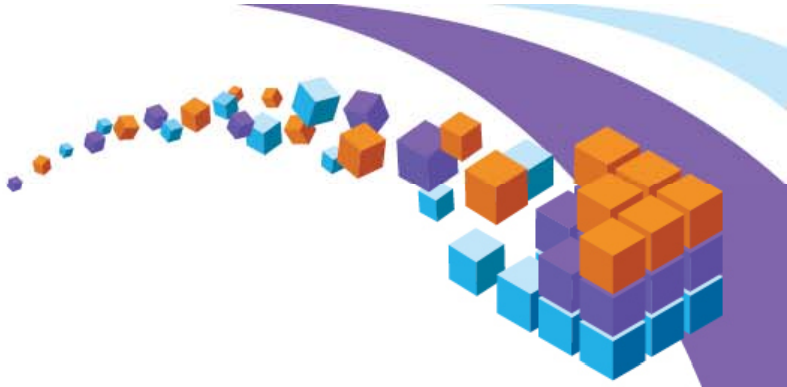

IBM INFORMATION INTEGRATION
& GOVERNANCE SYMPOSIUM 2012

Delivering Trusted Information for Smarter Business Decisions

Patient Identity Management in Singapore

Anthony O'Neill – Lead Architect

MOH Holdings Pty Ltd

01/05/2012

Singapore success factors

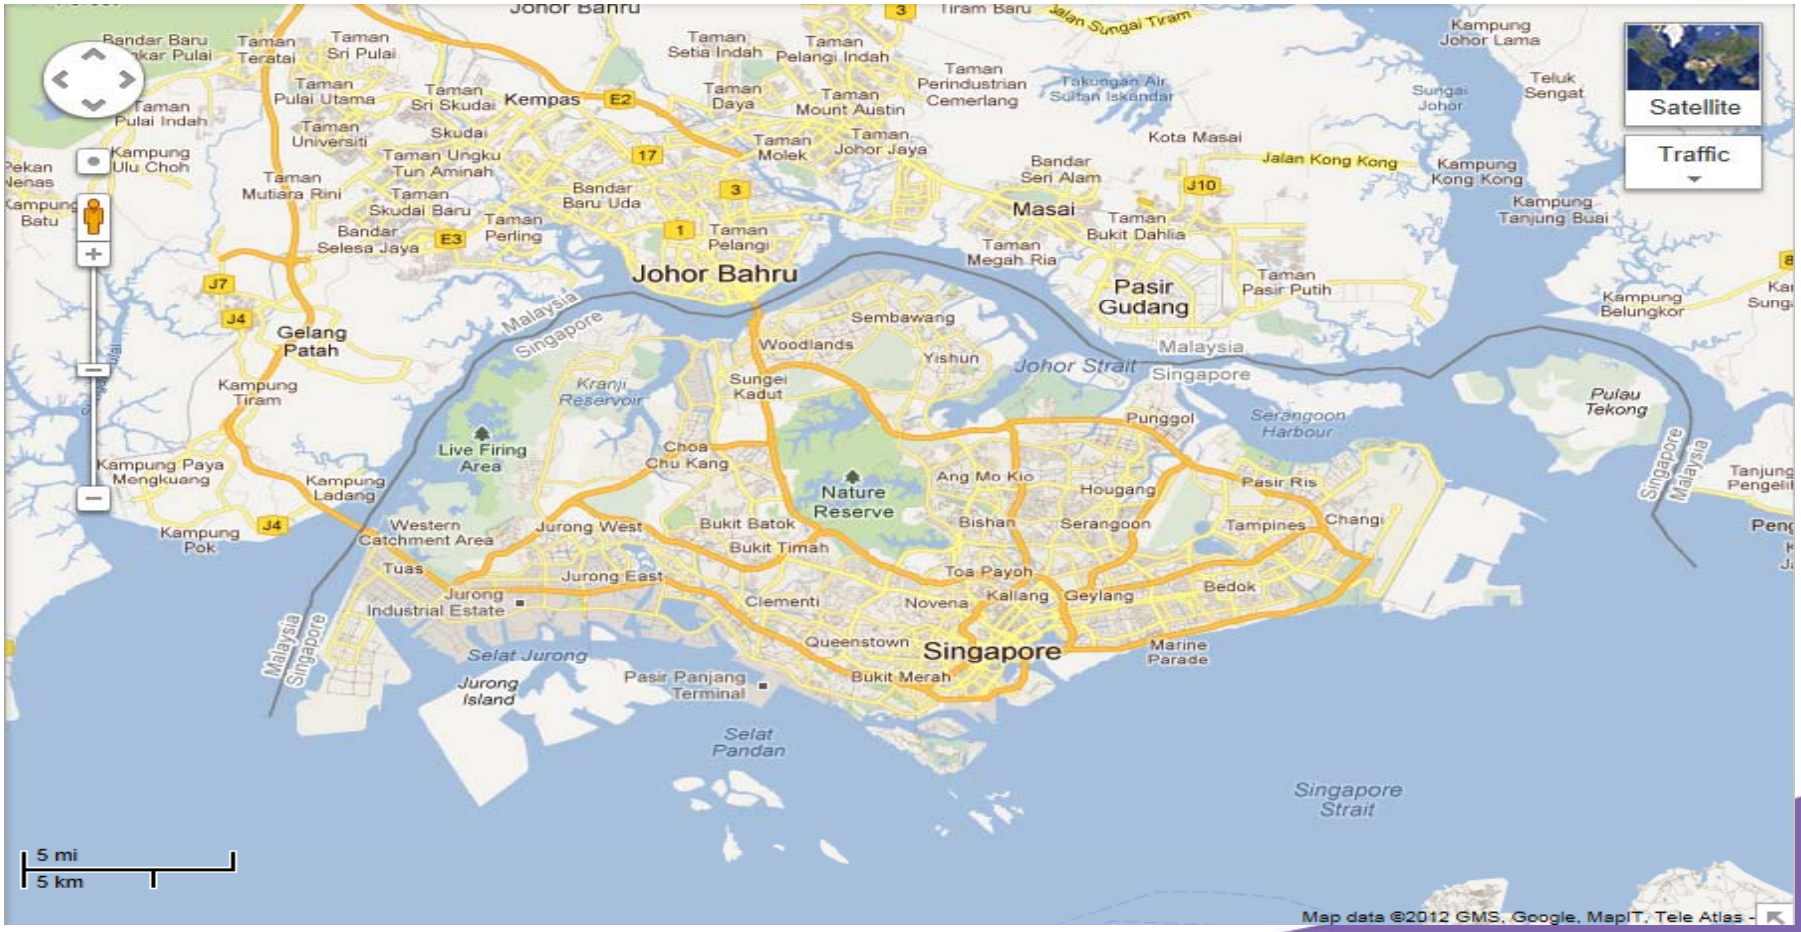

Singapore success factors

- Tiny
- Densely populated
- Highly urbanised
- Tech savvy
- Stable, single-tier government
- Foreign expertise
- Opt out

Information flow

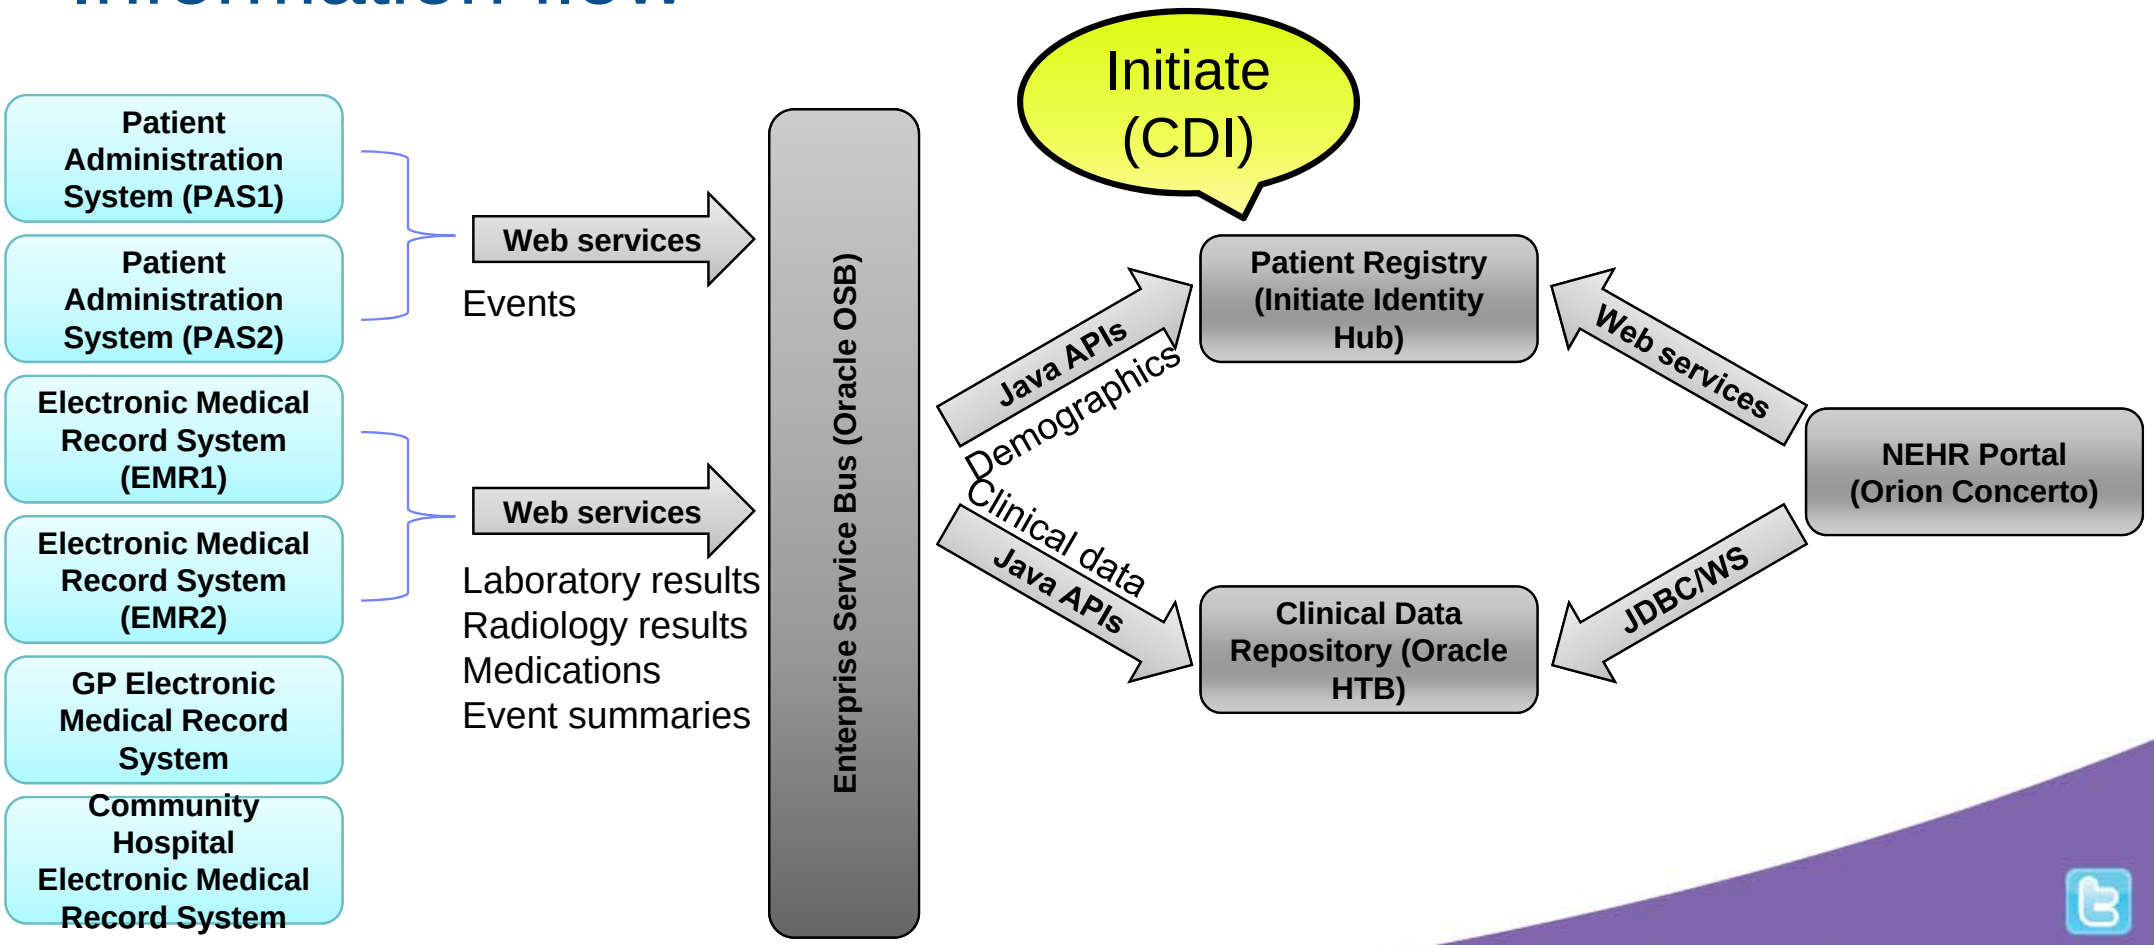

Singapore context

- All Singapore residents must be registered with the National Registry
  - 3,230,700 Citizens
  - 541,000 Permanent Residents
  - 1,305,000 Foreign Residents
  - 5,076,700 Total (2010 population census)
- All residents are issued with an Identity Card (IC)
  - card has unique identifier with check digit, plus portrait photo, plus thumbprint…
- Process is formal
  - pre-application followed by interview with proof of identity
- Penalties apply for:
  - failing to register within the prescribed period
  - failing to notify address changes within 28 days
  - providing false or misleading information
  - possessing a forged identity card or another person's identity card
- Healthcare institutions use IC number to identify patients

IC number penetration in healthcare records

- Almost 20% of all patient records do not have a valid IC number
  - Don’t have an IC number (e.g. tourists)
  - Left IC at home
  - ???

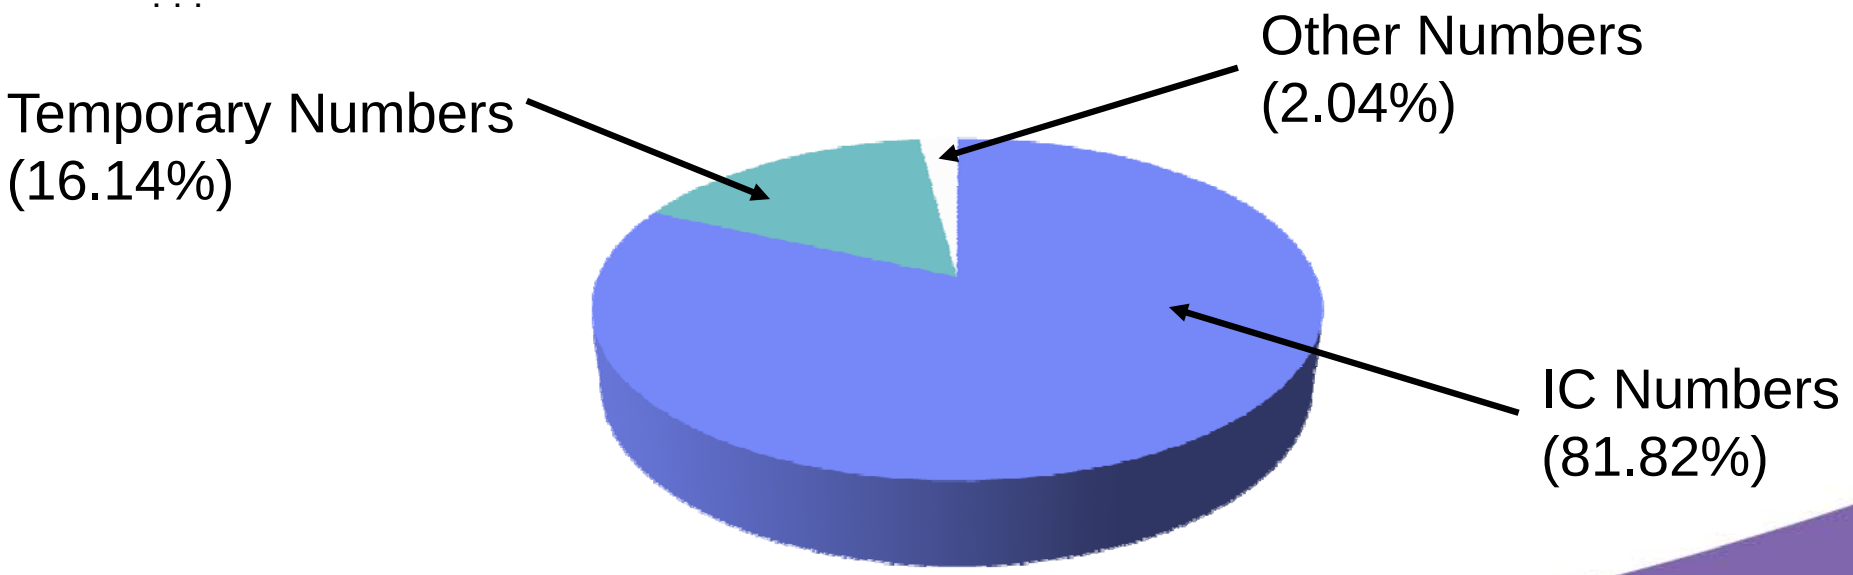

IC number penetration in healthcare records

- Records using IC number as MRN are duplicated within sources
  - IC number associated with primary key is not equal to primary key (foreign key without unique constraint)
  - Dependency on external system (tight coupling)

PAS Database		
Primary Key 	MRN	Name
90014578	S0475288Z	Lee
90023629	S1996256H	Tom
90039954	T0019879J	Mary
90041727	G6092662Y	Sue
90057321	S1996256H	Tom

IC number penetration in healthcare records

- Records using IC number as MRN are duplicated within sources
  - IC number associated with primary key is not equal to primary key (foreign key without unique constraint)
  - Dependency on external system (tight coupling)

PAS Database		
Primary Key 	MRN	Name
90014578	S0475288Z	Lee
90023629	S1996256H	Tom
90039954	T0019879J	Mary
90041727	G6092662Y	Sue
90057321	S1996256H	Tom

Which record?

IC number penetration in healthcare records

- Different patients share the same IC
  - ???

PAS Database		
Primary Key 	MRN	Name
90014578	S0475288Z	Lee
90023629	S1996256H	Tom
90039954	T0019879J	Mary
90041727	G6092662Y	Sue
90057321	S1996256H	Tom

IC number penetration in healthcare records

- Different patients share the same IC
  - ???

PAS Database		
Primary Key 	MRN	Name
90014578	S0475288Z	Lee
90023629	S1996256H	Tom
90039954	T0019879J	Mary
90041727	G6092662Y	Sue
90057321	S1996256H	Tom

Select from EHR where:

Source System = EMR and MRN = S1996256H

Benefits of NHIS

- 16.8 million records representing 6.3 million patients (30 April 2011)
  - Links possible using IC number – 13.7 million
  - Links made by NHIS – 14.7 million

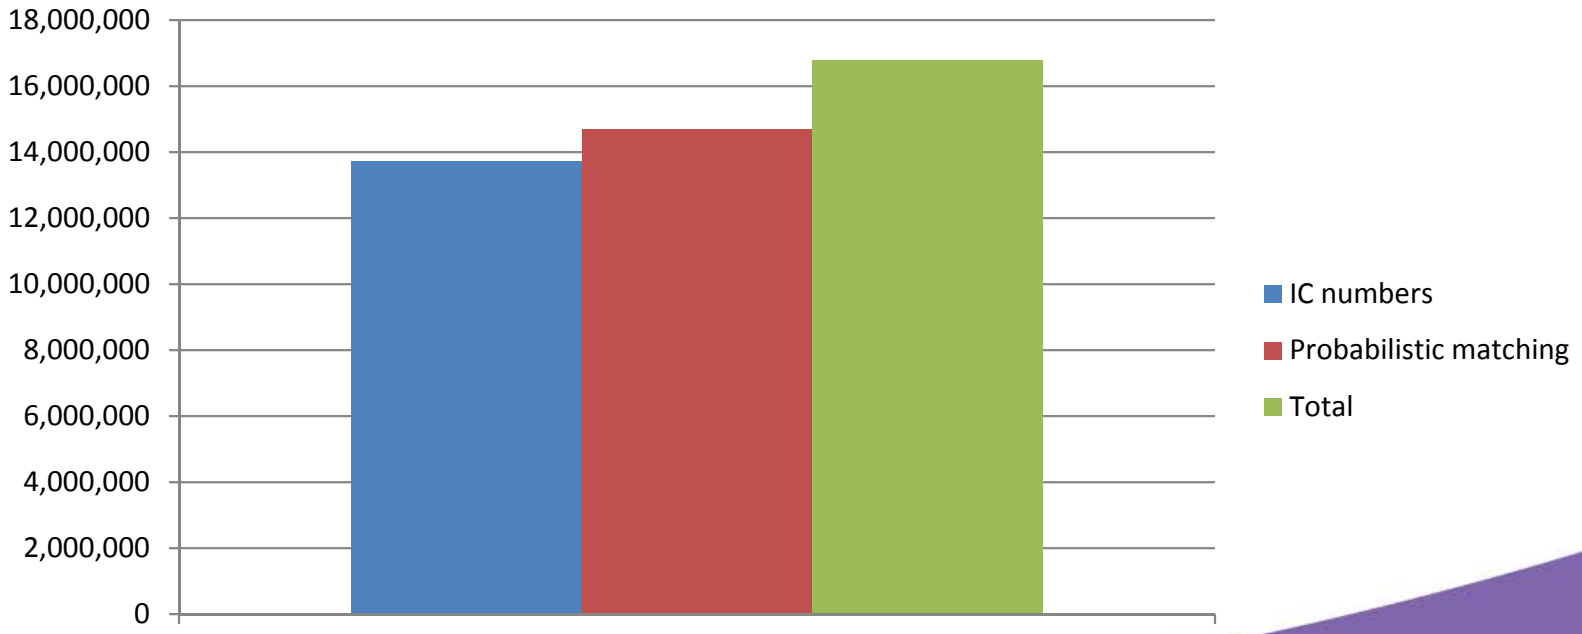

Challenges

- 1 million *Potential Duplicate* tasks
  - At $100 per merge = $100 million dollars
  - At 60 minutes per merge = 480 years (1 person, 40 hours per week, no holidays)
  - Clinical risk - until resolved, patient record may be incomplete
- 1.5 million *Potential Linkage* tasks
  - Resolved only through improvement in data quality in terms of accuracy, completeness and currency
  - Clinical risk - until resolved, patient record may be incomplete
- 100 thousand *Review Identifier* tasks
  - High clinical risk – until resolved, patient record may include clinical data from multiple patients

Possible approaches to improve linking

1. Standardise information models
   - e.g. AS4590 – Interchange of Client Information

2. MDM for CDI

Standard information models – a physical tactic

Name
Surname
First given name
Second given name
Other given name(s)

Name
Last name
First name
Other name(s)

Name
Surname
Given name(s)

Address
Unit number
Floor number
Street number
Street name
Suburb
Post code

Address
Street number
Street name
Street type
Suburb
Zip

Address
Line 1
Line 2
Suburb
Post code

Why standard approach can't succeed

- Standard needs to be more granular than typical information models
  - Physical issues - less granular can understand more granular. Reverse not always true.
  - Data issues – neither direction is always true (e.g. Chinese v Western names).

Name
Surname = Lee
First given name = Ching Mee
Other given name(s) = Bruce

Name
Name = Bruce Lee Ching Mee

Name
Name = Anthony O'Neill

- Historical data issues are not addressed
- Most healthcare system acquisitions are COTS products from *international* vendors
  - Compliance with national standards is improbable
- Other issues that prevent linking are not addressed
  - accuracy, completeness and currency of data
  - variation in processes for collecting data

MDM

- Recognises heterogeneity as a reality
  - Information models are different
  - Business processes are different
- Yields better results than identifier-based linking out of the box
  - 1 million more links in Singapore (87.5% v 81.6%)
- Provides tools to identify and manage data quality issues
  - Potential duplicates
  - Potential linkages
  - Potential overlays
  - Review identifiers
- Can facilitate access to authoritative sources of data at point of registration
  - IC number against IC database
  - Demographics belong to IC number
  - Address is valid

Lots of work to do

- Address existing data quality issues
  - Effort needs to be linked to clinical imperative
- Prevent new issues
  - Impacts existing processes and system capabilities
- The next area of focus in Singapore
  - Going beyond identity validation – how to prove the identity belongs to the patient?

Summary

1. Unique patient identifier efforts do not solve the identification challenge
   - *If it can't be done in Singapore, it can't be done anywhere*
2. MDM tools improve identification, but considerable effort remains
   - *Resolve known data quality issues*
   - *Prevent the creation of new data quality issues:*
     - *Improve identification at the point of registration*
     - *leverage authoritative sources to validate patient demographic information*
   - *Maintain data currency*
3. A proxy for a primary key is not the same as a primary key
   - *Referential integrity works within RDBMS, not across RDBMS*
   - *Australia's UHI will not be the primary key in most, if not all databases*
4. Use Web services, not health-specific messaging (e.g. HL7)
   - *Enables access to authoritative sources, especially those outside healthcare*

Questions?

Thank you

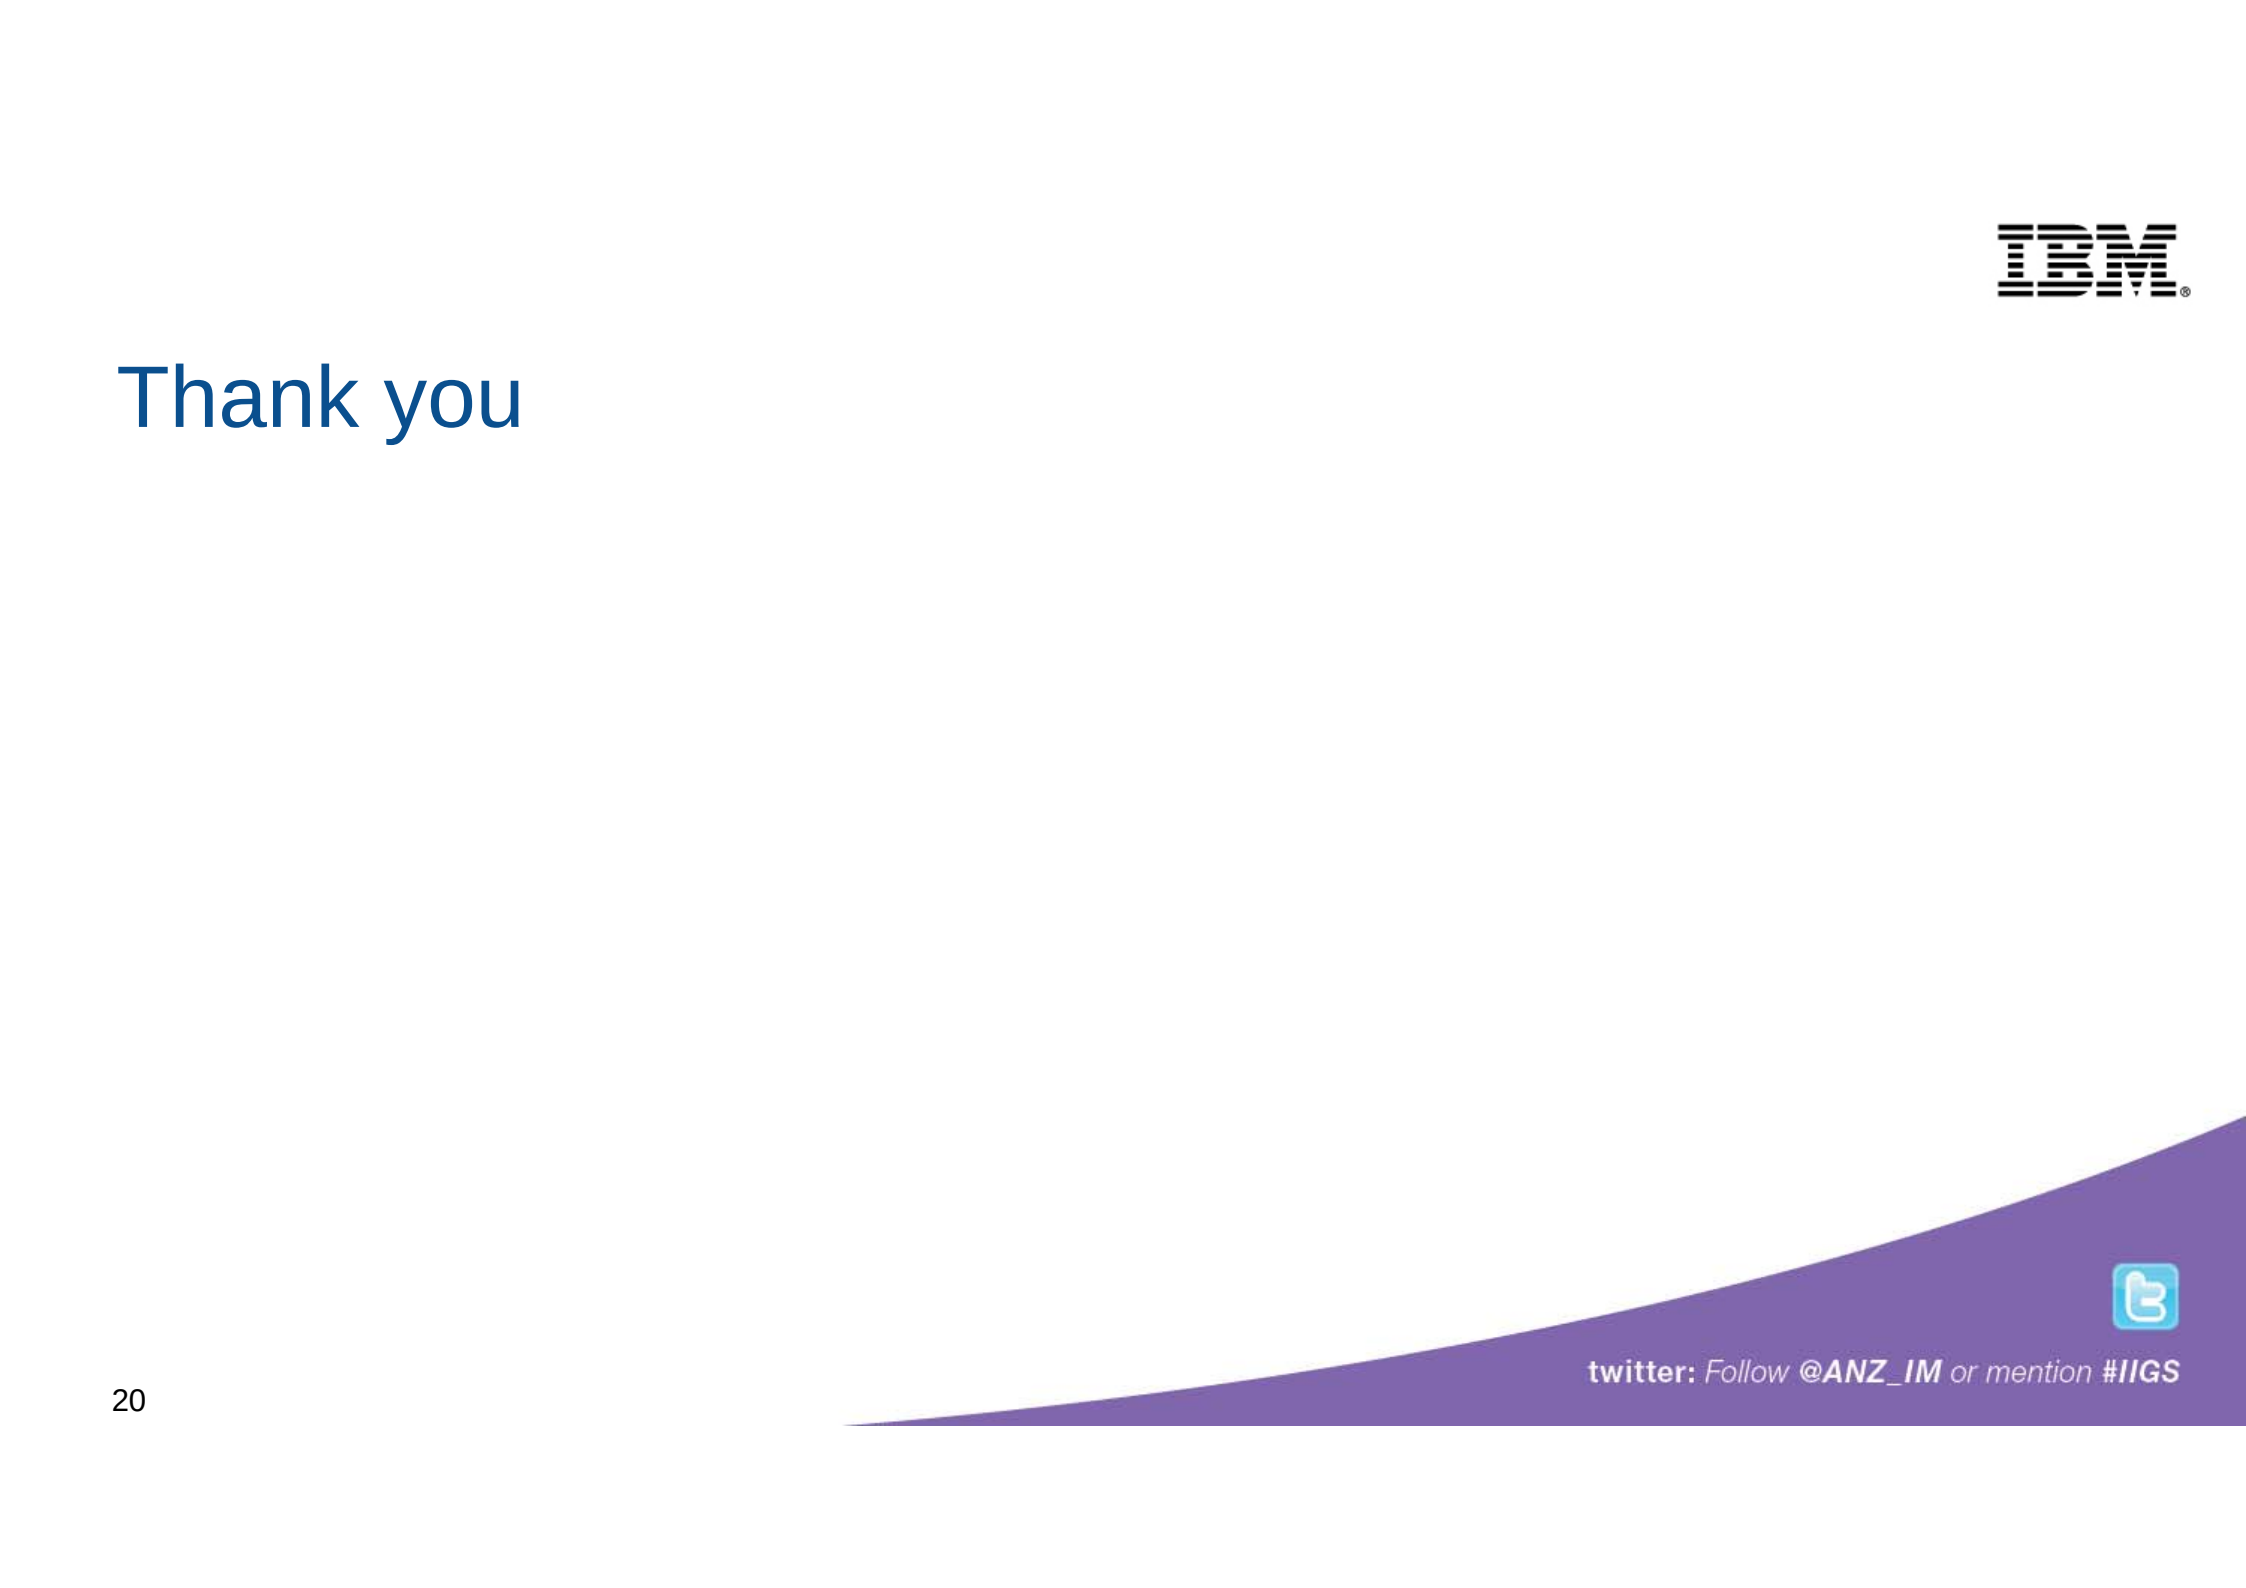

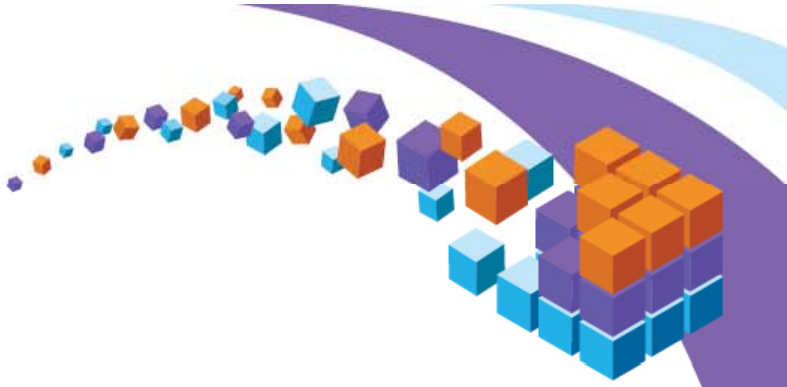

IBM INFORMATION INTEGRATION
& GOVERNANCE SYMPOSIUM 2012

Delivering Trusted Information for Smarter Business Decisions

IBM Data Management

Sven Hansen – Sales Lead, Data Management.

01/05/2012

The Business Value of DB2 for Linux, UNIX, and Windows

- ✓ **Lower administration requirements than other DBMSs**
  - Means less DBA resources for day-to-day activities
- ✓ **Better performance per core than other DBMSs**
  - Means less cores required for the same work
- ✓ **Better compression than other DBMSs**
  - Means less storage and improved performance
- ✓ **Better availability & Scalability than other DBMSs**
  - Means less down time, no loss of goodwill/opportunity
- ✓ **Better workload management than other DBMSs**
  - Means precise control of applications to ensure service level agreements and performance goals are delivered
- ✓ **Better development - Supports PL/SQL and T-SQL with little to no change**
  - If you want consolidate Oracle and/or Sybase applications, you can do that on DB2 with little or no change

Announcing DB2 10!!

Lowest TCO
Unparalleled Automation
Deep Compression
Lightning Fast

Simple to Use
Flexible Development
Industry leading XML support
Self-Managing

Most Reliable
World class audit & security
Easy High Availability
Workload Management

The Business Value of Informix Dynamic Server (IDS)

- ✓ **Reliable:** Provides one of the industry's widest sets of options for keeping data available at all times, including zero downtime for maintenance.
- ✓ **Low Cost:** Cuts development costs by nearly eliminating up-front licensing costs typically incurred during the development phase
- ✓ **Hassle Free:** Runs virtually unattended with self-configuring, self-managing and self-healing capabilities
- ✓ **Best-of-Breed Embeddability:** Provides a proven embedded data management platform for ISVs and OEMs to deliver integrated, world-class solutions, enabling platform independence.

Database

Warehouse

Analytics

Development

IBM solidDB: An In-Memory Relational DBMS

- ✓ **Extreme Speed:** Designed to achieve tens-to-hundreds of thousands of transactions per second. Keeps data in main memory at all times.
- ✓ **Extreme Availability:** Supports 99.9999% high availability. Provides instant application failover and transparency to users.
- ✓ **Low Cost:** Ability to run virtually unattended. Sustains high-throughput workloads with less hardware
than disk-based databases.
- ✓ **Data Caching:** Accelerate access to enterprise on-disk databases. DB2 LUW/I/Z, Informix, Oracle, SQL Server, Sybase.

Extreme Speed

Extreme Availability

Low Cost

Database Caching

IBM DB2 Delivers Price Performance Over Oracle

DB2 has lower staffing costs

- DB2 needs **43%** Less Staff on Power than Oracle Database – Solitaire Study
- **25%** to **35%** more DBA time is required for Oracle than for DB2 9 environments – ITG Report

DB2 has better performance

- **38%** better price/performance: DB2 pureScale on Power 780 vs. Oracle RAC on Nehalem
- Supports **3.3 times** more users on an SAP benchmark than Oracle running on Sun

ITG Report - 3-year costs for DB2 and Oracle Database***

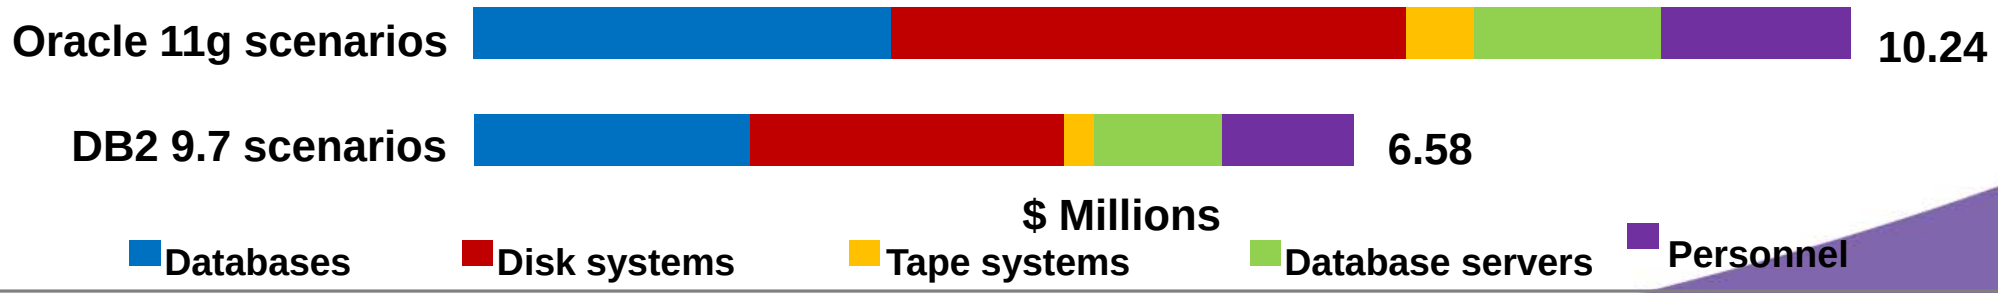

* Source: Oracle technology global price list. Based on comparison of US Prices of single processor core, equivalent of 100 PVU's.
** Source: IBM internal tests. Configuration detail: IBM WebSphere Application Server 7 - 1 JVM , AIX TL4, 64 bit, 16 threads. Competitive Application Server - 1 JVM, Windows 64 bit, 16 threads
*** Source: ITG whitepaper: VALUE PROPOSITION FOR IBM DB2 9.7: Cost Savings Potential Compared to Oracle Database 11g

Endorsements from Organisations Including….

twitter: Follow @ANZ_IM or mention #IIGS

Introducing….

Randall Ibbott, Technical Lead Database Administrat

- Been working for QBE for 24 years
- Started at QBE as Computer Operations
- Moved into Production Support
- Now the Senior Database Manager for DB2.
- Started using DB2 on System z (Mainframe) in 1993.
- Working with DB2 LUW for the last 10 years.
- QBE chose DB2 because of there fantastic experience with using it on the Sys z.

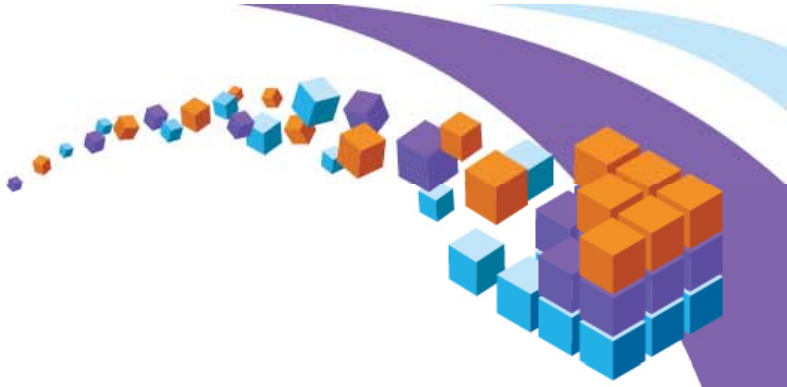

IBM INFORMATION INTEGRATION
& GOVERNANCE SYMPOSIUM 2012

Delivering Trusted Information for Smarter Business Decisions

Moving With the Times…

Database Technology in Action

Randall Ibbott – Lead Technical Consultant

05/01/2012

Disclaimer…

This presentation describes QBE's experience from a technical perspective only and does not necessarily reflect QBE's official position with regard to any operational or management decisions.

Two Important Principles

- “Horses for Courses” - Appropriate Use of Technology
- “Bang for Bucks” - Cost Effective Use of Technology
- DB2 Ticks Both Boxes…

QBE

IBM DB2

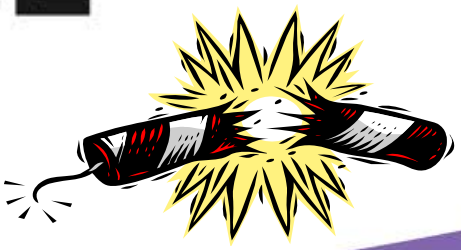

Historical Overview - Life was Simple

- MVS/SP
- Supra Network Database Engine
- CICS
- Green Screen
- SLA 8:00 – 20:00. Monday to Friday

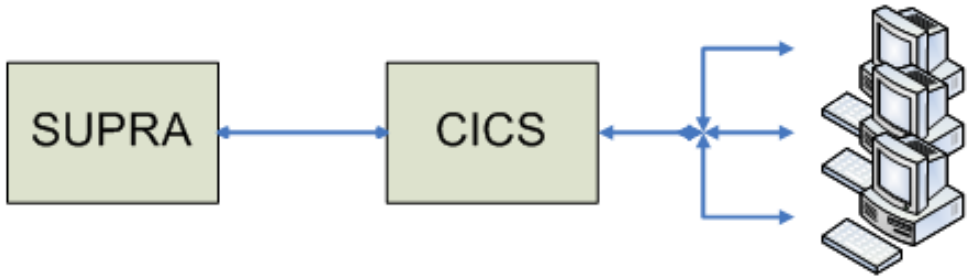

Impacting Factors

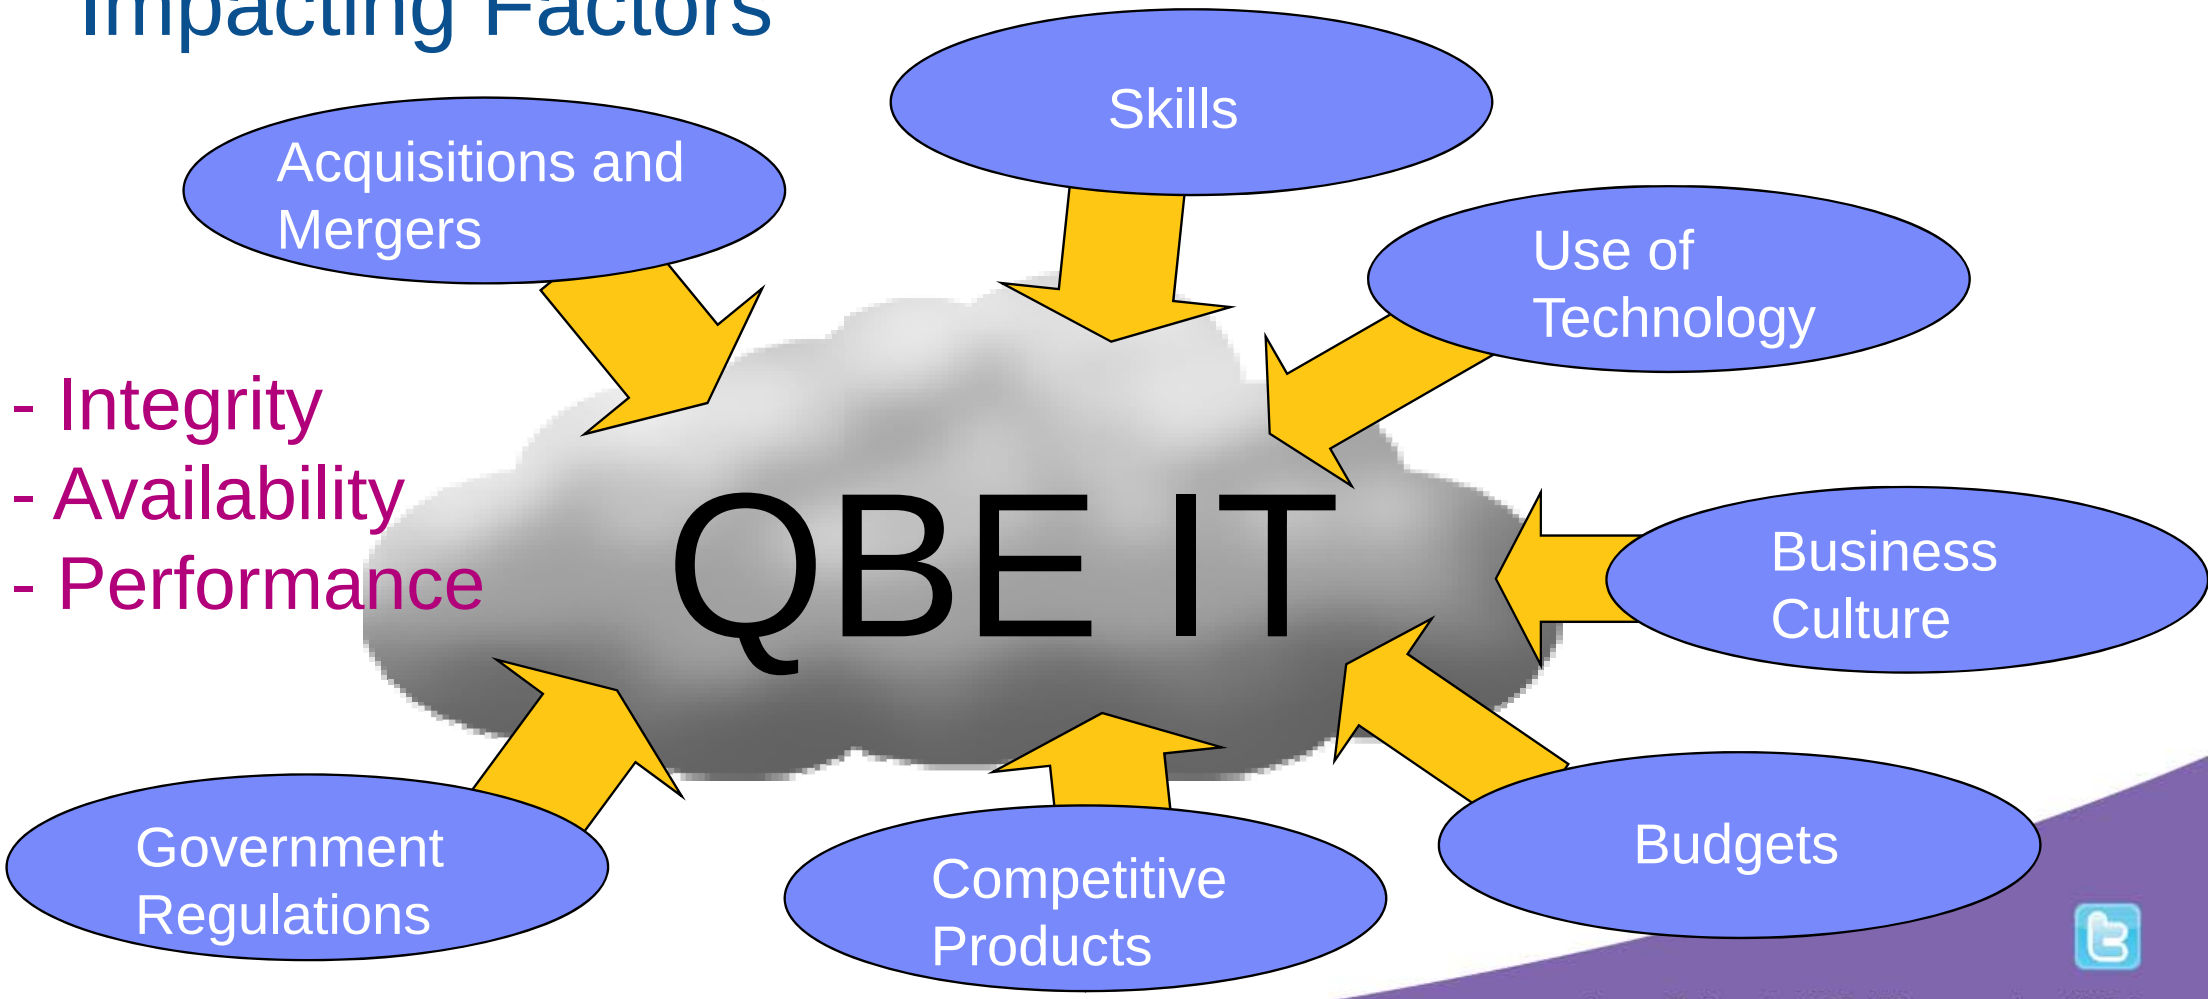

- Integrity
- Availability
- Performance

First Steps Engaging DB2 Technology

- Moving to the Web – c.change
- Changing Face of QBE's User Interfaces
  - DB2 for Z/OS (Dynamic SQL)
  - CICS to MQ-Series to Legacy Applications

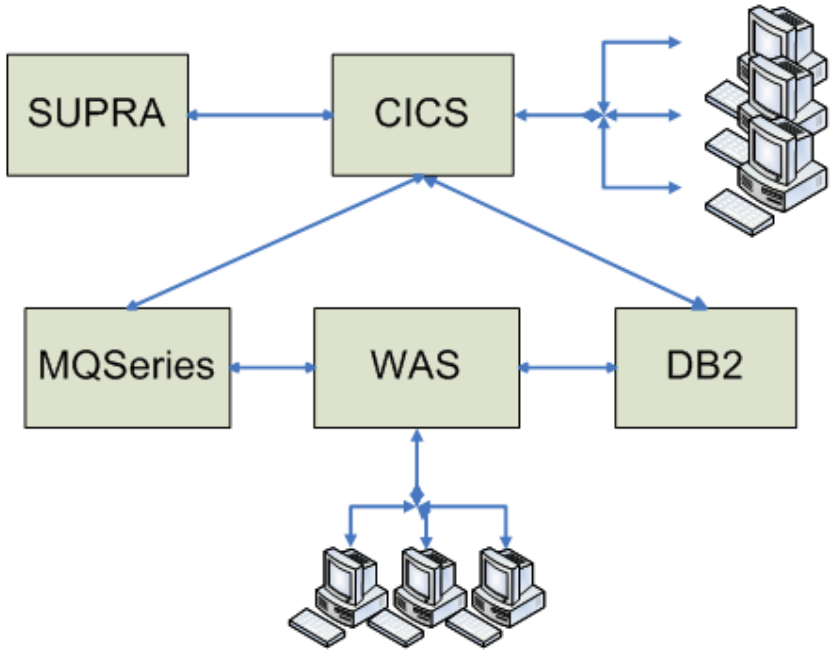

OLTP Consolidation

- Multi-Faceted Business Processes
  - Leveraging Functionality from Diverse Components (Pega, OCS, Claims Centre, Dialogue, etc.)
  - DB2, Oracle, SQL Server, Legacy Applications - all need to communicate seamlessly to provide a business transaction
- MQ, Message Broker - Enterprise Services Bus (ESB)
- Migrating Functionality from Legacy Systems to DB2
- Warehouse Feeds from Multiple Disparate Applications

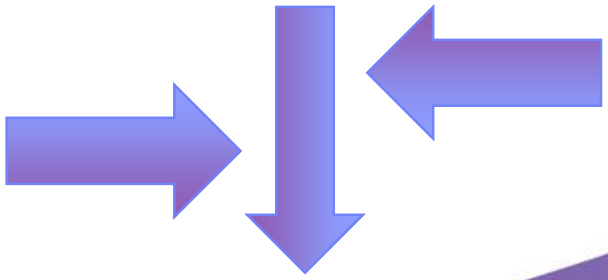

Historical Overview – Now

- 20 Years Later…
  - System Z, i90, P770, EMC
  - Z/OS, AIX, Linux, i90, Windows Server
  - Supra, DB2 for Z/OS, DB2 for LUW, Oracle, SQL Server, DB400
  - MQSeries, WebSphere Application Server, Message Broker
  - DataStage
  - Pega Systems (Business Rules)
  - Guidewire Claimcentre
  - BMC Remedy

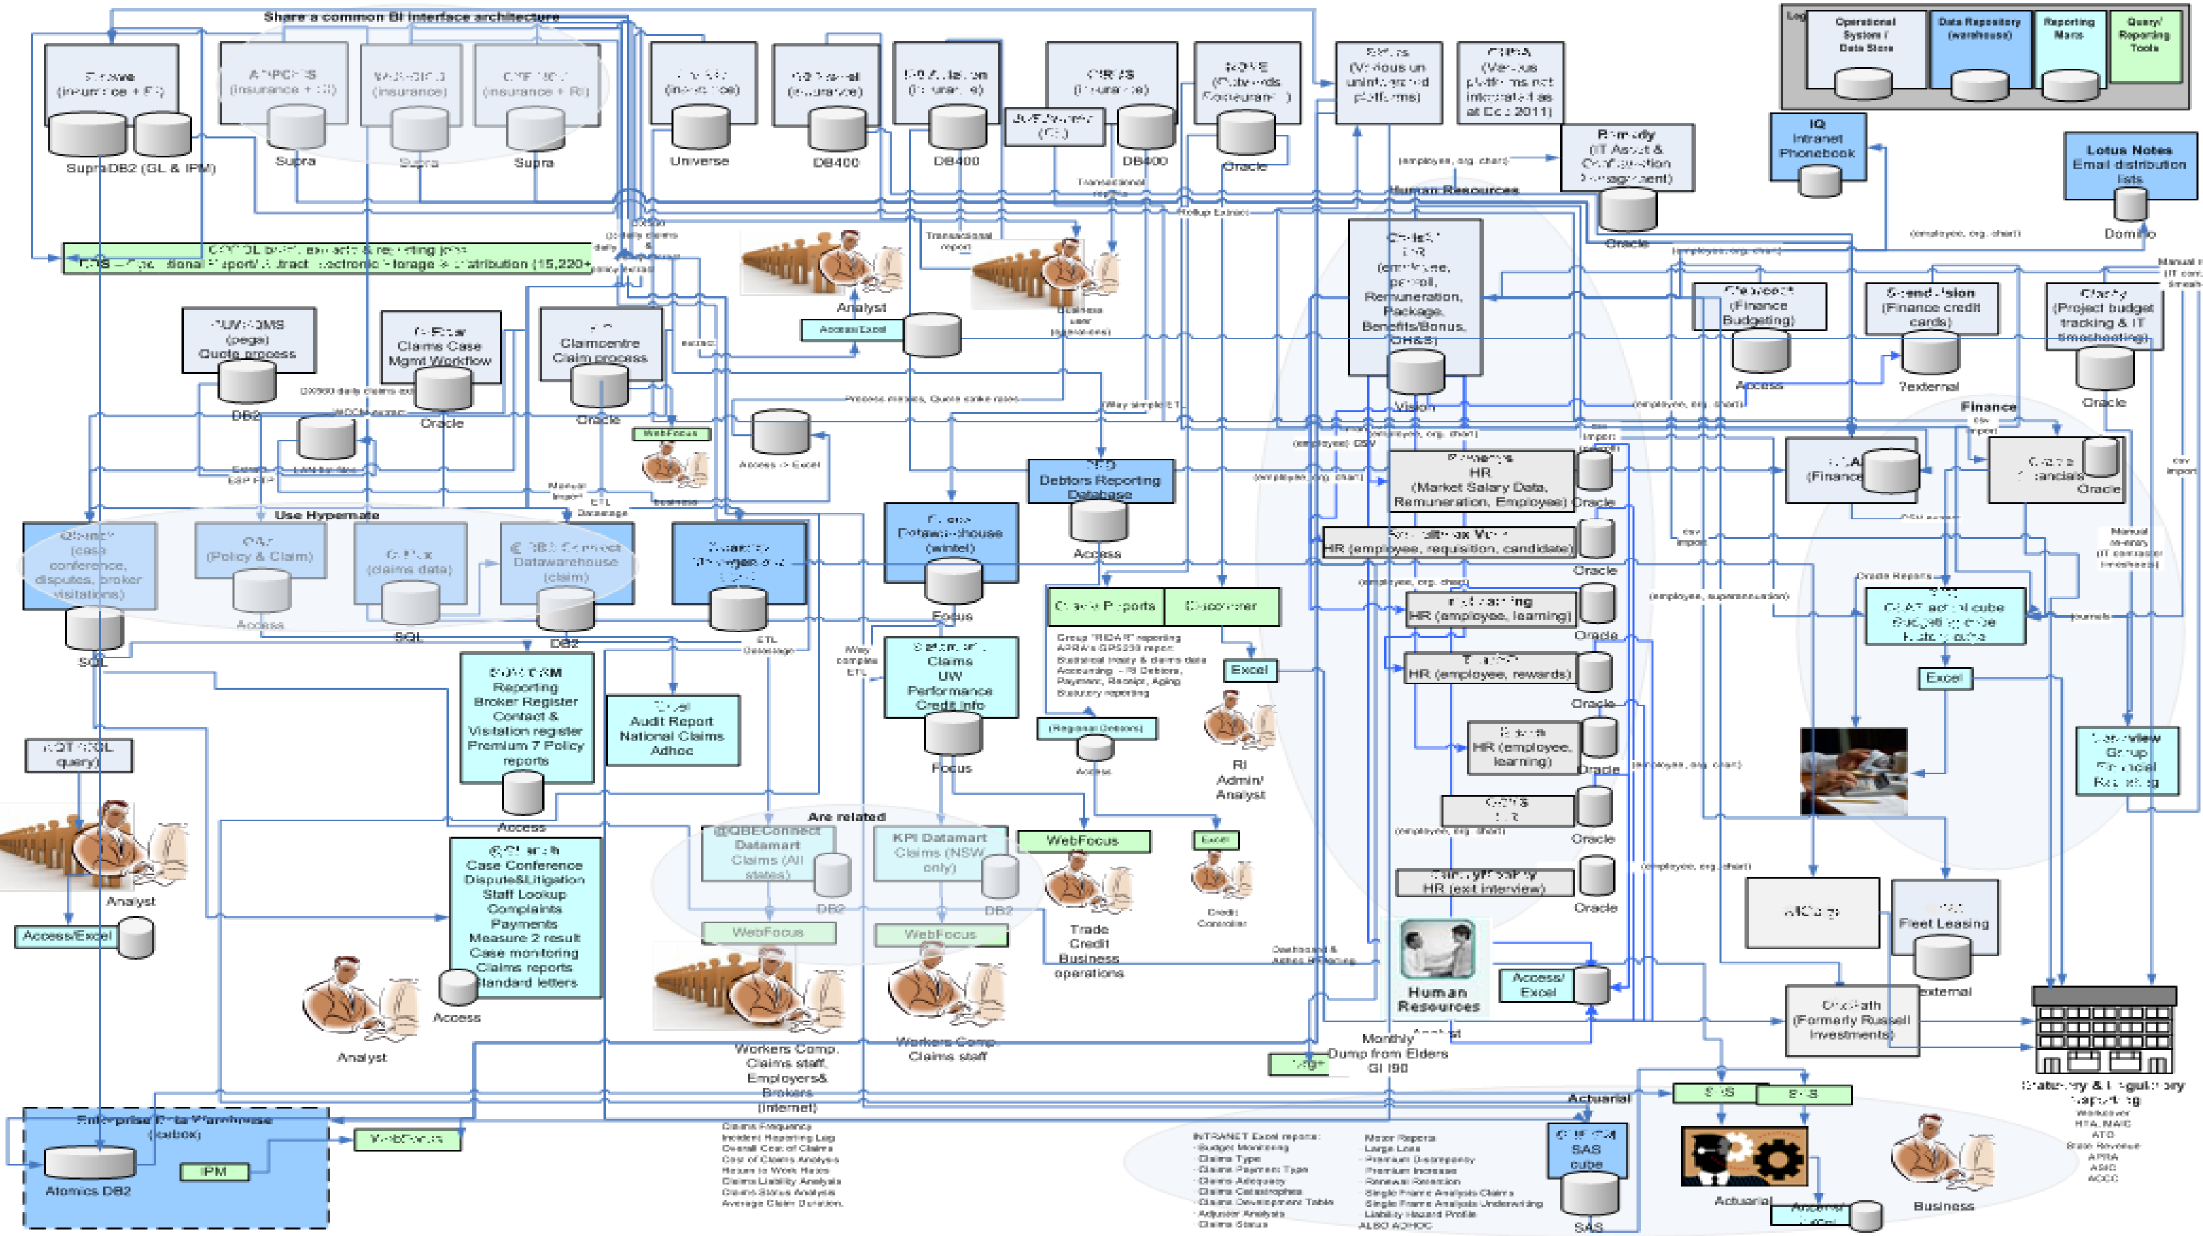

Share a common BI interface architecture
Operational System / Data Store
Data Repository (warehouse)
Reporting Marts
Query/ Reporting Tools
SupraDB2 (GL & IPM)
Supra
Universe
DB400
Oracle
IQ
Intranet
Phonebook
Lotus Notes
Email distribution lists
Domino
Human Resources
(employee, org. chart)
Analyst
Access/Excel
Claimcentre
Claim process
DB2
Access
Vision
Finance
Oracle
Use Hyperion
(Policy & Claim)
(claims data)
SQL
Datawarehouse
(claim)
Debtors Reporting Database
Focus
Oracle Reports
Discoverer
Excel
HR (employee, requisition, candidate)
HR (employee, learning)
HR (employee, rewards)
HR (exit interview)
(employee, superannuation)
Audit Report
National Claims
Adhoc
Reporting
Broker Register
Contact &
Visitation register
Premium 7 Policy
reports
Case Conference
Dispute&Litigation
Staff Lookup
Complaints
Payments
Measure 2 result
Case monitoring
Claims reports
Standard letters
Are related
@QBEConnect
Datamart
KPI Datamart
WebFocus
Trade
Credit
Business
operations
RI
Admin/
Analyst
Credit
Controller
Human
Resources
Monthly
Dump from Elders
GI I90
Fleet Leasing
external
(Formerly Russell
Investments)
Actuarial
SAS
cube
Business
Enterprise Data Warehouse
Atomics DB2
IPM
Workers Comp.
Claims staff,
Employers&
Brokers
(internet)
Workers Comp.
Claims staff
Claims Frequency
Incident Reporting Log
Overall Cost of Claims
Cost of Claims Analysis
Claims Liability Analysis
Claims Status Analysis
Average Claim Duration
INTRANET Excel reports:
Budget Monitoring
Claims Type
Claims Adequacy
Claims Catastrophes
Claims Development Table
Adjuster Analysis
Claims Status
Large Loss
Premium Increase
Renewal Retention
Single Frame Analysis Claims
Single Frame Analysis Underwriting
Liability Hazard Profile
ALSO ADHOC
APRA
ASIC
ATO
ACCC

Delivering More with Less – DB2

- zSeries - DB2 for Z/OS
  - Large Binary Objects
  - XML
  - Compression
  - Enhanced SQL Functionality (MQT, Star Joins)
- pSeries - DB2 for LUW
  - Virtualised Power Processors (P6, P7 and Blade Centre)
  - HADR – High Availability
- xSeries – DB2 for LUW
  - Intel Processors (xSeries)
  - Data Partitioning Feature
  - Active/Active High Availability Cluster
- iSeries

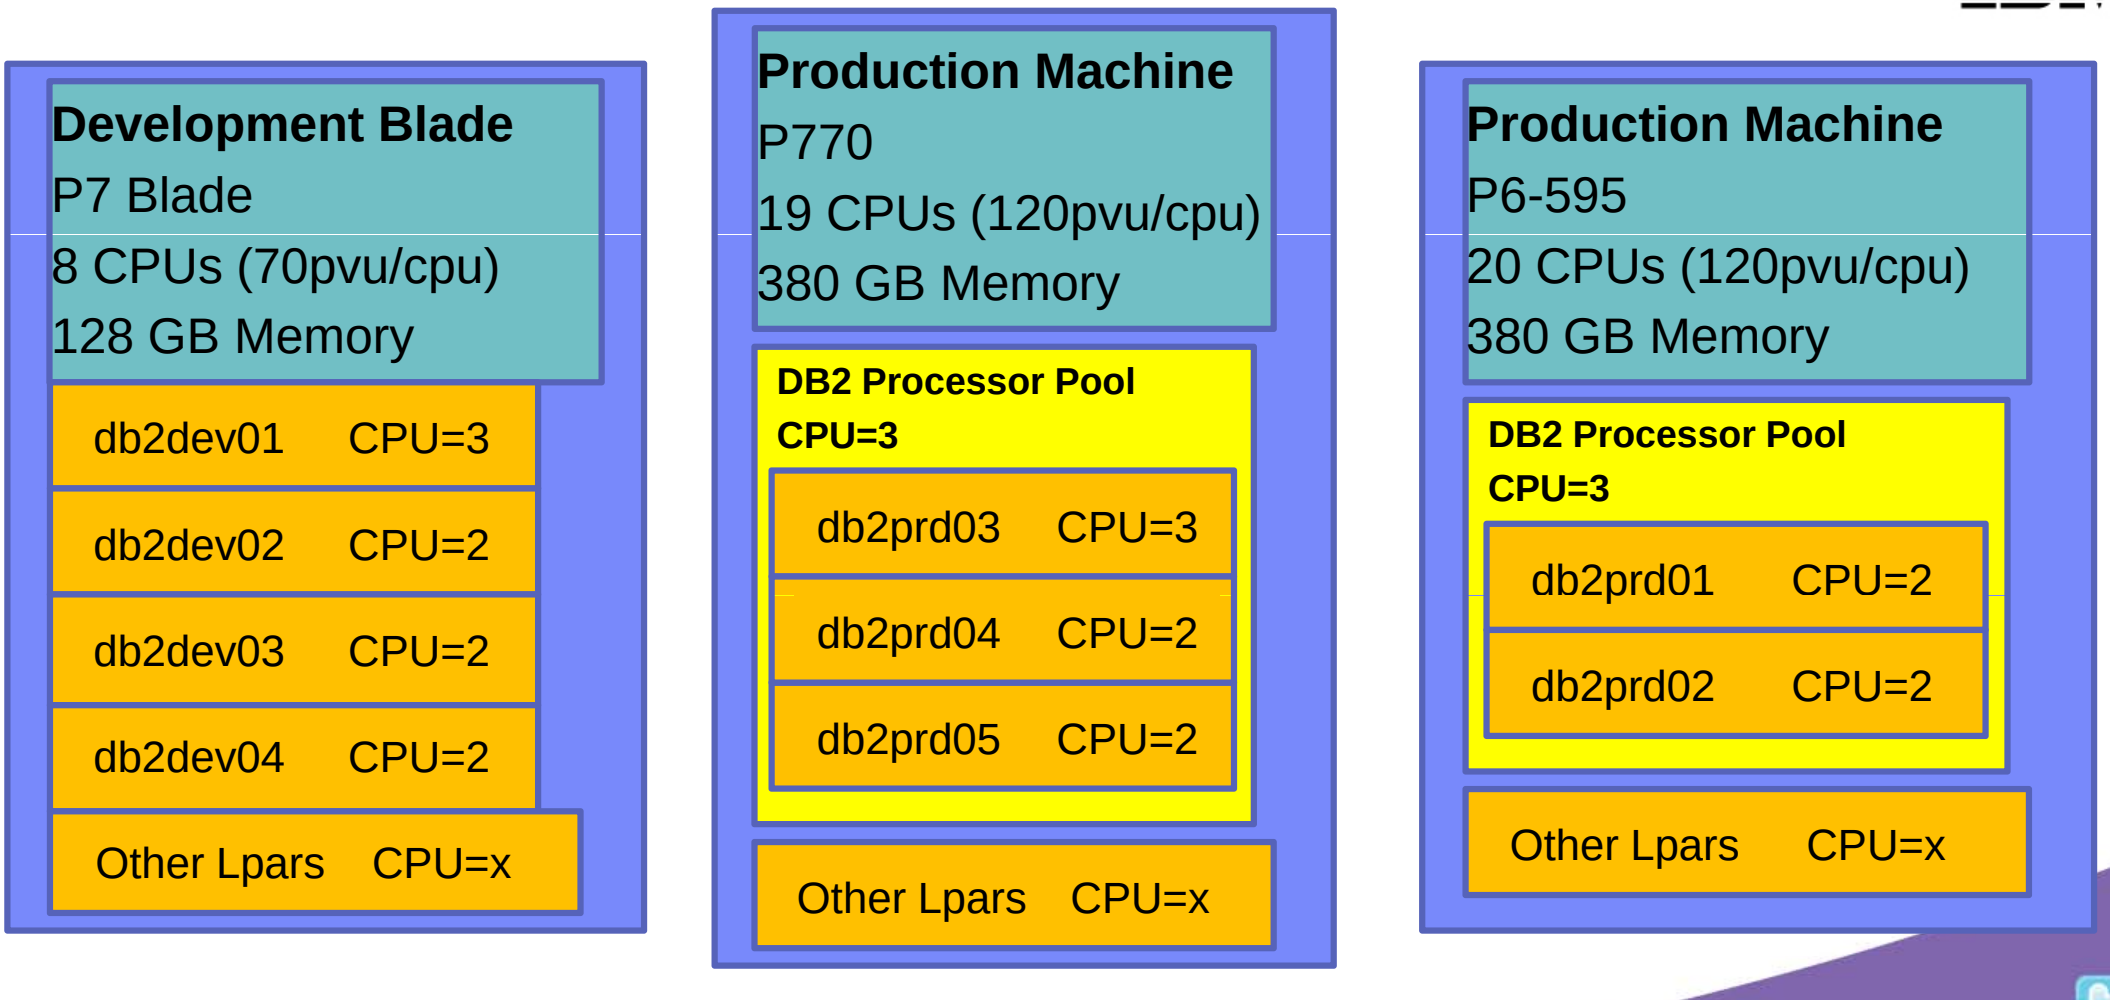

QBE – V-MAX Storage Topology

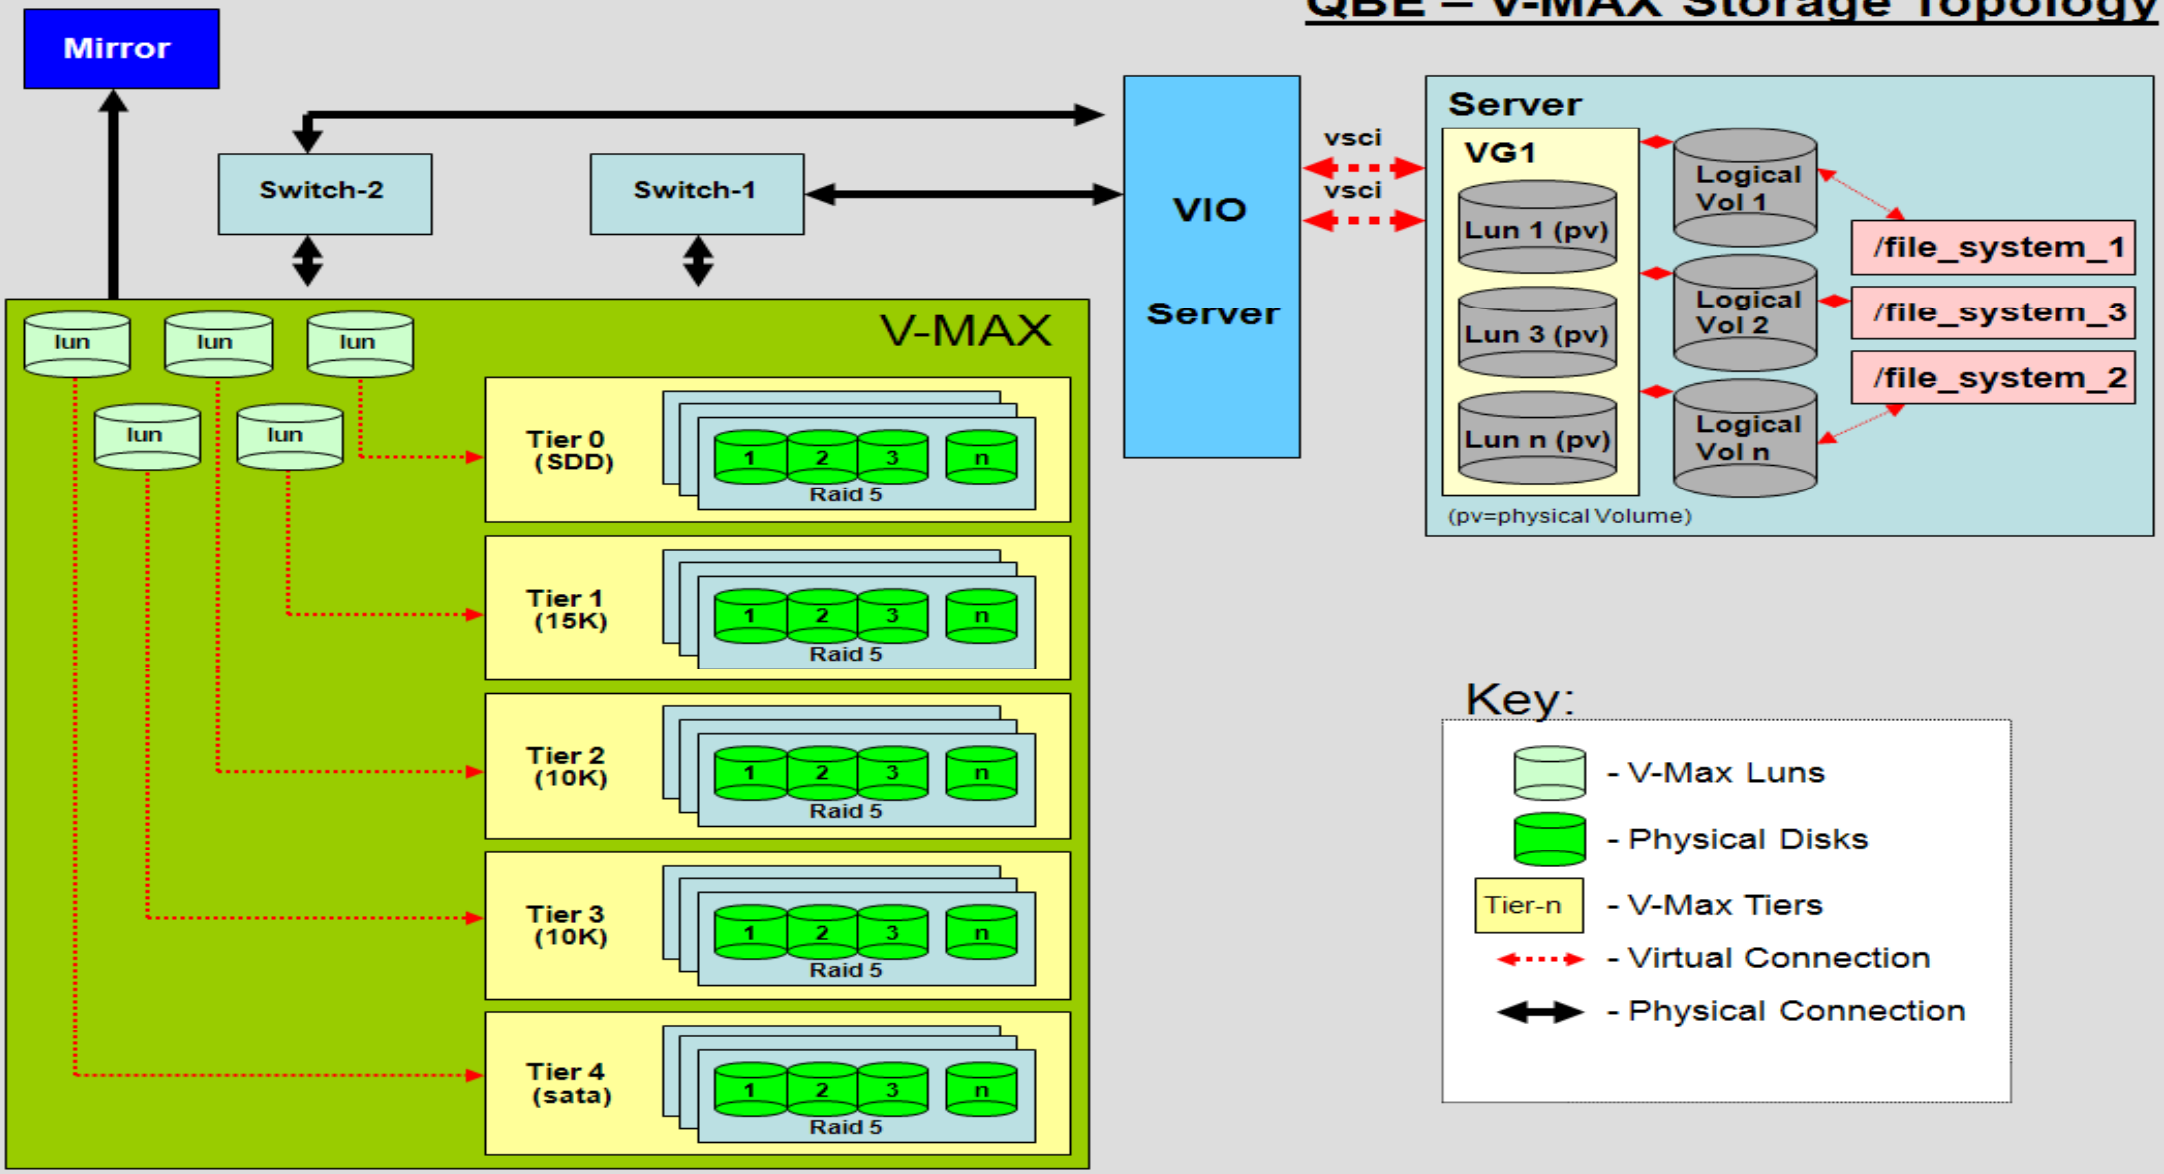

OLTP Consolidation – $oftware $avings...

- Enterprise Software Agreement
  - Virtualisation of CPUs Allows Leveraging of Pricing
  - Enterprise Server Edition to Advanced Enterprise Server Edition
    - •Workload Manager
    - •Compression

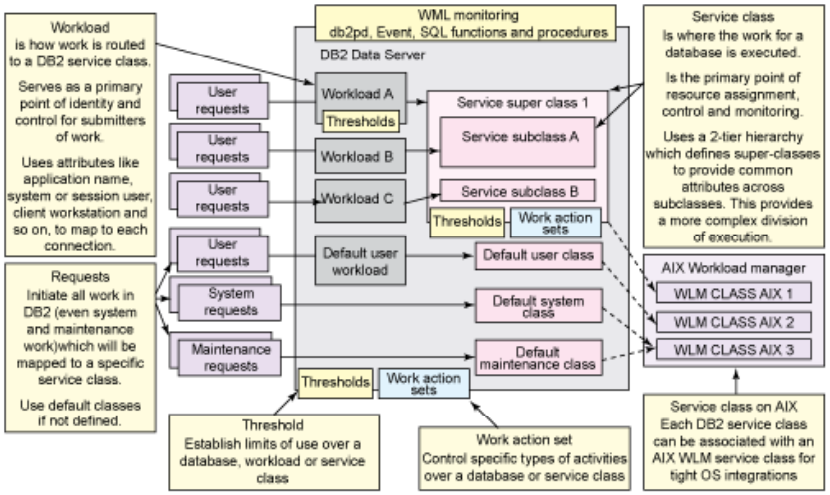

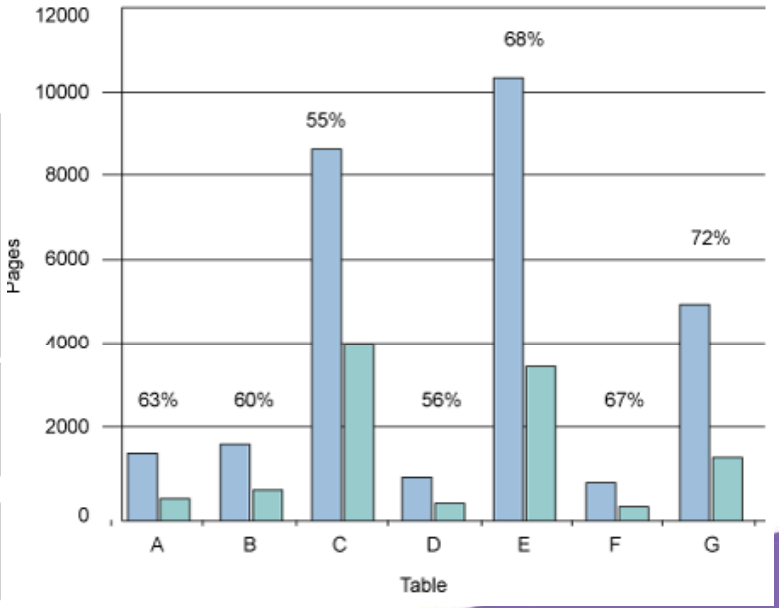

OLTP – Benefits

- Leveraged DB2 Product Licenses via AIX Virtualisation and Processor Pools
- DB2 Compression
  - Reduction of Disk Space
  - Reduction of Backup Times
  - Increased Transaction Throughput
  - More Effective Use of Memory
  - Reduction in Maintenance Window

- Utilising *WebSphere MQ* and *Message Broker* – Enabled Seamless Application Wide Connectivity
- Enables Business Interaction to Systems in a Cost Effective Manner

Does it Work Well Enough?

QBE wins brokers' award Number 10

19 September 2011

QBE has won the National Insurance Brokers Association (NIBA) General Insurer of the Year award for the 10th consecutive year.

The keenly contested award was made at the NIBA Convention opening ceremony in Sydney yesterday.

NIBA
National Insurance Brokers Association

QBE wins top insurer for 10th year in a row

The West Australian
September 19, 2011, 7:28 am

Recommend Tweet Email Print

The West Australian © QBE wins top insurer for 10th year in a row

QBE Insurance has been named top insurer by the National Insurance Brokers Association for a record 10th consecutive year.

The company was voted best general insurer by NIBA Qualified Practising Insurance Brokers.

NIBA president David Duffield said brokers were asked to nominate the general insurer they considered performed best overall on the 10 most important service and product factors.

"Only national insurers operating in the broker market were eligible for inclusion in the award," he said.

"The ratings were analysed independently with the winner being the general insurer which received the highest number of nominations.

Warehouse Management with IBM

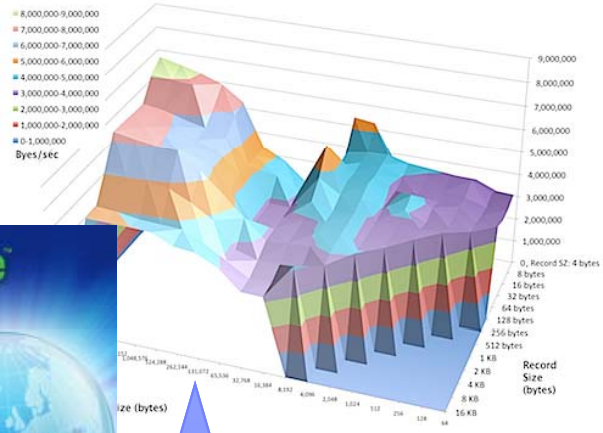

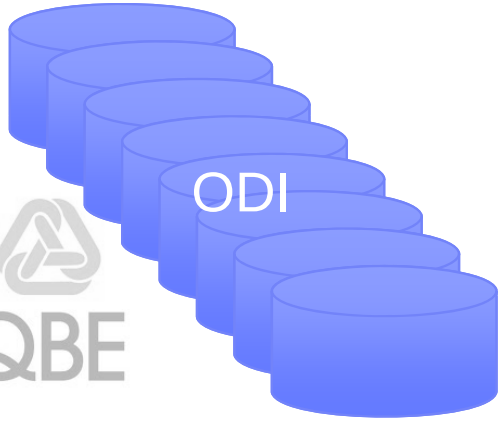

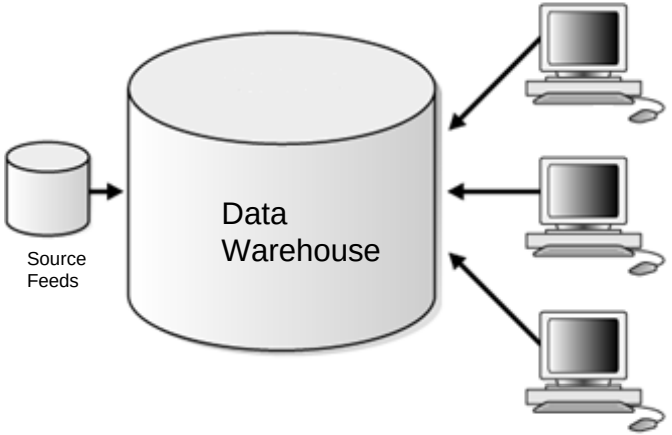

Star
Joins

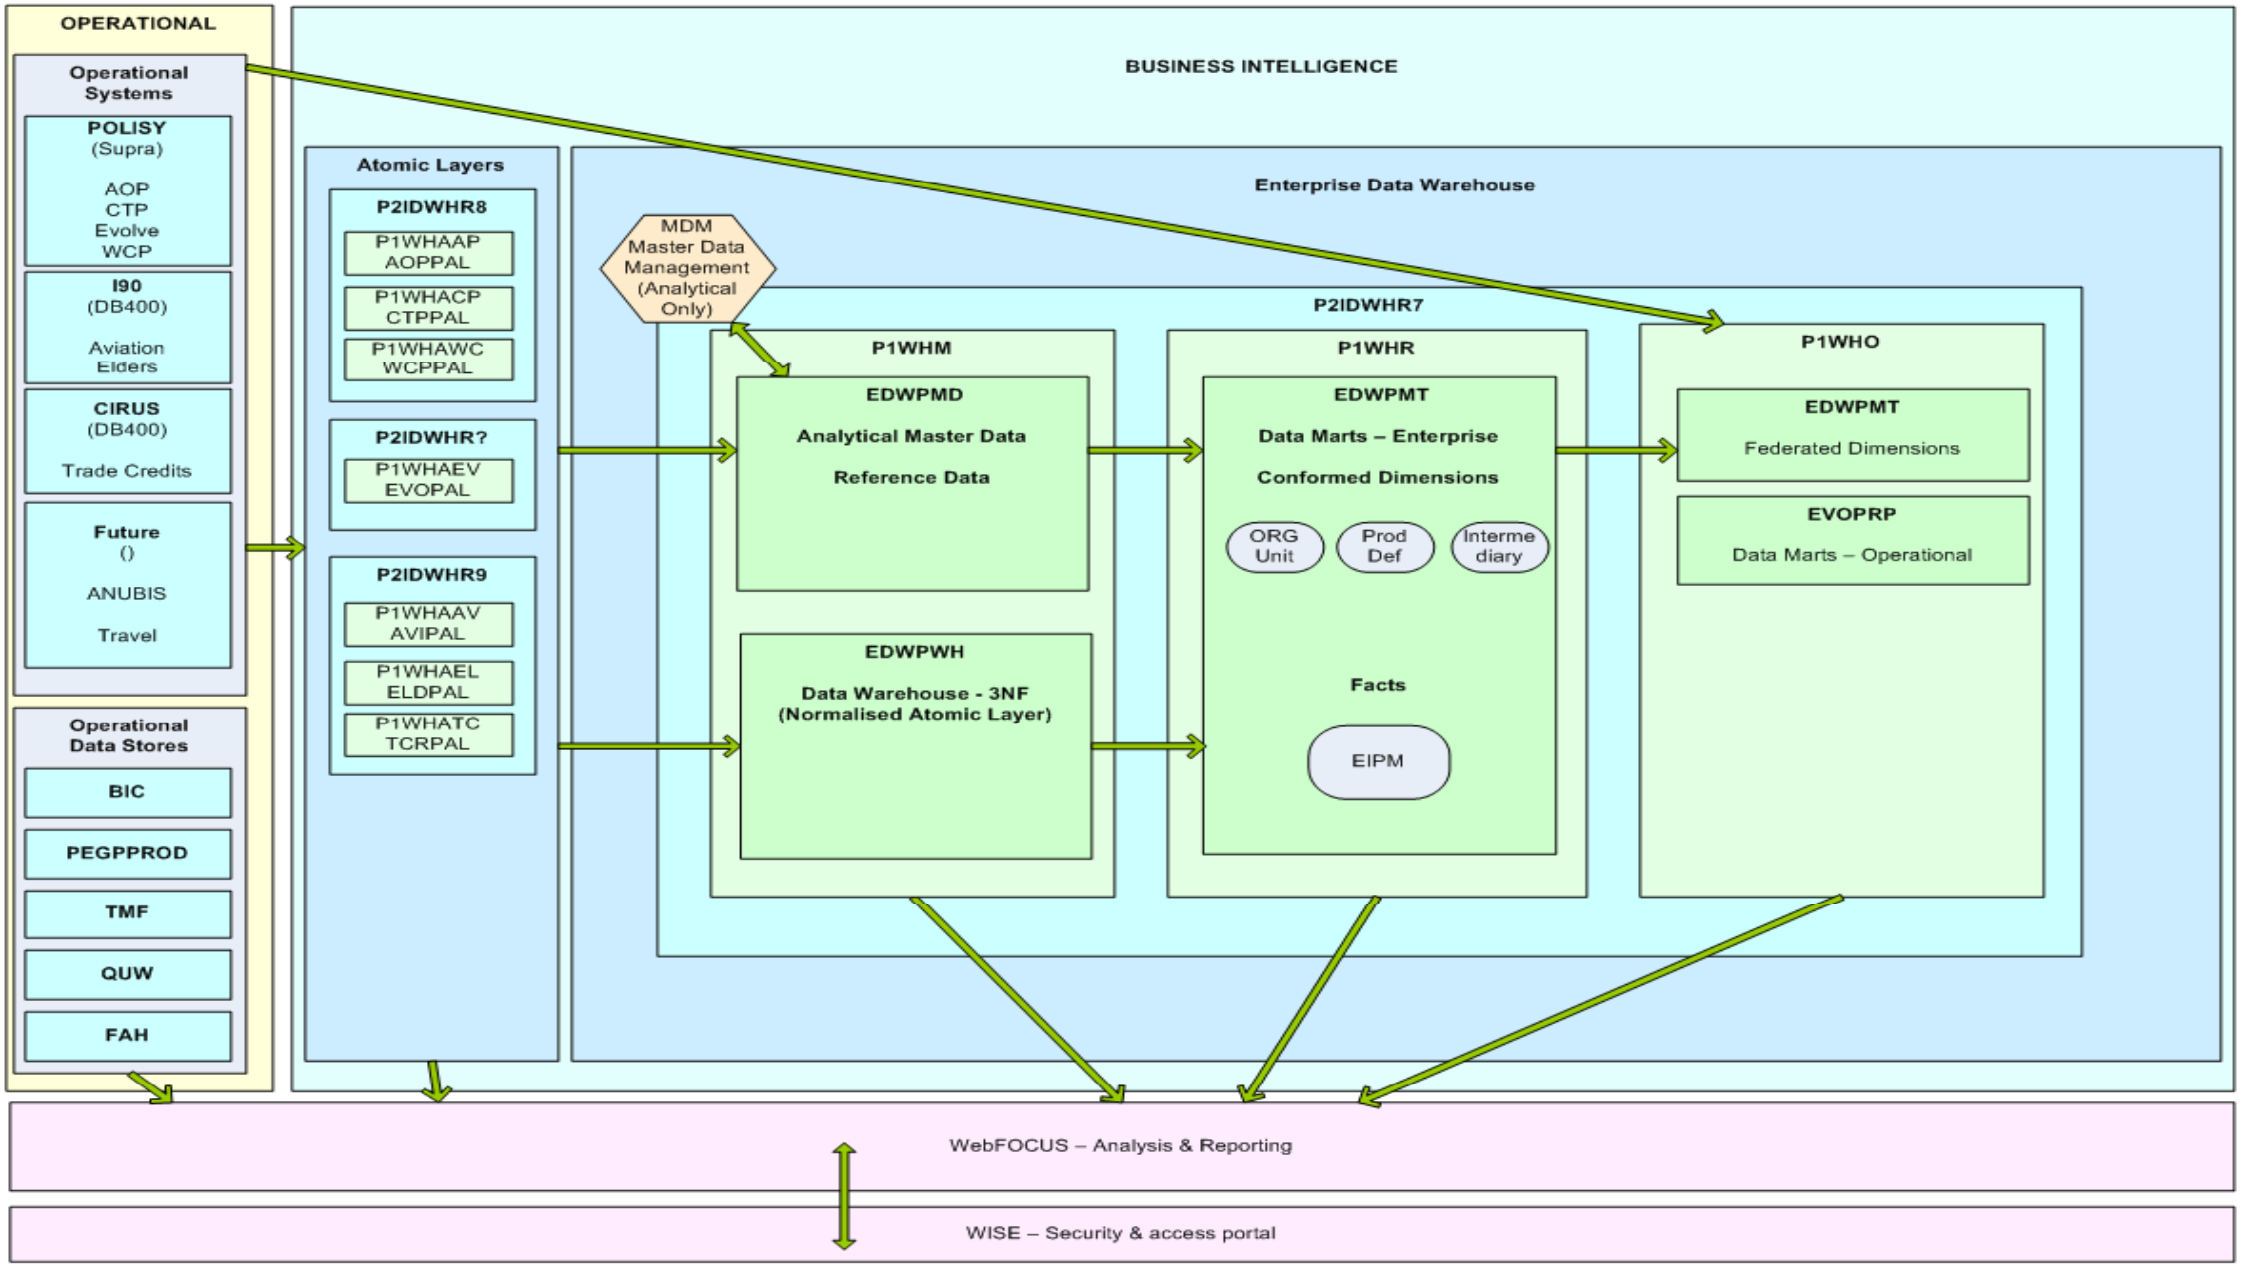
OPERATIONAL
Operational Systems
POLISY
(Supra)
AOP
CTP
Evolve
WCP
I90
(DB400)
Aviation
Elders
CIRUS
(DB400)
Trade Credits
Future
()
ANUBIS
Travel
Operational Data Stores
BIC
PEGPPROD
TMF
QUW
FAH
BUSINESS INTELLIGENCE
Atomic Layers
P2IDWHR8
P1WHAAP
AOPPAL
P1WHACP
CTPPAL
P1WHAWC
WCPPAL
P2IDWHR?
P1WHAEV
EVOPAL
P2IDWHR9
P1WHAAV
AVIPAL
P1WHAEL
ELDPAL
P1WHATC
TCRPAL
Enterprise Data Warehouse
MDM
Master Data
Management
(Analytical
Only)
P2IDWHR7
P1WHM
EDWPMD
Analytical Master Data
Reference Data
EDWPWH
Data Warehouse - 3NF
(Normalised Atomic Layer)
P1WHR
EDWPMT
Data Marts – Enterprise
Conformed Dimensions
ORG
Unit
Prod
Def
Interme
diary
Facts
EIPM
P1WHO
EDWPMT
Federated Dimensions
EVOPRP
Data Marts – Operational
WebFOCUS – Analysis & Reporting
WISE – Security & access portal

Leveraging DB2 Technology – zSeries

- SUPRA General Ledger in OLTP
  - Actual & Budgeted Figures
  - Viewed by Location, Branch, Accounts and Class of Business
  - 25 GB Data
  - Batch Process 18 Hours to Run – Inflexible and not Dynamic
- DB2 Solution Objectives Were to:
  - Reduce Size of Database
  - Remove the “18 Hour” Batch Process
  - Allow Dynamic Structure Changes
  - OLTP Interface to be as Quick as SUPRA (pre calculated!)
  - A Reasonably Big Ask!!!

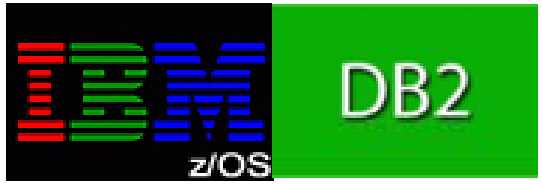

Leveraging DB2 Technology – zSeries

- Utilised DB2 Z/OS
  - Remodelled Table Structure
  - Utilised DB2 Z/OS Star Join Technology
  - Included Hardware Compression on the ZSeries Server

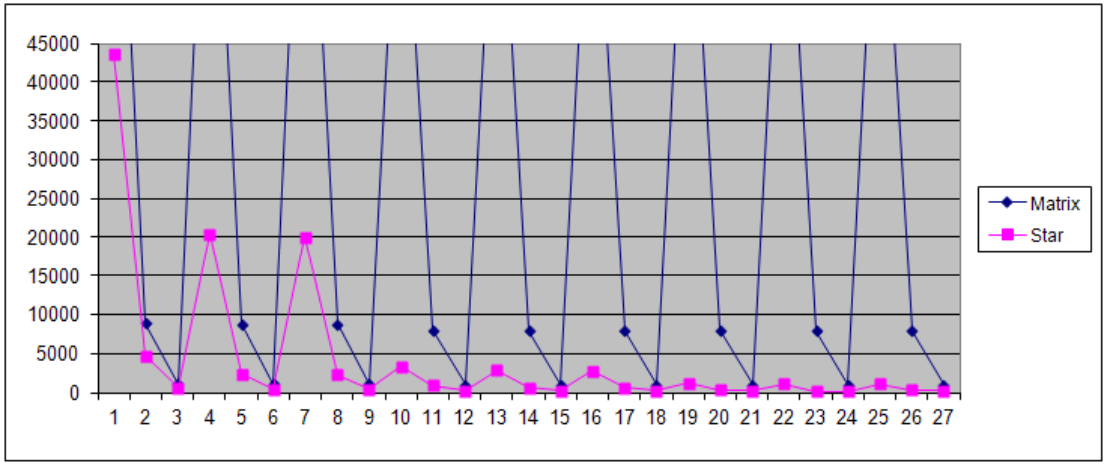

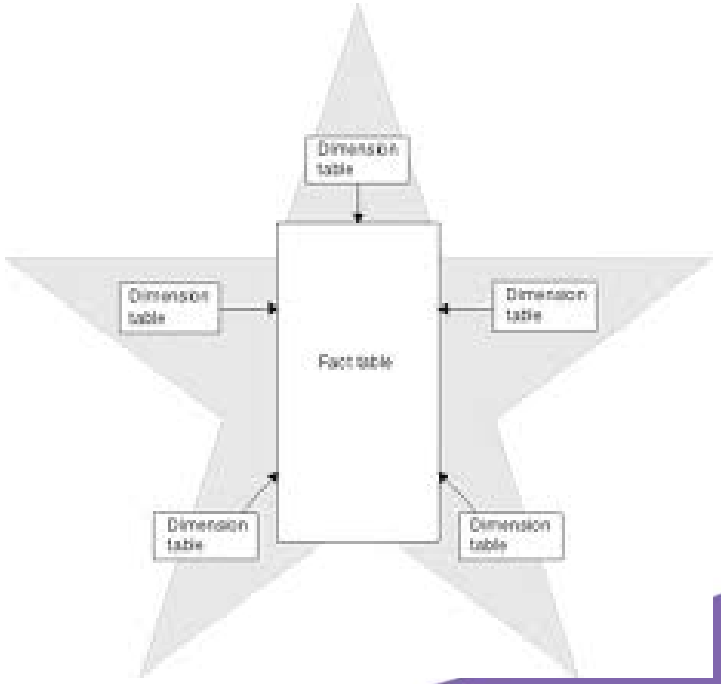

How it Performed:

Elapsed

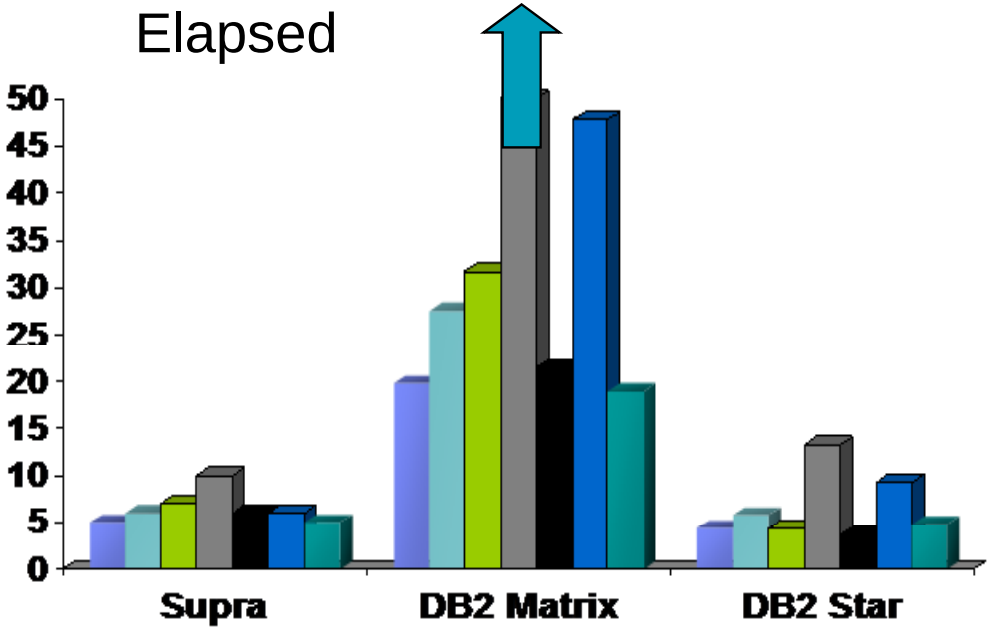

Getpages

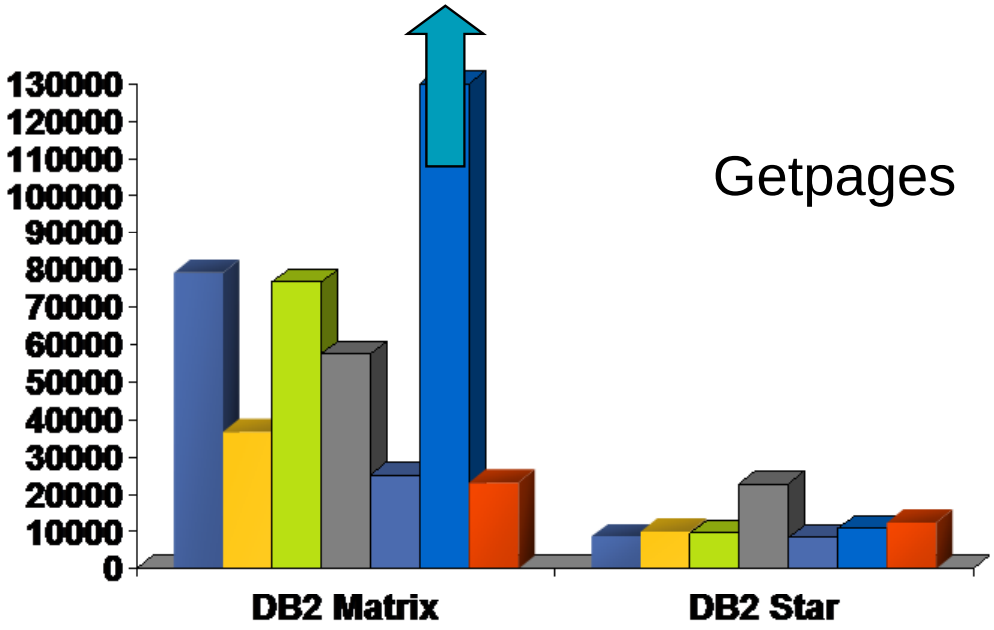

Final Outcome

- Database Size – from ~25 GB to ~2 GB
- Online Structure Updates (No long batch processing)
- Very flexible – Multiple Year Views (not fixed)
- In Most Areas it Outperformed the SUPRA Legacy Solution

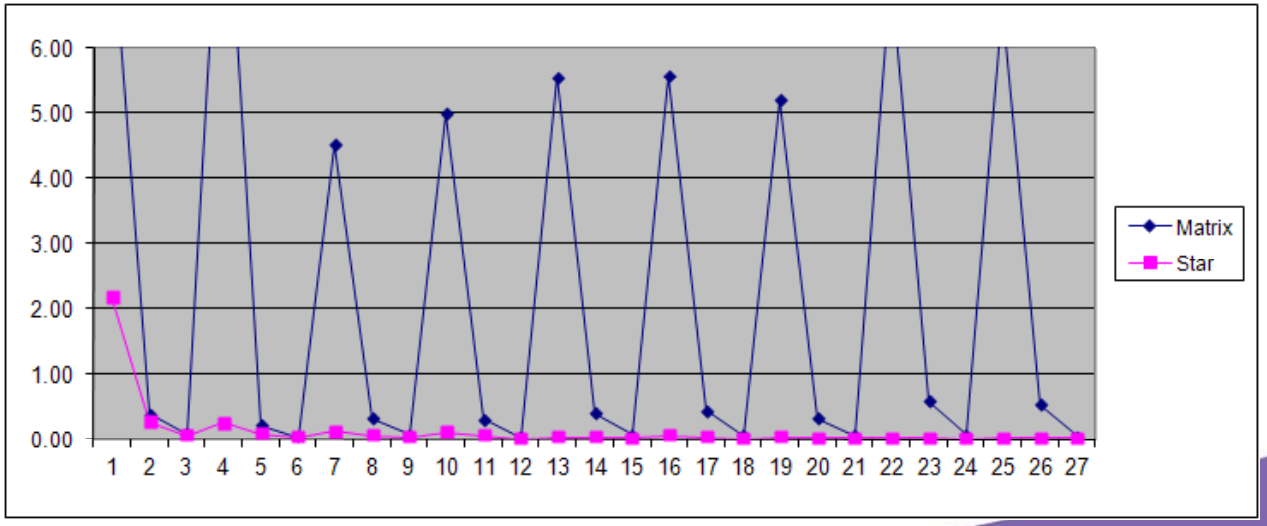

Leveraging DB2 Technology – xSeries

- **Build Ourselves" Option was a Non Starter**
  - High Profile Project : "Someone Else Can Build it"
  - SAN Wasn't Set Up for Dedicated Operation (Big Bucks…)
- **BCU for AIX Running on pSeries**
  - Required 2 x 4 Way BCUs
  - Less Flexible for Our Technology Fit
  - SLA Didn't Demand Higher Availability of pSeries
- **Our Choice : BCU for Linux Running on xSeries**
  - Less Expensive Option (Came in Well Below pSeries BCU)
  - Came in Considerably Below Our Own Build

DB2 on xSeries and DPF -

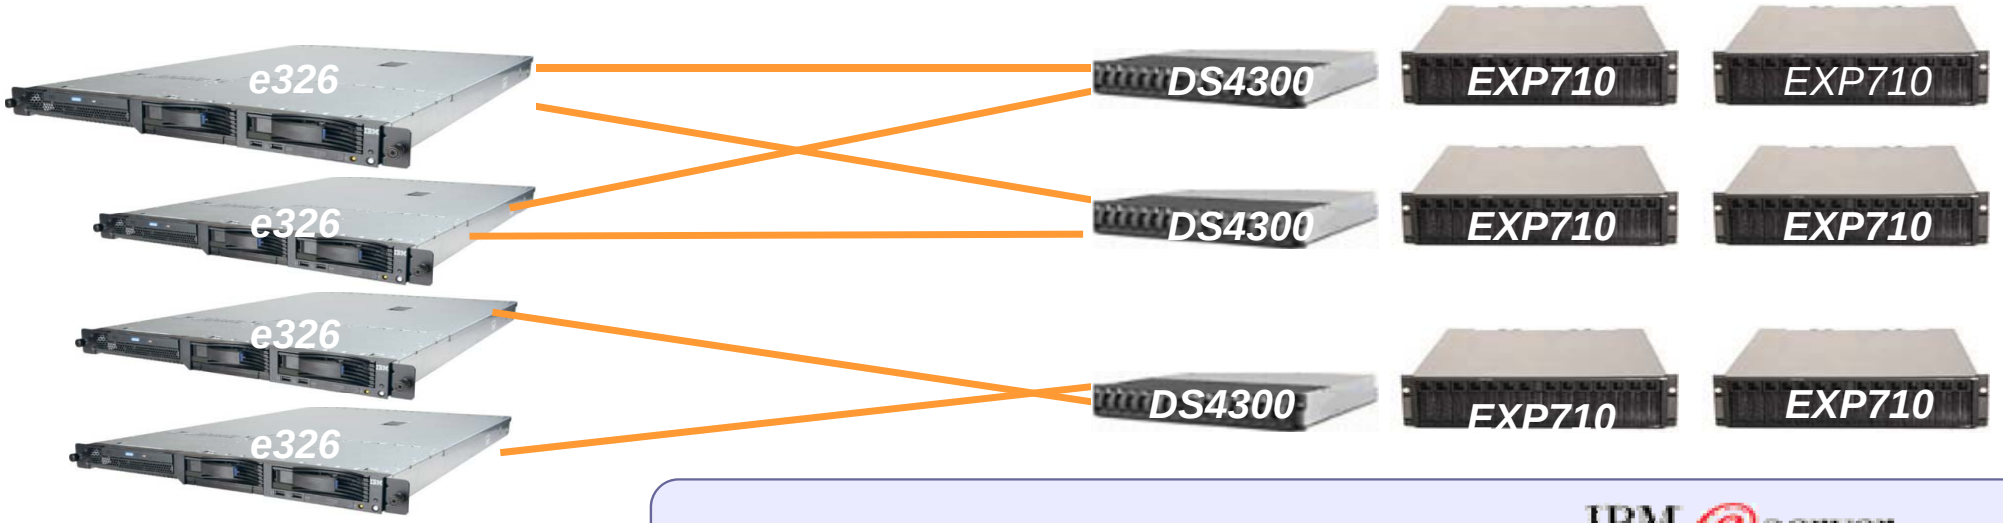

IBM eserver

Processors	2 x 2.4GHz dual core AMD Opteron
Memory	2 x 2GB
Disk Controller	1 x Qlogic 2-Gbps 4-port Fibre Channel to PCI-X HBA
Disk Drives	2 x 73.4GB, 10K RPM, Ultra320 SCSI Hot-swap HDD

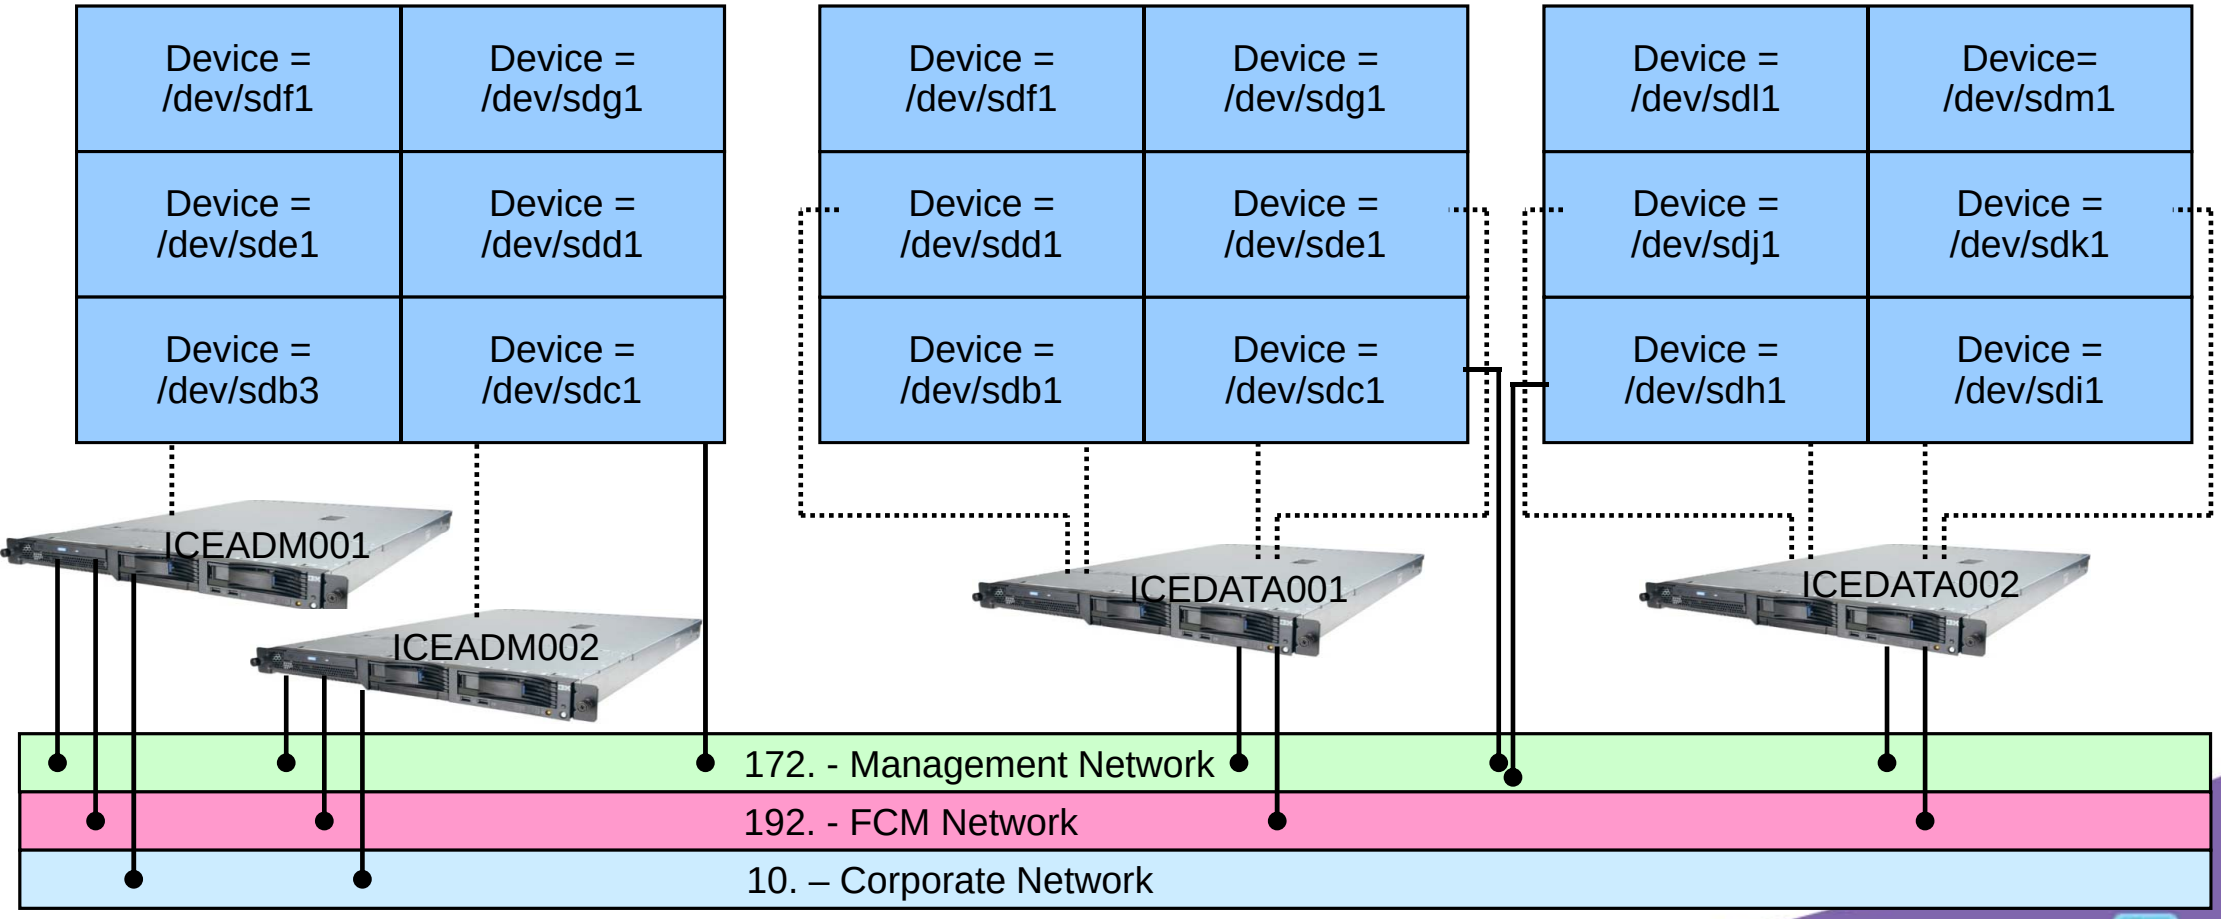

Comparative Query Performance

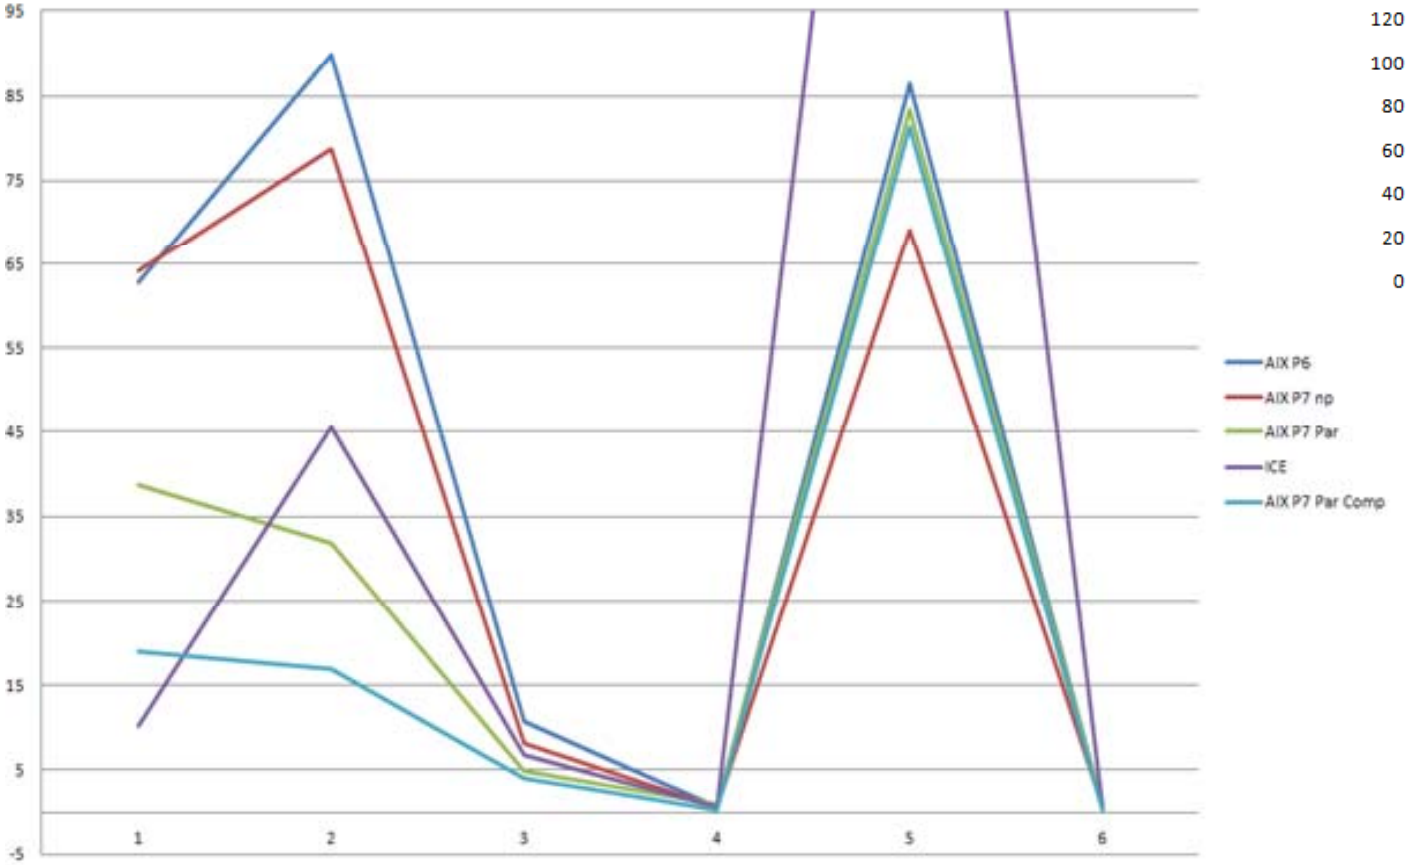

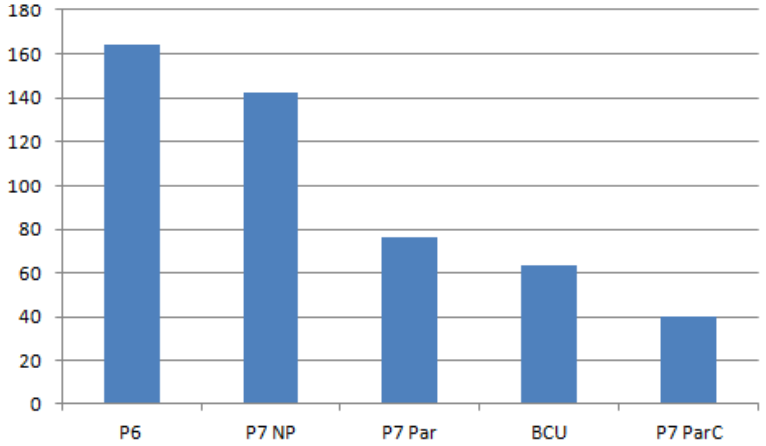

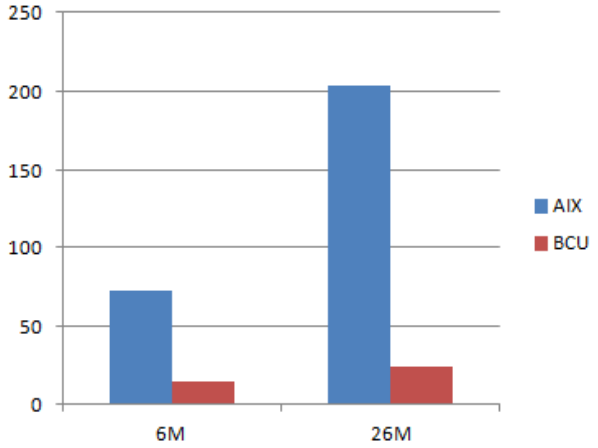

Disk I/O Performance - Gotcha

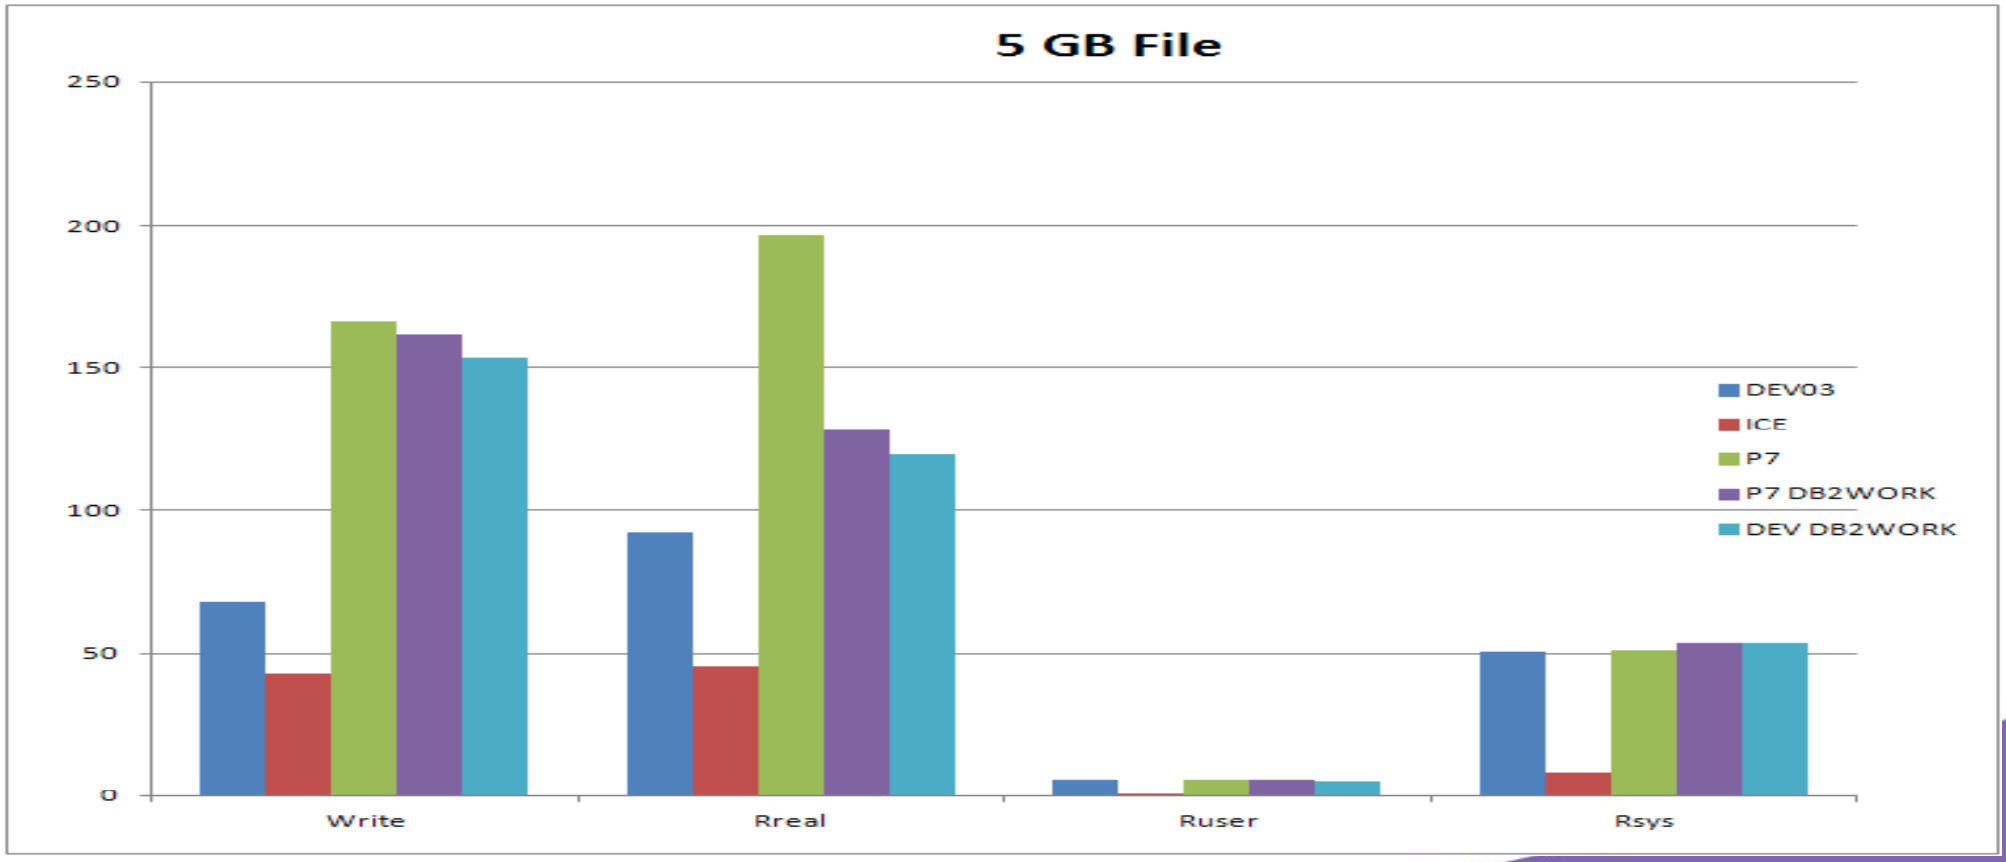

BCU Benefits

- Applied Both Horses for Courses and Bang for Bucks Principles
- No Virtualisation of Storage or CPU
- Highly Available Active/Active Cluster
- Scalable
- Exceptionally Cost Effective Utilising Commodity Hardware
- IBM Came, Installed it and it SIMPLY WORKED
- Fantastic Performance Straight Away

OLAP Consolidation – Terabyte Pricing

- Traditional PVU Pricing Does Not Leverage Business Value in the Warehouse Space
- Pricing based on Compressed Warehouse Data (excluding indexes!)
- Price Mapped to Data not the Hardware
- Allows Greater Predictability
- Simple to Understand and Transparent
- InfoSphere Warehouse
  - Enterprise Server Edition
  - Workload Manager
  - Compression
  - DPF
- A Smarter Way of Using Technology

OLAP Consolidation – Terabyte Pricing

- IBM x3850 M5
  - 32 Core 2.13GHz processors
  - 256GB Memory (can be expandable to 2TB)
  - 1.2TB Mirrored High IOPS Bus Adapters
- IBM DS3500 Storage
  - DS3524 plus EXP Shelf
  - 48 x 300GB 2.5in 15K 6Gb SAS
- Leveraging DB2 Technology with DB2 TB Pricing

DB2 High Availability

- HADR
  - HADR Setup Straight Forward
  - Cost Effective Approach to High Availability
- Log Shipping
  - How we Moved a 100GB Production Database to a Separate Server and Upgraded it in 10 minutes
  - Easy to Setup and Easy to Manage

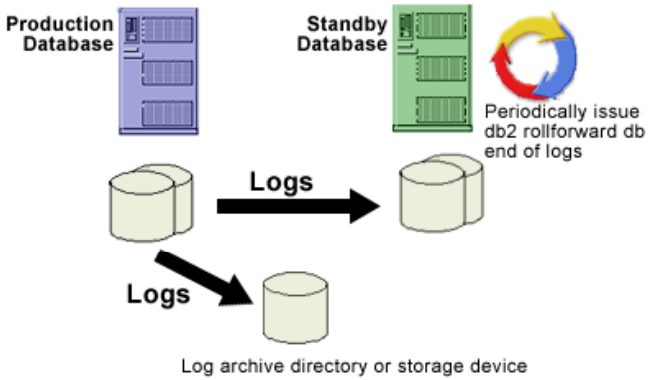

Conclusions – Administration and Management

- DB2:
  - 260 Databases on LUW
  - Two DB2 Z/OS Subsystems
  - 7TB of Data
  - 3 Full Time
- Oracle:
  - 60 Databases
  - 26TB of Data (22TB of which is development data)
  - 4 Full Time

Conclusions

- DB2 Provides Cost Effective Data Management
- It is Stable, Flexible and Robust
- Multiple Flavours from Rock Solid Z/OS to Commodity xSeries
- Easy to Manage – Patching is a Breeze with Traditionally Less Vulnerabilities Than Oracle
- Simple Licensing Model
- Tuning is a breeze – no reliance on code change (MQT's, Access Paths, etc)
- Native Integration and Management of XML and SOAP
- Role Based Security and Easy Integration to LDAP
- Multipath Install and Independent of Operating System

Questions???

Randall Ibbott
Technical Lead Database Administration
IT Services
QBE Australia
randall.ibbott@qbe.com

IBM INFORMATION INTEGRATION
& GOVERNANCE SYMPOSIUM 2012

Delivering Trusted Information for Smarter Business Decisions

twitter: *Follow* ***@ANZ_IM*** *or mention* ***#IIGS***