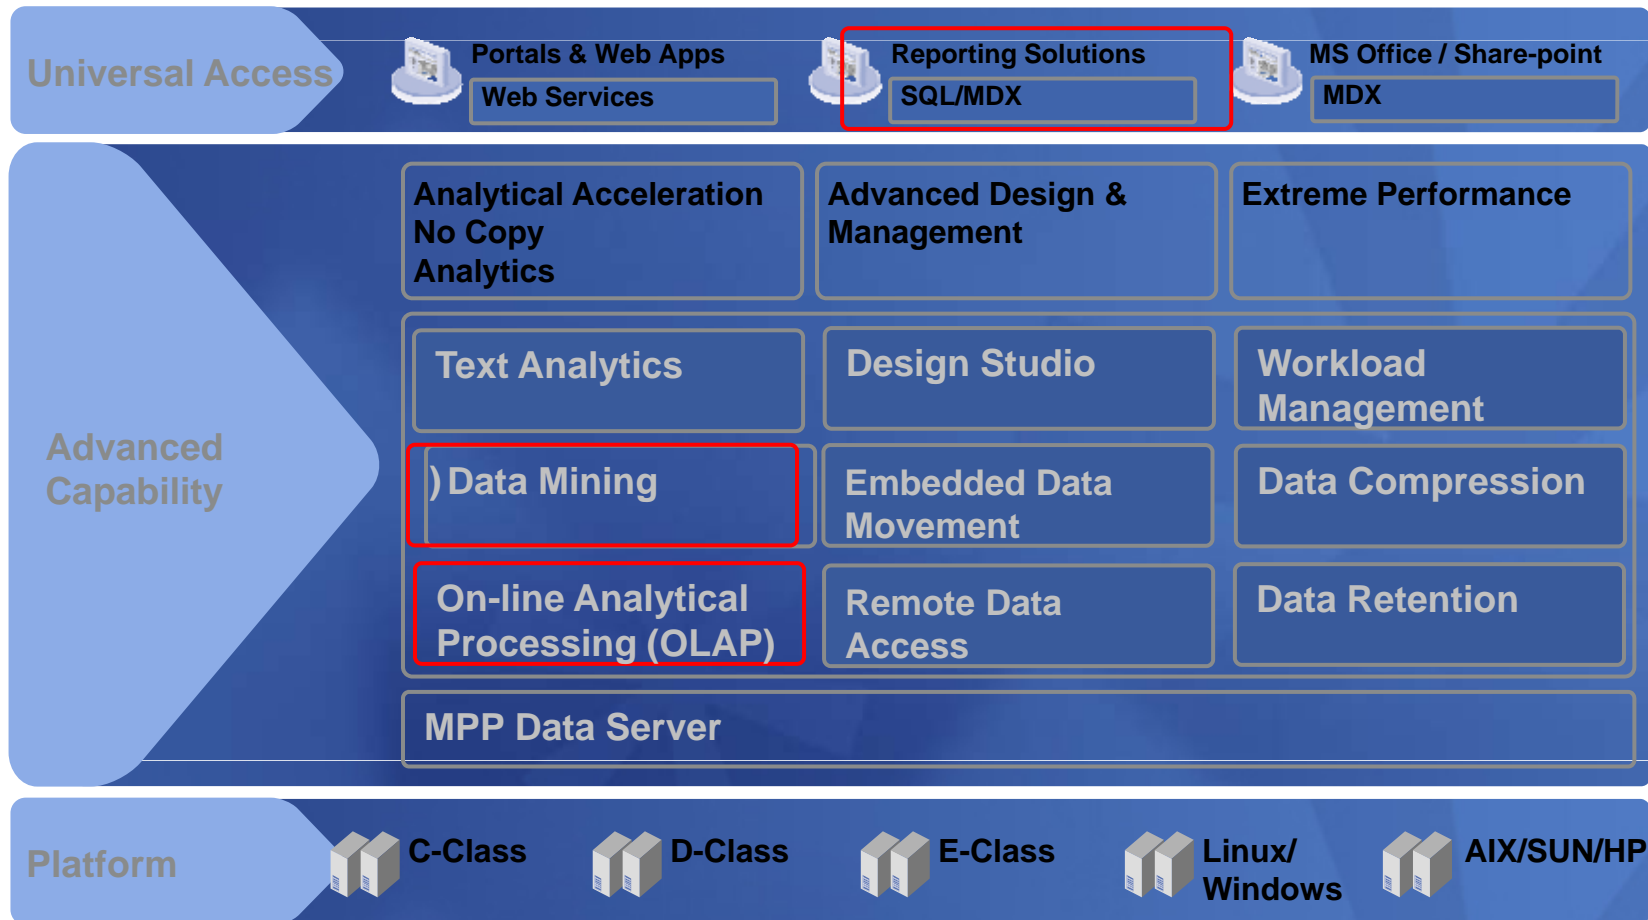
- **Performance and scalability features to support exponential data growth, and demanding users.**
- **Embedded analytics for smarter decision-making in real-time.**
- **Flexibility to address the needs of organisations of different sizes, usage requirements and environments.**

Thank You

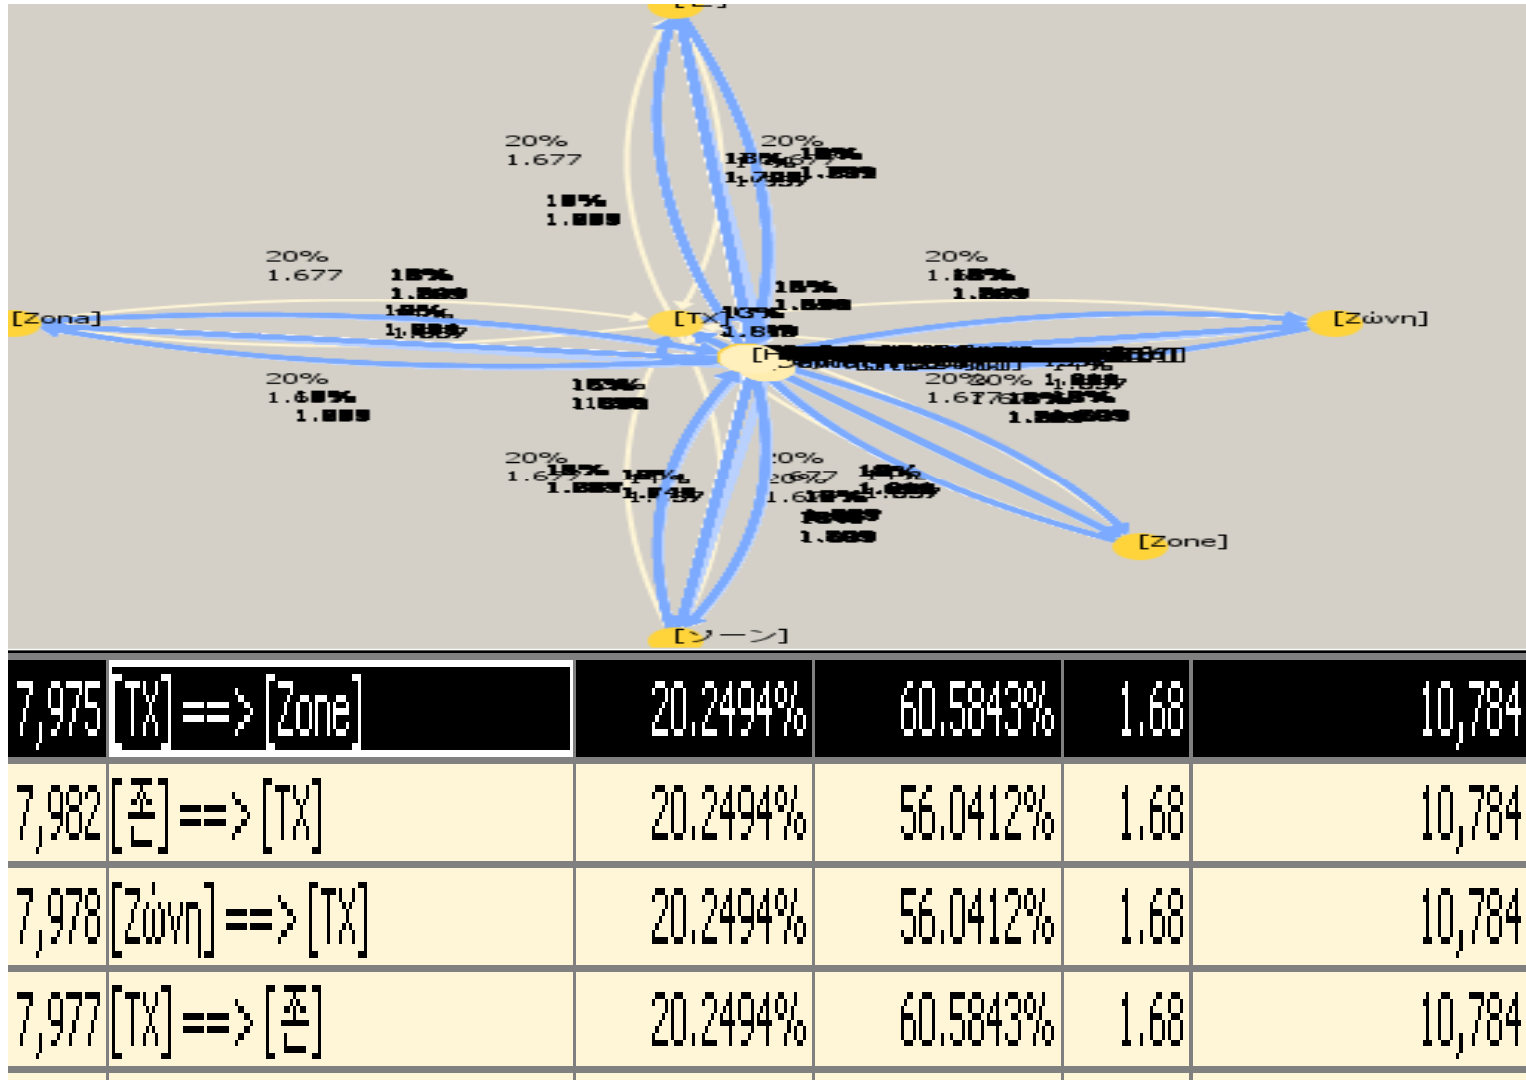
DEMONSTRATION

Ronnie Chan
Predictive Analytics Solution Architect
ronniech@au1.ibm.com


InfoSphere Warehouse : Supporting Advanced Analytics



Demo : Data Mining : Insight -> Action




Demo : CRISP-DM



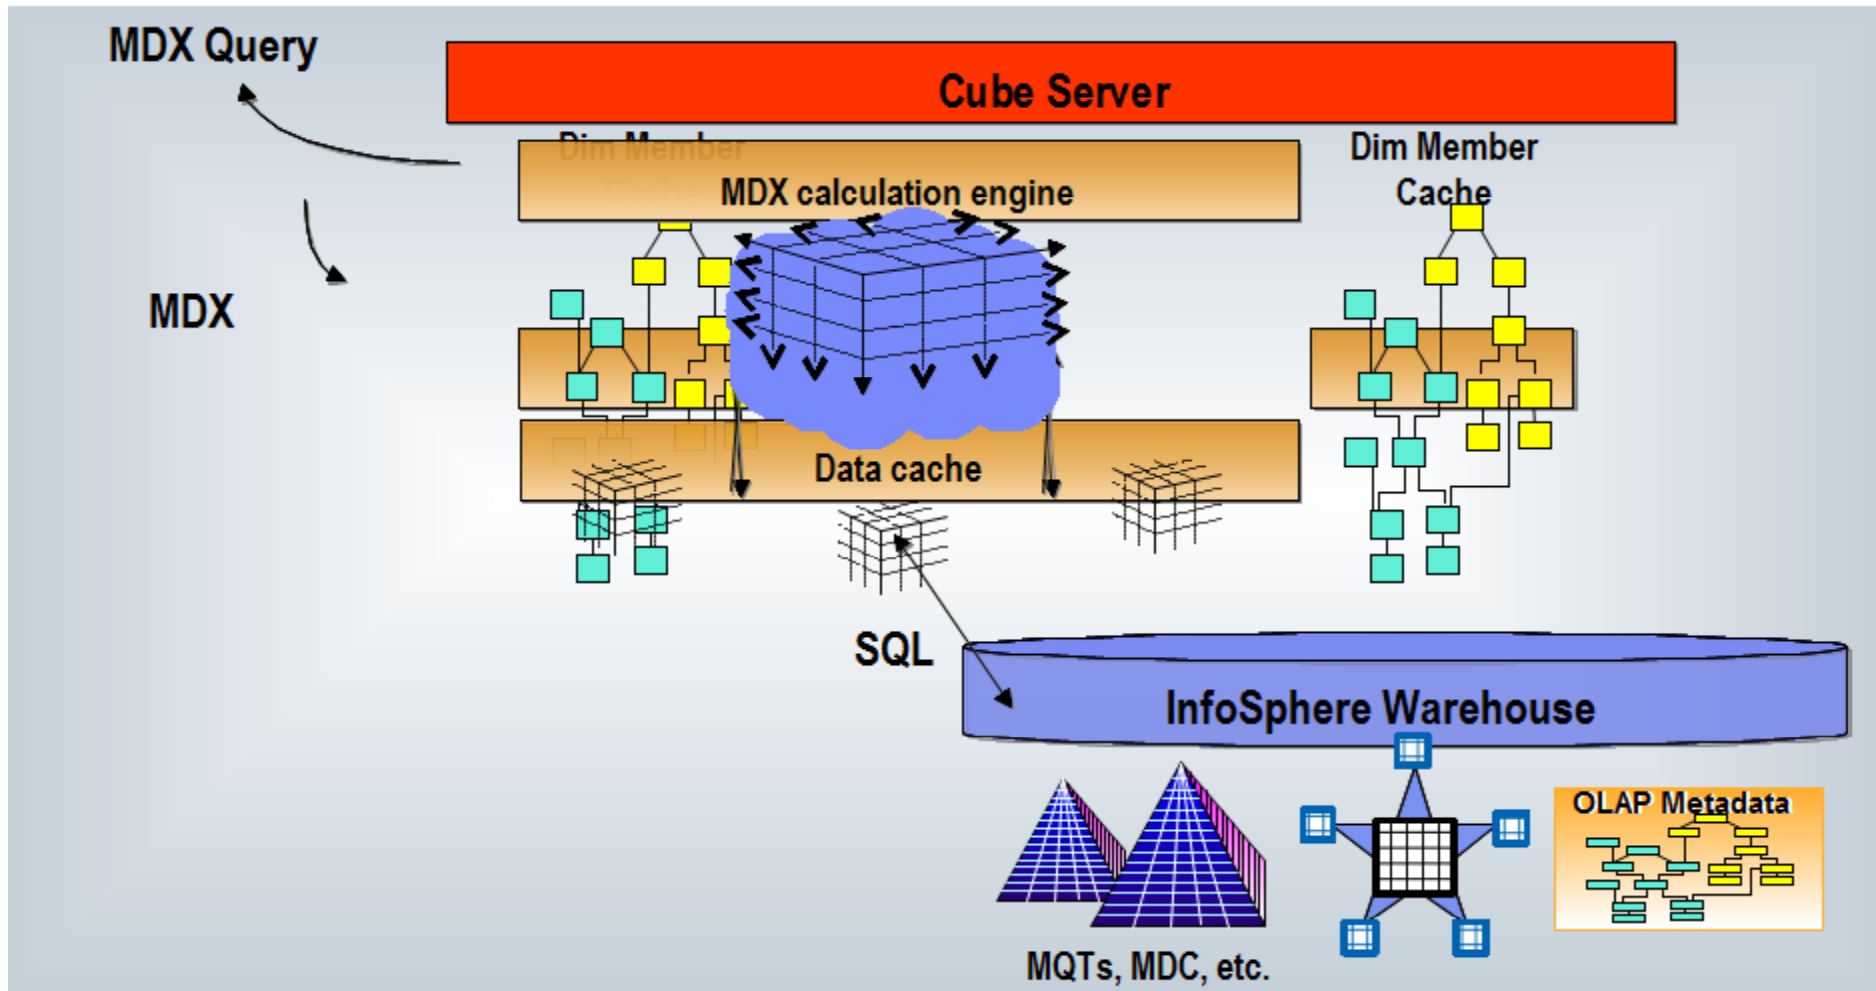
Demo: Cubing Services

Enables OLAP applications to access large data volumes stored inside the warehouse



Performance Method Number Three

- OLAP Acceleration = Cubing Services



InfoSphere Warehouse : Supporting Advanced Analytics



