

Henry Cook



InfoSphere Warehouse UKI 11/09/2010

IBM Smart Analytics System

IBM Smarter Systems Tour 2010



Customer Case Study



Challenges:

- More detailed predictive models
- Speed of development
- SAS support
- Scaling

Benefits:

- Faster model development
- Move to using 100's of detailed models
- Performance, TCO

IBM Smart Analytics System

What smart organizations want and need for faster results

EVERYTHING you need for Business Analytics
– not *just* a data warehouse appliance...

Analytics Software

- Business Intelligence Software
- Cubing Services
- Text Analytics & Data Mining
- Integrated Development Environment

Powerful Data Warehouse

- Warehousing Platform
- Optimisations
- Advanced Workload Management
- System Automation

Hardware & Services

- Flexible Server Platform Options
- Modular Storage Capacity
- Build, Deploy, Health Check and Premium Support Services



*Transforming
information into
business insight*


3x *Faster - Workload optimized analytics run
business intelligence 3x faster*

50% *less floor space*
Data compression reduces storage cost

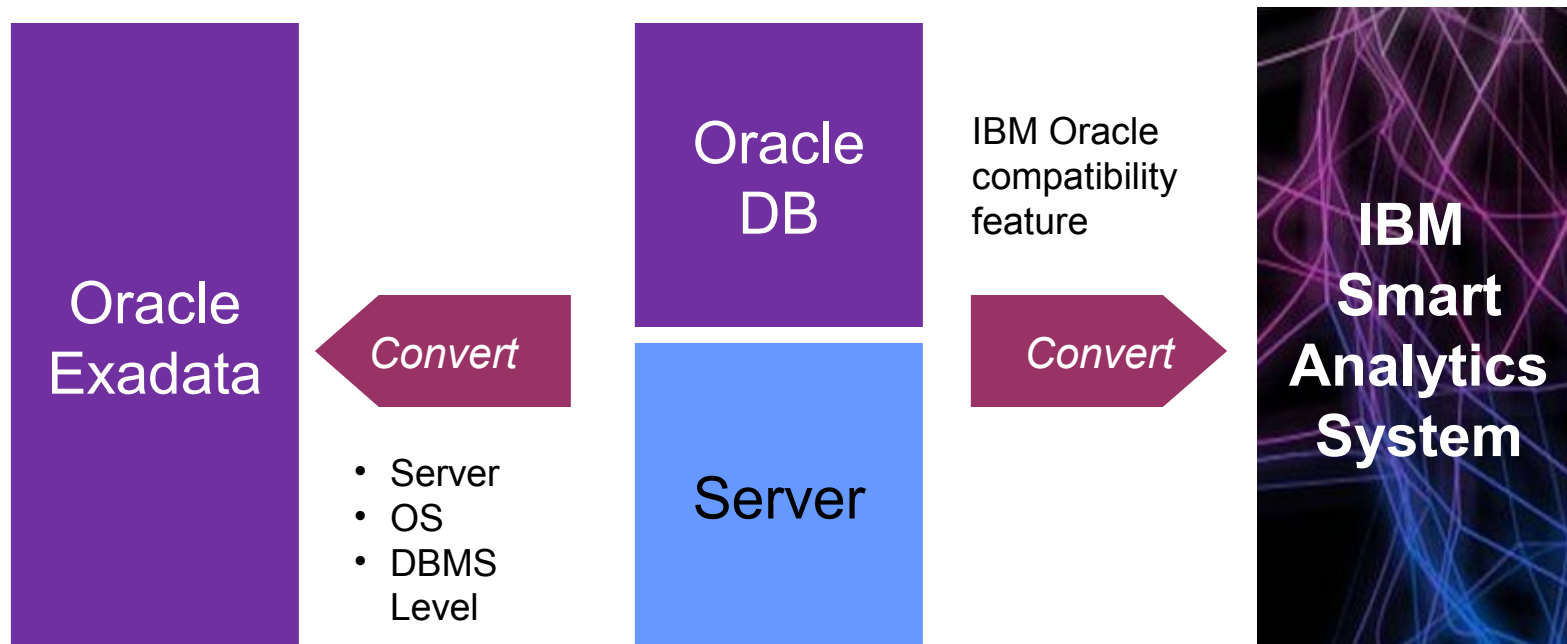
World record performance

The Evolution of IBM Warehousing Solutions

- IBM had been developing our analytics / warehousing capabilities since the early 90's.
- We were one of the first companies to produce a 'parallel edition' of our core DB2 database to provide scalability and performance for very large volumes of data. This has been enhanced and added to over many years
- The current solution is much more than just a core DBMS, and includes comprehensive features specific to analytics / data warehouse. These include specific optimisations, development tools, workload management, mining, data modelling, cubing support, embedded ETL, administrative functions, performance monitoring – all targeted at the analytics / warehouse problem
- In recent years this has included a range of complete hardware / software integrated appliances



IBM's Oracle Compatibility is a game changer. It may be no more effort to convert to IBM than to adopt Oracle Exadata




- Moving to Oracle Exadata means switching platform. You actually now have a choice.
- It may be just as easy to convert to IBM!
- You may obtain more value by switching to IBM, performance, cost etc


IBM Oracle Compatibility

ISV: “A conversion previously estimated at 2 years, took 1 week”

- Oracle compatible SQL
- Native PL/SQL
- Minimal impact on developers
- Simple administrator training
- “Like learning a new release of Oracle”



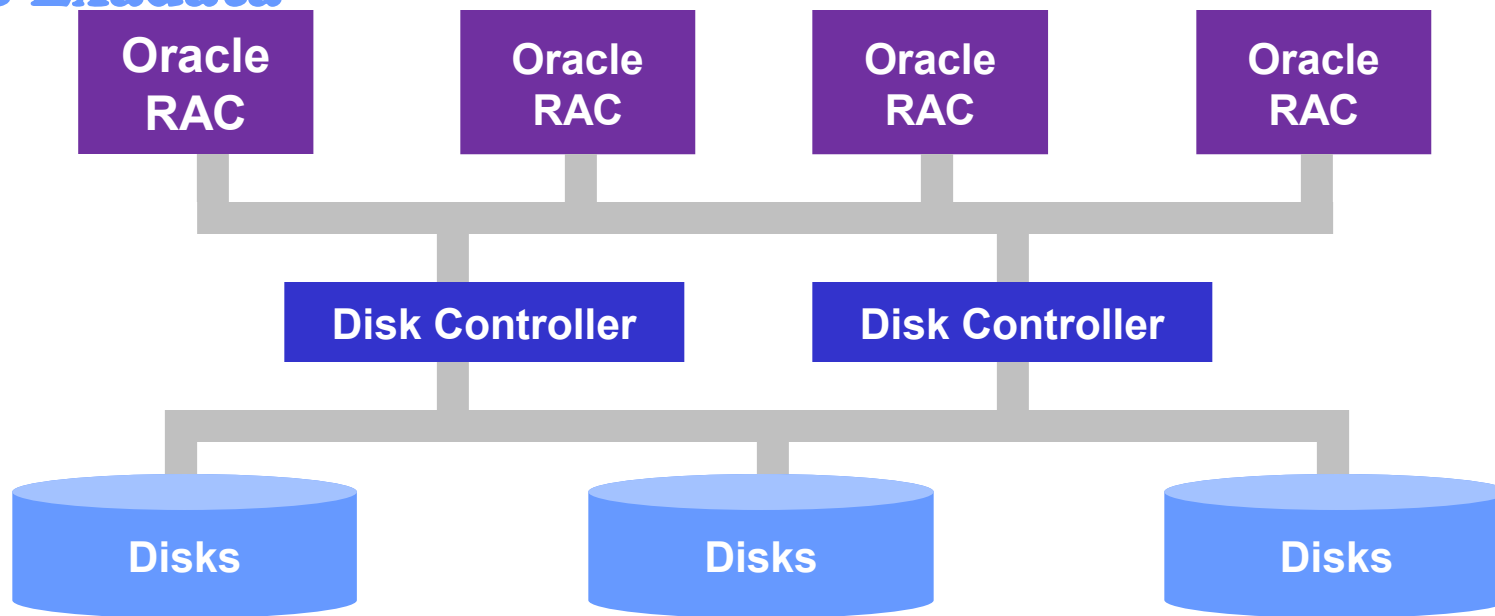
IBM Smart Analytics System uses a 'Shared Nothing' Architecture, the most suitable for Warehousing and Analytics



- True Shared Nothing architecture since 1995
- Single Layer of Servers
 - Less complex
 - All processors contribute to all processing
- Optimized for data warehouse workloads – Multi Dimensional Clustering etc
- Storage bandwidth is always added with each data server, True Linear Scaling


We can compare the architectural approach taken by Oracle

1) Pre-Exadata



- RAC Cluster
- Shared Disk Architecture
- Distributed Lock Manager

2) With Exadata. The approach taken for Exadata simply adds whole second layer of servers to assist with scanning



- RAC Cluster remains
- Storage Cells are SMP servers with attached storage
- Special extra disk software \$10k per disk !!!
- Two Levels of server
 - Complexity
 - Cost
- **The lower level only does basic scanning** all complex processing remains at the higher level
- **2½ times the hardware, more complex**

Characteristics of Online Transaction Processing vs. Data Warehouse

These workloads are diametrically opposed

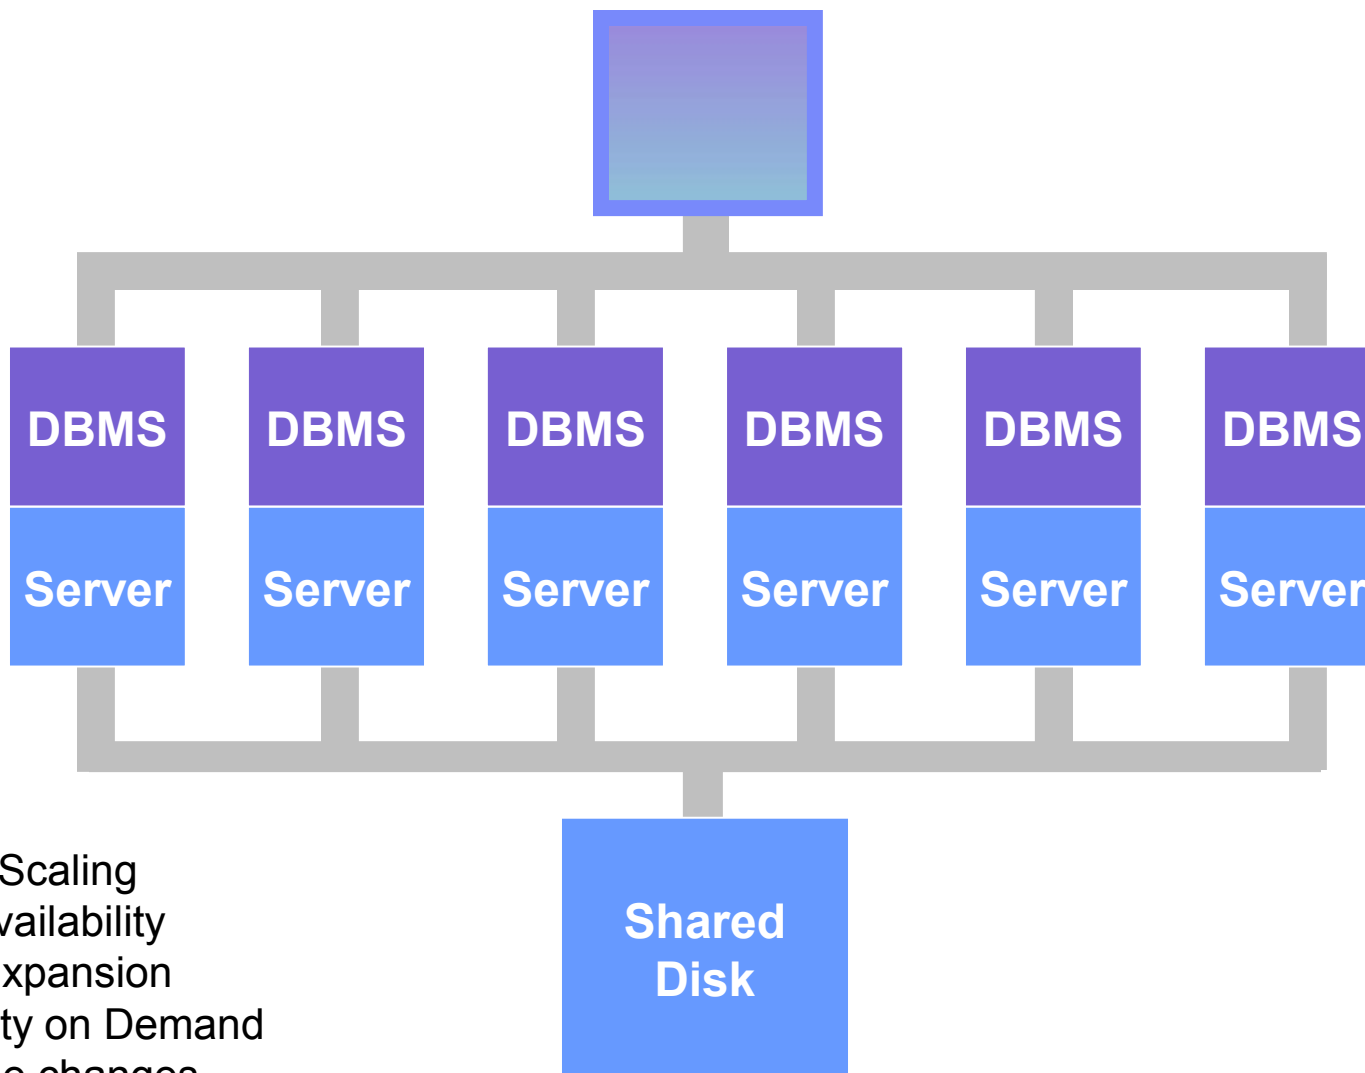
OLTP characteristics:

- Smaller database sizes ranging from a few Gigabytes to Terabytes
- I/O requirements that are primarily random I/O
- Data is more likely to benefit from in memory cache
- Transactions should minimize system resource usage as much as possible
- Need to scale compute power independent of storage and I/O
- Shared data architecture that can scale up and scale out is ideally suited for transaction based OLTP workloads

Data Warehouse characteristics:

- Larger database sizes ranging from hundreds of Gigabytes to Petabytes
- I/O requirements that are primarily sequential I/O
- Data is less likely to benefit from in memory cache
- Complex queries should maximize system resource usage as much as possible through parallelism
- Need to scale compute power with increase in storage and I/O
- Shared nothing architecture that can scale up and scale out is ideally suited for parallel query processing data warehouse workloads


IBM pureScale Application System



- OLTP
- Linear Scaling
- High Availability
- Easy Expansion
- Capacity on Demand
- No code changes

Smarter Systems: IBM Compresses *all* Aspects of the Database, and for *all* Workloads

Database Size Comparison



Efficient compression and XML data support

2x

Better Compression

77x

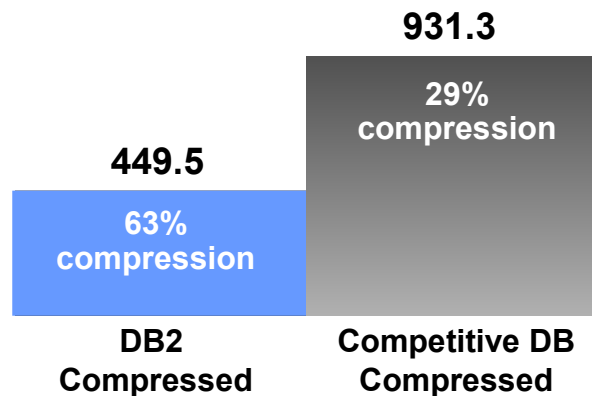
Better XML Query Performance

50%

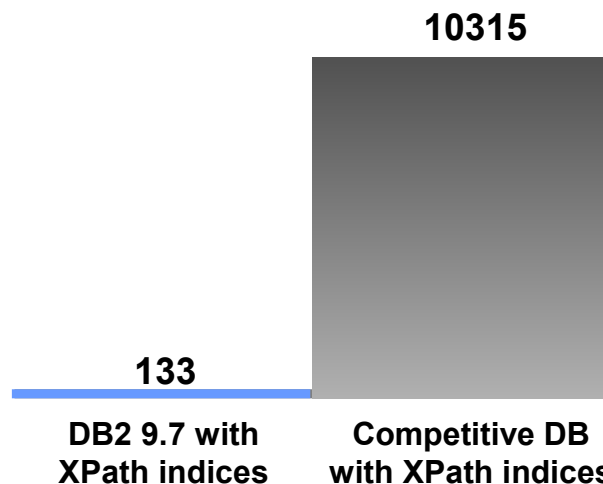
Less Storage

Source: IBM internal studies - compression based on TPC-H warehouse data, XML based on test cases from <http://tpox.sourceforge.net/>

Storage for 1200 MB Warehouse Benchmark



Time (sec) Per 1000 X Queries



Travel and leisure

"We searched for a solution that grows along with business requirements while keeping IT costs flat. Scalable and high performance functions have been a basic requirement to realise continuous, fast and knowledge based decisions"

Matthias Wunderlich, Head of BI

- **Business Challenge**
 - Tour operators are facing intense competition, need more agility
 - Consumers demand more choice
 - Fast and focused decisions required

- **What's Smart?**
 - New business model supported by more agile, more near real-time BI
 - Dynamic warehousing platform crunches higher volumes of data faster for operational and tactical decisions

- **Smarter Business Outcome**
 - Improve business agility, scalability and performance and keep IT costs flat

Case Study



TUI enters the future of travel with new intelligence from IBM

Overview

■ **Challenge**

To meet a range of requirements: faster, more stable information management for a new business model; web-based access to cross-system, standardised information; integration of business intelligence in operational applications; optimal decision-making foundation for a market-driven offering, CRM, yield management and strategic control; support for all the company's decision-making processes.

■ **Why IBM?**

Key factors included the IBM IT specialist's approach that was based on an integrative information agenda, and an end-to-end solution that offered an integrated and scaleable hardware package, a dynamic warehousing concept and software that provided Web-based access to all data.

■ **Solution**

IBM Cognos 8 BI, version 8.5, IBM InfoSphere Balanced Warehouse and IBM InfoSphere Datawarehouse

■ **Benefits**

Greater agility thanks to comprehensive and consolidated near real-time information; shorter product cycles; greater flexibility; lower costs thanks to an integrated, company-wide, streamlined system; well-founded decisions and simulations; involvement of all employees in company intelligence.



With a turnover of 5.2 billion euros and a market share of around 22 percent, TUI Germany is the undisputed market leader in Germany. The tour operator is a subsidiary of TUI Travel PLC. The company brings together tour operators, airlines, travel agents and incoming agencies to serve more than 30 million customers. TUI Germany is responsible for the source market of Central Europe (Germany, Austria, Switzerland and Eastern Europe). It manages a host of well-known brand names, including TUI, aitours, 1-2-Fly, Dr. Tigges, Gebeco, Itur und Berge & Meer, to name but a few.

Challenges faced

"Improve or perish." This maxim holds particularly true in the volatile tourism market. TUI, a global organization, quickly realised that the key to agility lies in integrated information management. TUI Germany implemented a Business Intelligence solution (BI) back in 1996. Since then, the company has continued to evolve, taking tourist travel into the future with the introduction of a new and more flexible business model. And this concept calls for efficient and near real-time analysis of business data. TUI Germany called upon IBM® Software Group Services to advise it

Information Management

Cognos.
software