IBM System z Technology Summit

zEnterprise – The First System Of Systems

Data Serving And Business Analytics On A Single Platform

Dave Hayslett

25th May 2011

Data Plays A Key Role In Smarter Planet Solutions

1 Trillion connected intelligent devices

4 Billion mobile phone subscribers worldwide

2 Billion people on the web


Smart retail


Smart supply chains


Smart healthcare


Smart food systems


Smart countries

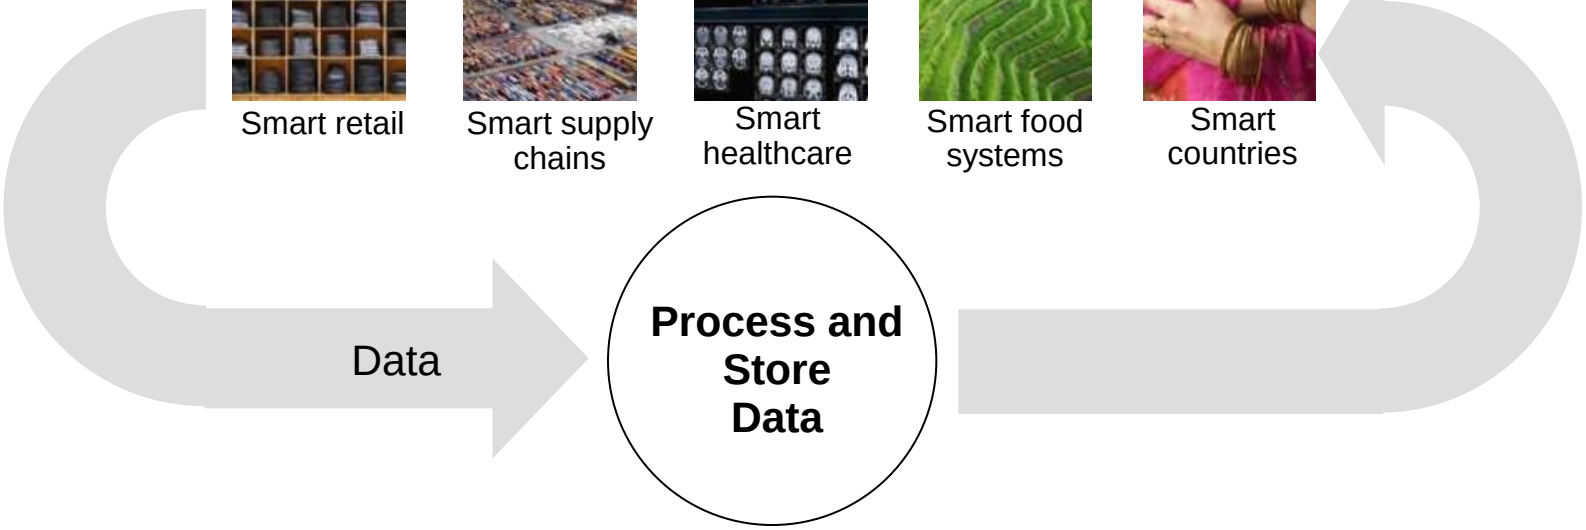

15 petabytes of new information generated daily
64 billion credit card transactions per year (up 35% YTY)
30 billion embedded RFID tags by 2010

Having A Reliable And Scalable Database Is Vital

- Databases help you store, organize, and access information in an efficient manner

- A good database should:
  - Scale as your data processing needs grow
  - Be highly available to ensure access to information
  - Protect the integrity of the data stored
  - Support new data types and access methods
  - Protect the security of the data
  - Compression data to save disk space
  - Be cost competitive

Let us explain why DB2 on z/OS is the best choice

A History Of Database Workload Optimization

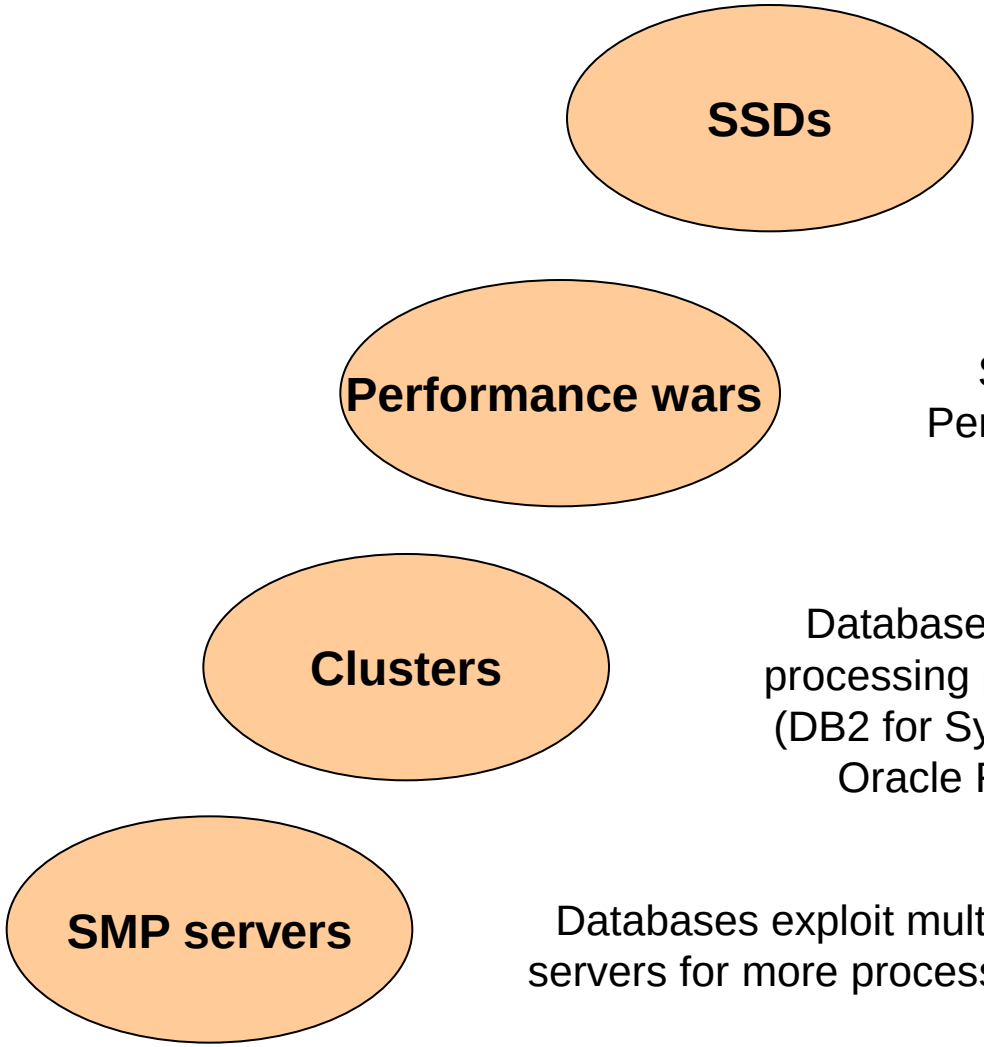

Dramatic Increase In Database Performance Requires Two Tricks

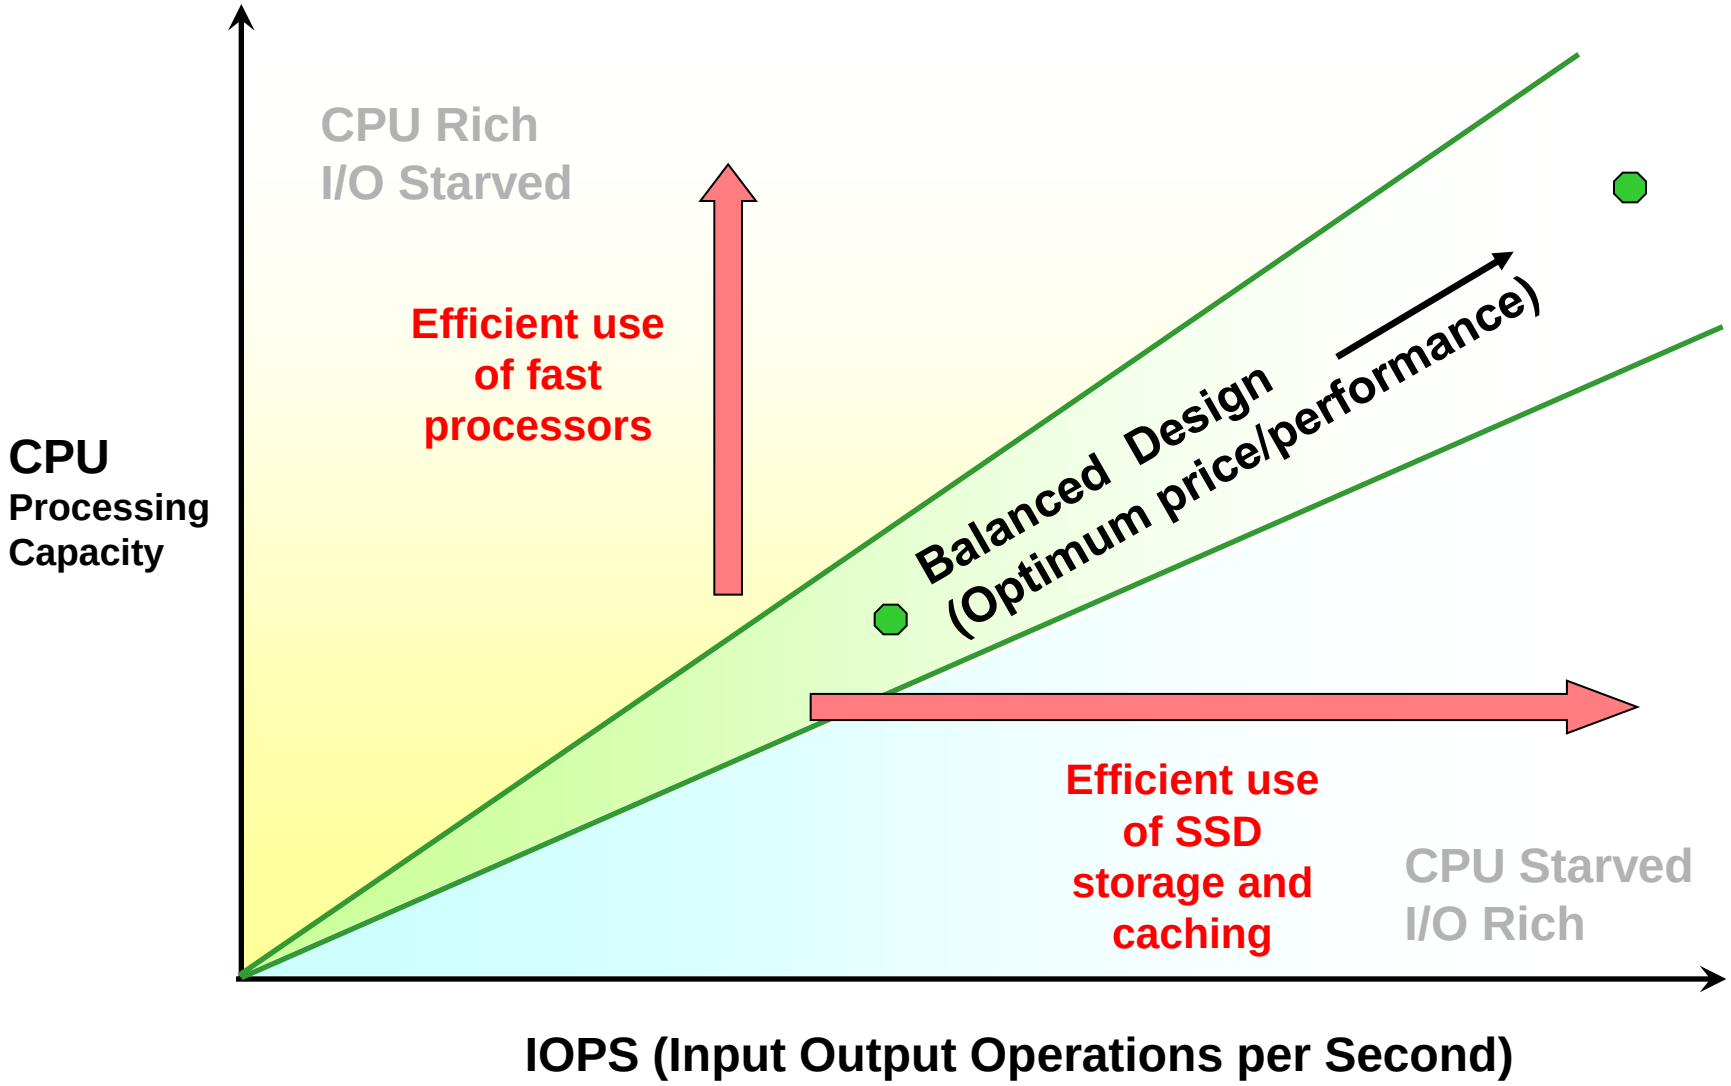

DB2 For z/OS Supports Parallel Sysplex Shared Data Clustering

- Shared data across nodes
- Hardware-based centralized lock and cache management
  - Provided by Coupling Facility
  - Supports near linear scalability
- Recovery from a node failure without a freeze
- Supports rolling upgrades with up to two different releases in a data sharing group

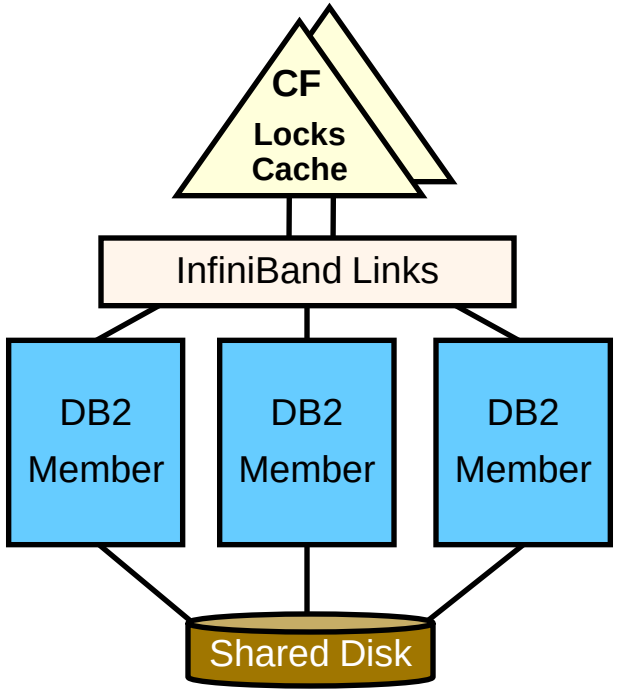

Now Available DB2 10 for z/OS:

Up to 20% faster performance
Hash access for faster OLTP
Automatic snapshots of changing data
Improvements in DB2 QMF and Tools suite
10x more concurrent users
More online administration
SQL and pureXML enhancements

DB2 For z/OS Grows Database Capabilities Significantly Beyond SMP Solutions

- Lets you **Add Capacity** as you need it
  - Each z196 is capable of executing 50 billion instructions per second and you can cluster up to 32 z196 systems
  - New members automatically process new transactions
  - Load balancing is automatic
- Provides **Continuous Availability**
  - Available during unplanned outages
    - Redistribute workload to surviving members automatically
    - In-flight transactions automatically rolled back in as little as 15 seconds
  - Available during planned outages
    - Remaining members handle workload as a member goes offline
    - Cluster operates continuously during rolling software updates

DB2 For z/OS Demonstrates Near Linear Scalability For Complex Banking Workload

Kookmin Bank
- **IBM System z9 and DB2**
- TCS BaNCS
- **15,353** Transactions/second
- **50** Million Accounts
- IBM benchmark for customer
- DB2 V9, CICS 3.1, z/OS V1.8

Bank of China [1]
- **IBM System z9 and DB2**
- **TCS BaNCS**
- **9,445**[2] Transactions/second
- **380** Million Accounts
- IBM benchmark for customer

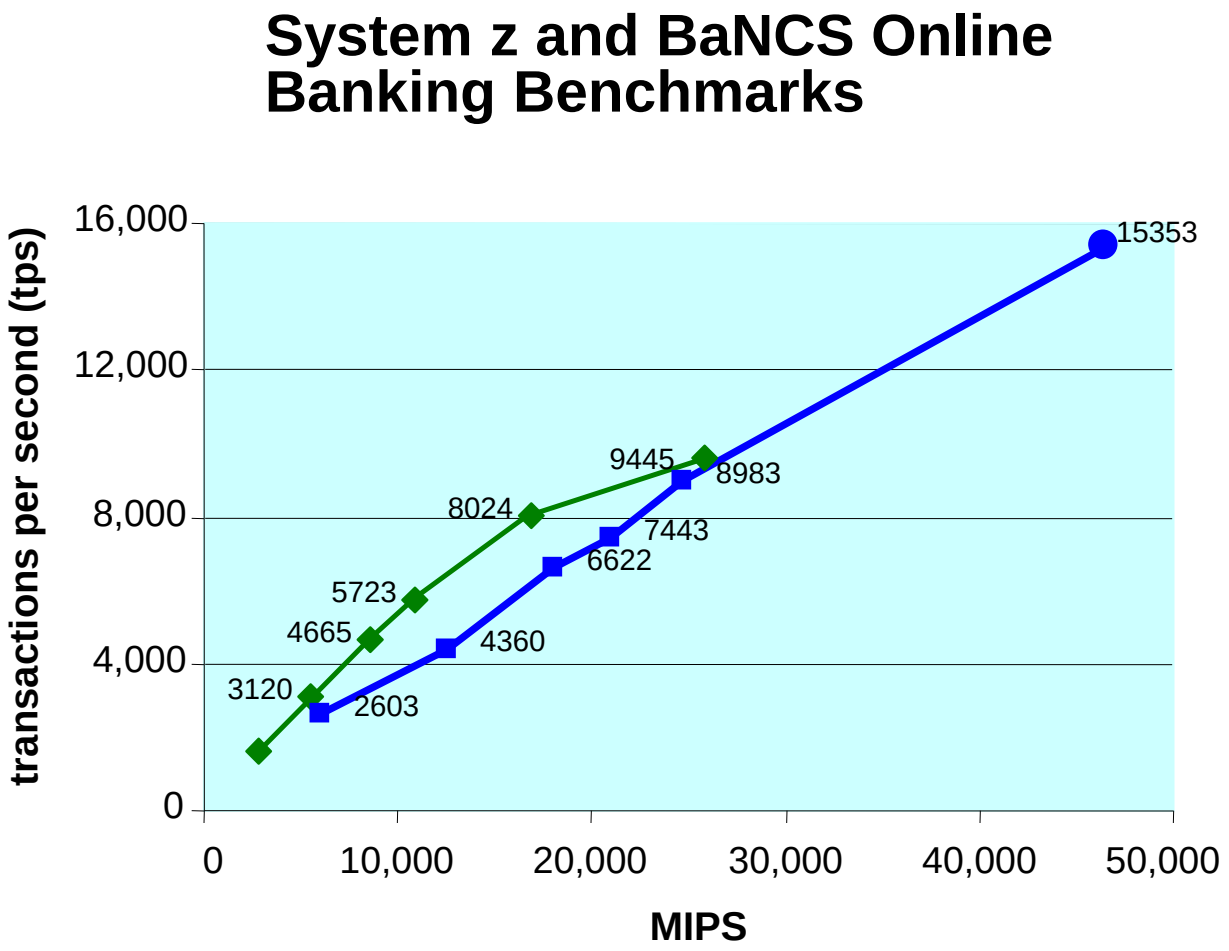

[1] **Source**:http://www.enterprisenetworksandservers.com/monthly/art.php?2976 and *InfoSizing FNS BANCS Scalability on IBM System z – Report Date: September 20, 2006*
2 Standard benchmark configuration reached 8,024 tps, a modified prototype reached 9,445 tps

Oracle Real Application Cluster (RAC) Supports A Distributed Network Clustering Architecture

- Shared disk across nodes
- Software-based distributed lock and cache management
- Poor scalability due to increasing interconnect traffic as cluster grows
- Upon node failure clusters can freeze during lock remastering process

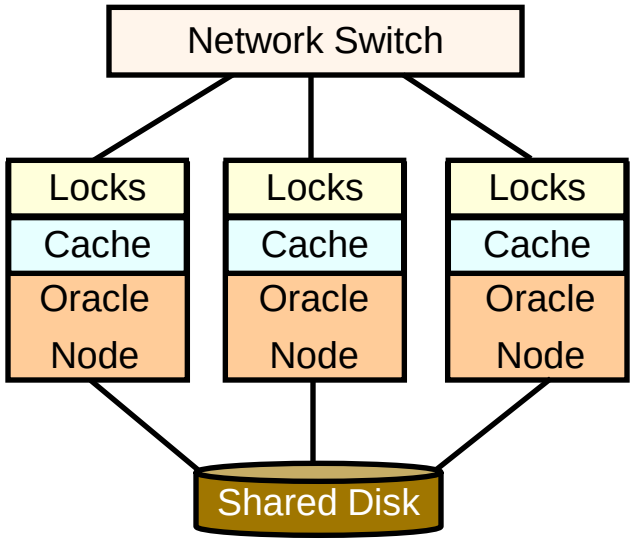

- Database software has to be of same release on all nodes
- No mirroring

There Are Expensive Consequences For Oracle RAC's Distributed Design

- Cluster members constantly share lock and cache data
  - In a 4-member cluster, obtaining a write lock could require 6 separate network communications
  - This is referred to as interconnect traffic
  - Interconnect traffic grows as the cluster grows and queries disperse across the cluster

Per-member performance declines as the cluster grows.

- During an unplanned outage Oracle RAC "freezes" I/O and lock requests
  - While re-mastering data blocks to surviving members
  - While locking pages that require recovery

Oracle RAC does not offer continuous availability.

Compare Techniques To Provide Sufficient I/O Bandwidth

- DB2 on System z
  - I/O sub-system hardware is dedicated for I/O processing
    - Up to 14 SAP processors
    - Up to 336 channel processors
  - Typically attached to a DS8000 class storage sub-system with disk caches and large I/O bandwidth
- Oracle RAC on distributed server
  - No dedicated I/O sub-system
    - I/O operations executed by general purpose processors
  - Typically attached to mid tier storage

What Equivalent DB2 For z/OS Performance Can You Get Out Of Oracle Running On Sun Servers?

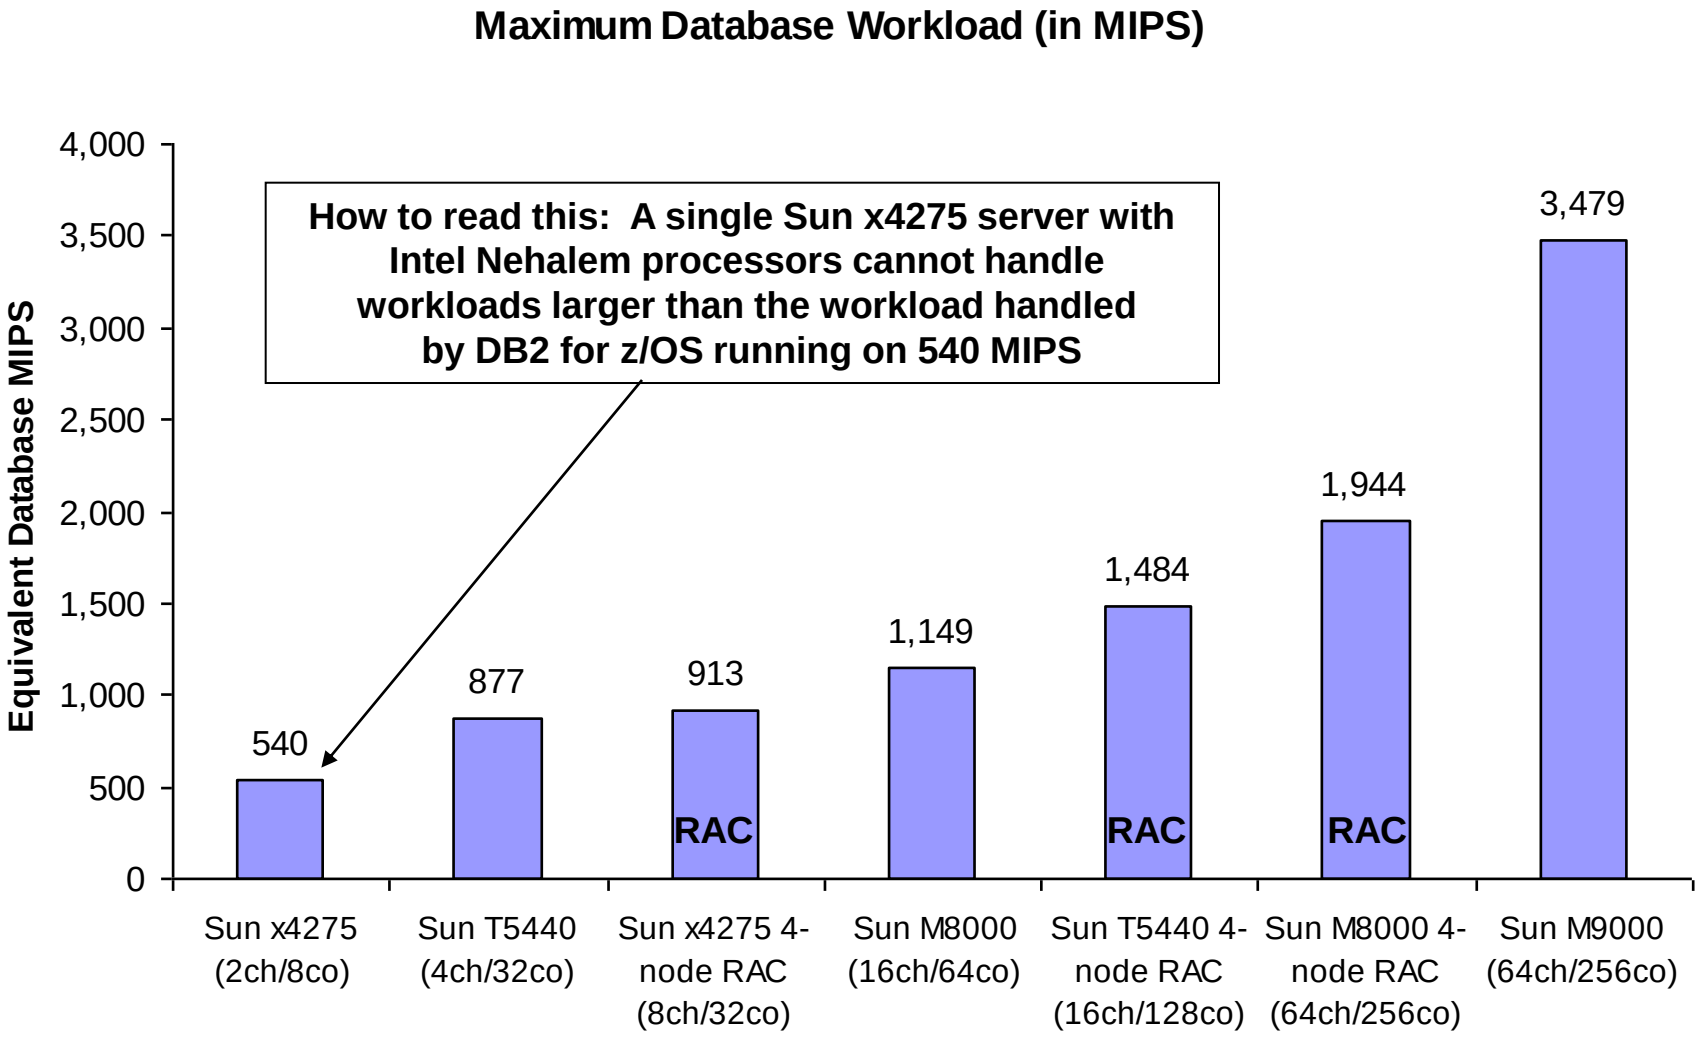

Large Databases Are Best Fit On DB2 For z/OS

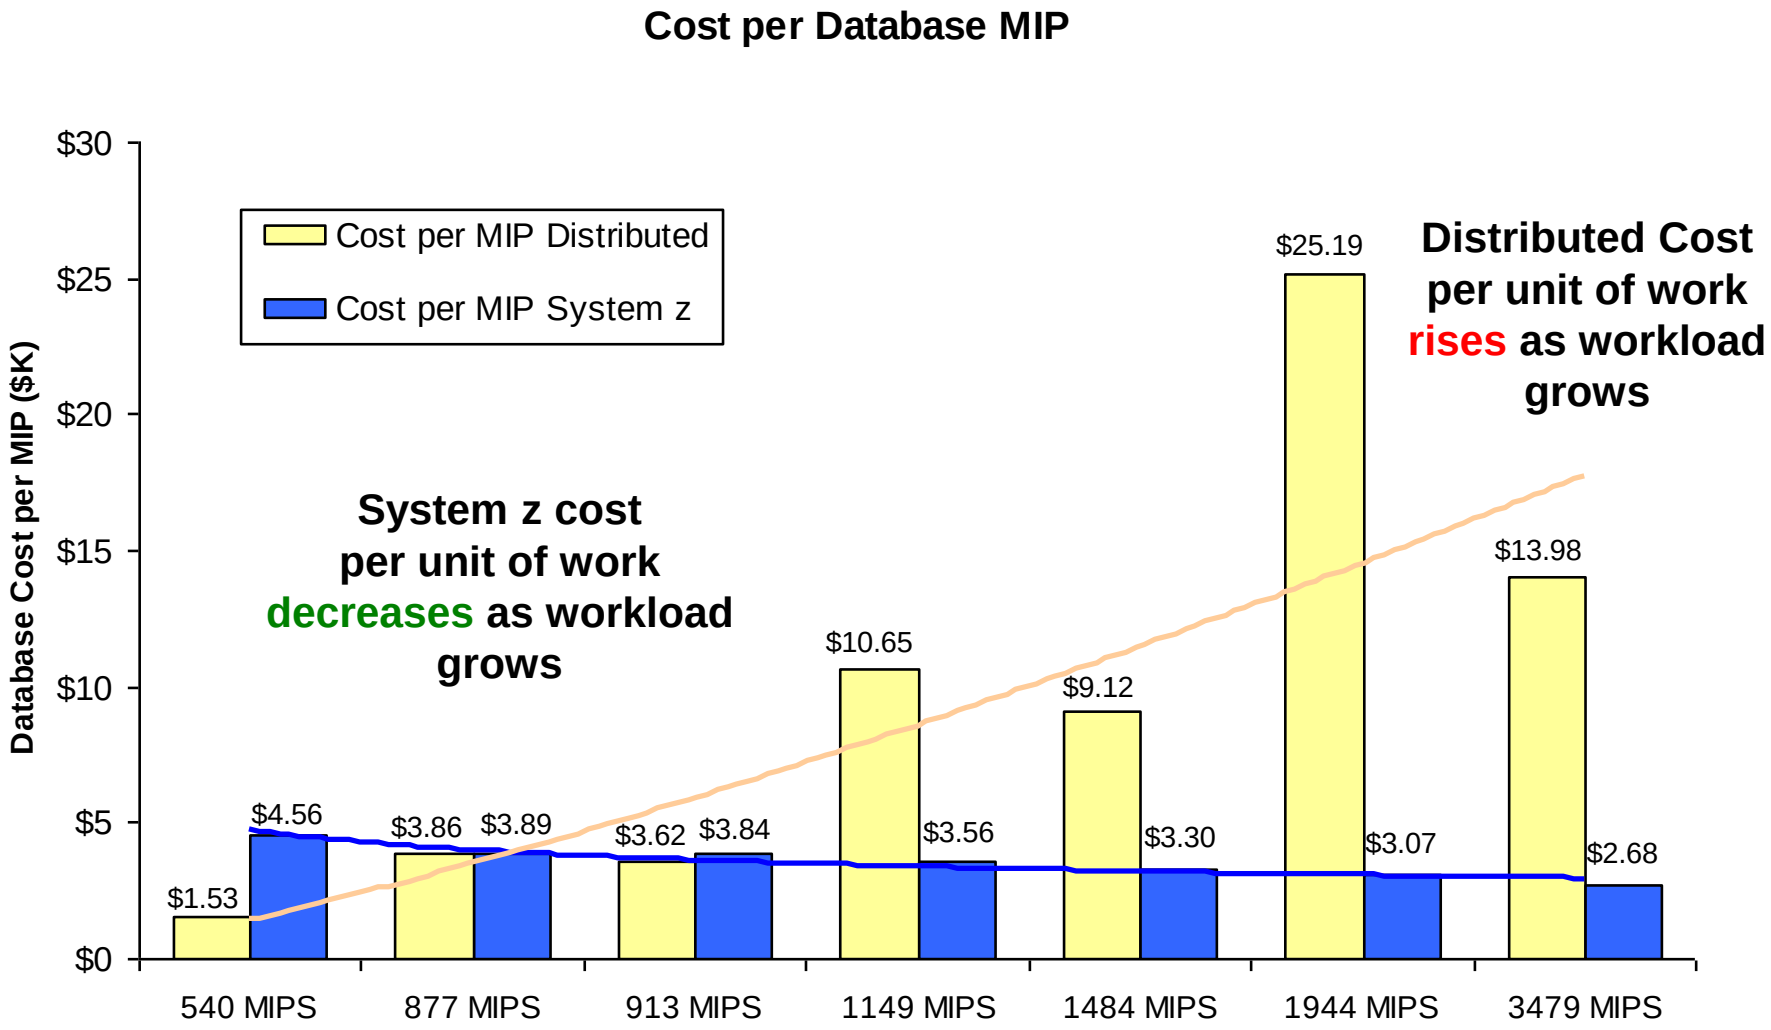

zEnterprise Delivers Over 3x Price Performance vs Oracle/Sun For Deploying SAP Workloads

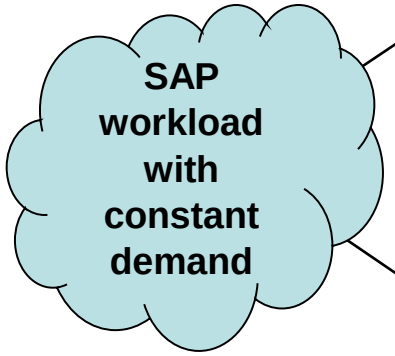

SAP Oracle

Solaris

100% Utilization

M9000 SPARC VII 2.88Ghz 64ch/256co

Sun M9000 Enterprise Server

34,186 users (DB+App)

$311 per user

SAP

DB2

AIX

z/OS

100% Utilization

z196-720 zBX 9 POWER7 Blades

zEnterprise

34,484 users (DB+App)

$94 per user

Based on SAP 2-tier SD Benchmark Sizing for Sun and ISICC sizing for zEnterprise

3-year TCA with Infrastructure (HW, OS and Hypervisor) for App and HW+SW for DB

SAP Solution Edition pricing used for DB and List prices for the rest

Over 3x better price/performance

Solid State Disk Drives Are Here To Revolutionize Storage

- Response times is around 0.8 milliseconds in contrast to 6 milliseconds for a typical hard disk drive
  - 5-10x improvement in throughput and queries
  - SSD drives can sustain I/O rates from 6,000 to 20,000 compared to spinning disk rates of 250 - 300
  - Cost per TB is 10x cost of spinning disks
  - Reduce the “batch window”
- Semiconductor storage available in DS8700 storage sub-systems
  - Random access solid state storage – no moving parts
- Benefits
  - 75% reduction in space
  - 80+% reduction in power and cooling
  - Reduce RAM requirements

DS8700 Easy Tier Capability Automatically Migrates Frequently Accessed Data To SSD

- Automated hotspot detection and migration of data between SSD and HDD
  - Transparent to applications, no code changes required
- Easy Tier maximizes SSD performance gains while minimizing costs
  - Increase performance by up to 300%
  - Relocating just 5% of data to SSDs can reduce response time by 78%
- No charge feature – Microcode update to DS8700

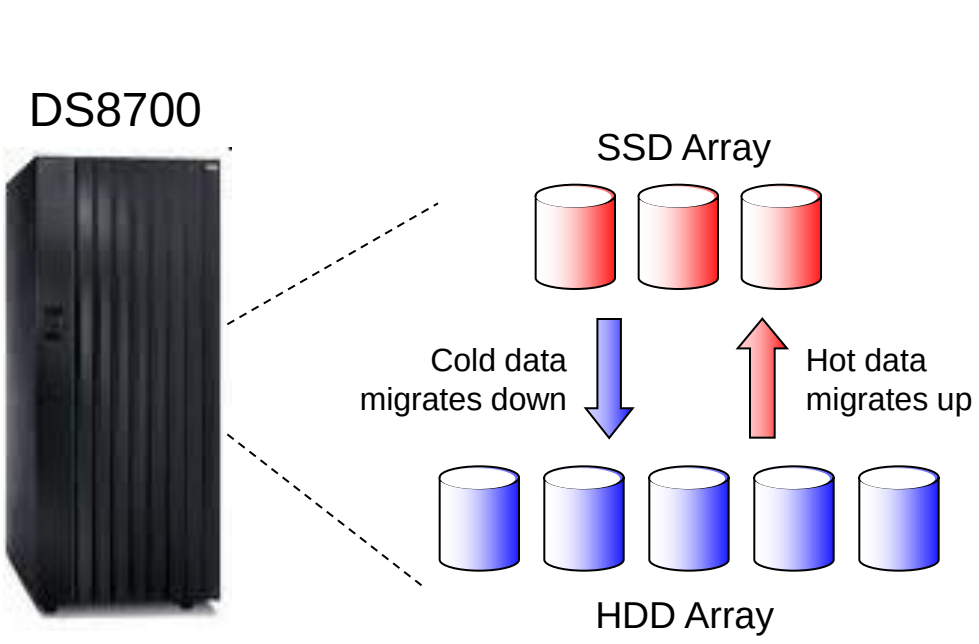

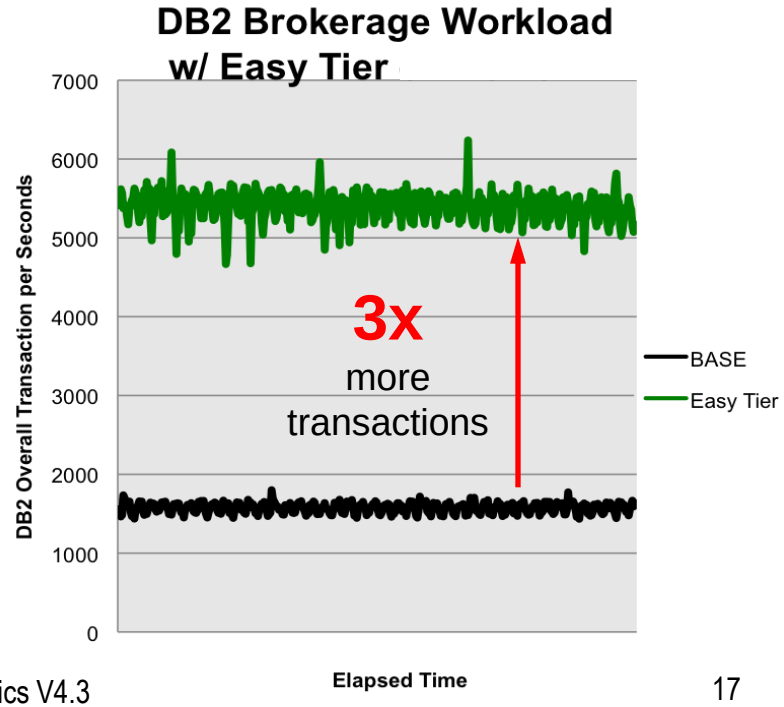

Oracle Database Downtime Can Be Significant With 35 Security Patches Issued Last Year

Oracle's Security Exposures

- Oracle.com – January 2011
  66 security patches, including **6** for the database
- Oracle.com – October 2010
  85 security patches, including **9** for the database
- Oracle.com – July 2010
  59 security patches, including **13** for the database
- Oracle.com – April 2010
  47 security patches, including **7** for the database
- In the last year, Oracle has issued 257 security patches, **35** for the database
  - 35 patches x 4 nodes = 140 possible System Freezes for 4-node Oracle RAC

DB2 For z/OS Provides Rock Solid Security

- Proven granular Multi Level Security leveraging RACF
  - Secures access of tables, views, rows, columns
- End-to-end encryption via hardware assist
  - CP Assist for Cryptographic Function (CPACF) and Crypto Express3
- Optim Test Data Management
  - Ensures anonymous access to data necessary for testing
- DB2 Data Archive Expert
  - Allows customers to easily archive and access data
- DB2 Audit Management Expert
  - Supports compliance requirements
  - Tivoli zSecure Audit for enterprise wide audit

DB2 For System z Provides A Balanced Database For All Smarter Planet Solutions

- Exploits System z Parallel Sysplex for availability and scale
  - RAC's distributed design difficult to scale, freezes during outages
- Dedicated I/O Hardware and Solid State Storage drives I/O throughput
  - Oracle running on distributed servers use precious general purpose processor cycles for I/O
- Superior Qualities of Service
- Competitive Price

Data Plays A Key Role In Smarter Planet Solutions


Smart retail


Smart supply chains


Smart healthcare


Smart food systems


Smart countries

Data

Process Data

- Process operational data at increasing rates
- Perform sophisticated analysis on massive amounts of historical data in a data warehouse

Captured Data

DB2 z/OS

New Intelligence

Load Operational Data From Silos Into A Unified Data Warehouse For New Intelligence

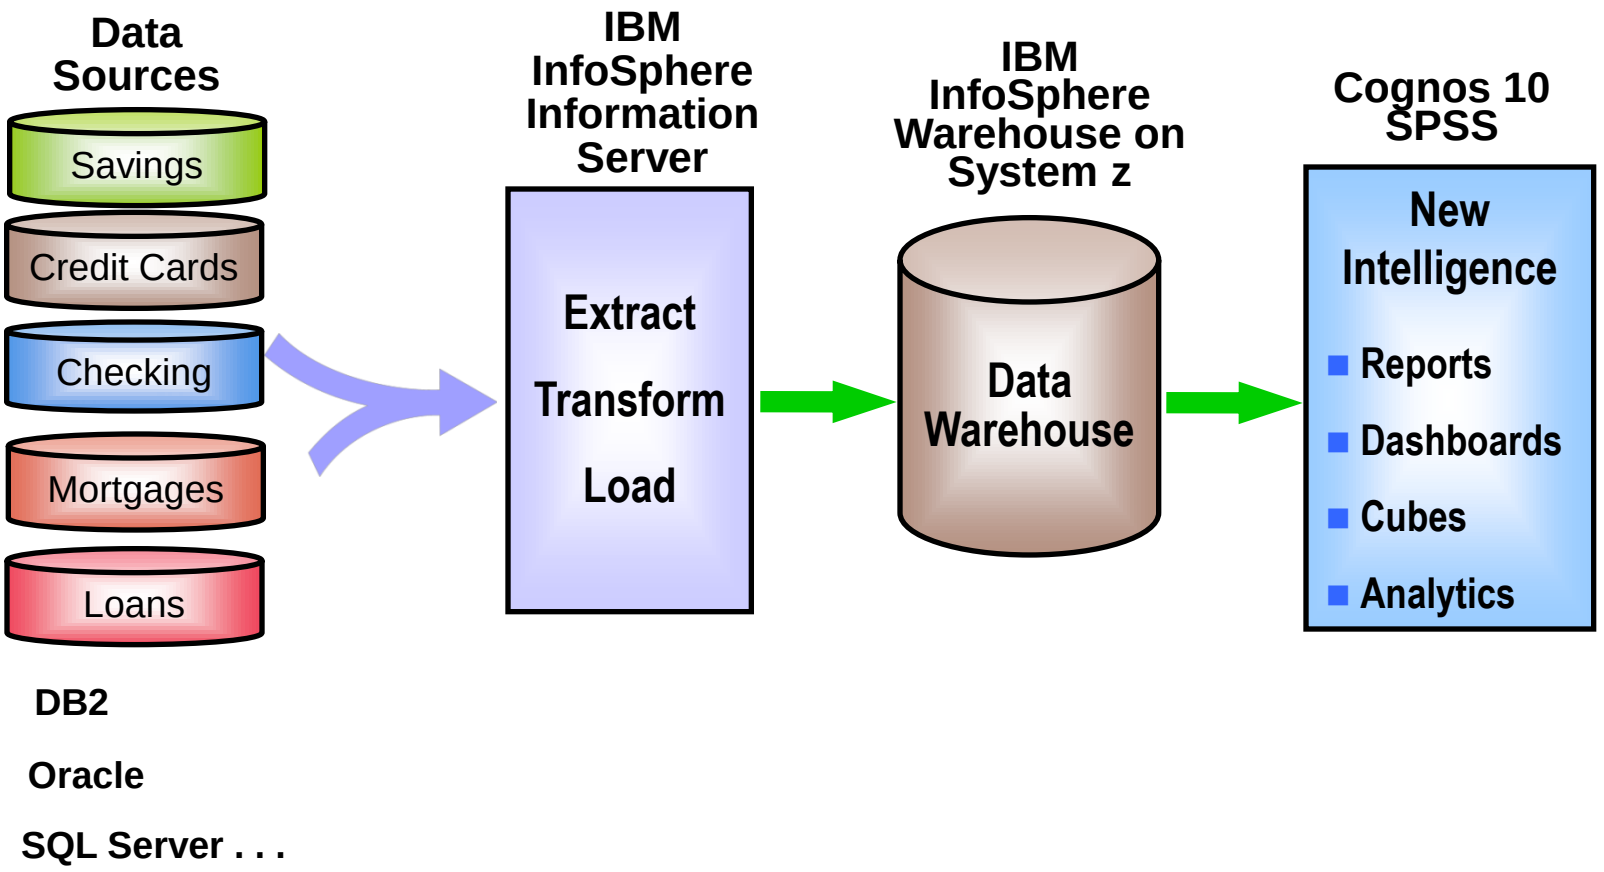

InfoSphere FastTrack Creates Data Maps And Specifications For Your ETL Jobs

- Create simplified data maps and transformations using drag and drop
  - Automatically discover source and target columns
    - Uses database introspection and Web 2.0-style tagging
    - Use business terms to accurately match source to target
- Data analysts and developers share project specifications
  - Collaboration and reuse improve productivity
  - Use metadata common to all Information Server tools
  - Standard formats and centralized management for governance
    - Synchronize work across global teams
- Generate ETL code directly from job specifications
  - Reduces costs and errors in ETL job development

Oracle doesn't offer any of these capabilities

InfoSphere DataStage Implements The ETL Jobs

- Graphical description of ETL jobs using hundreds of pre-built transformation and data quality functions
- Stores and retrieves metadata from Information Server
  - Allows easy reuse of integration work between projects
- Advanced parallel processing capabilities
  - Dynamic partitioning and pipelining
  - Scale jobs across additional hardware without modification

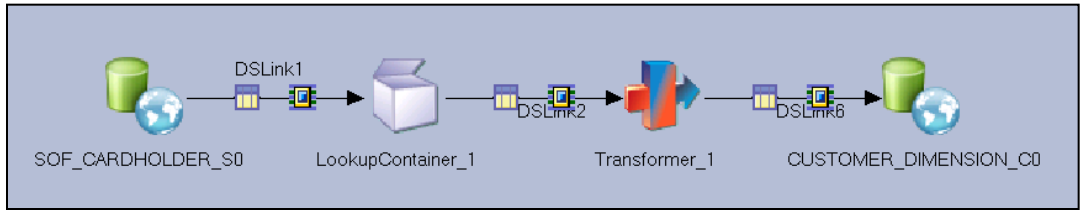

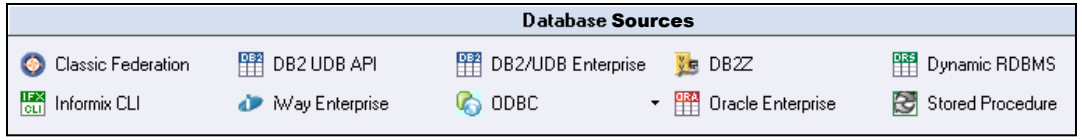

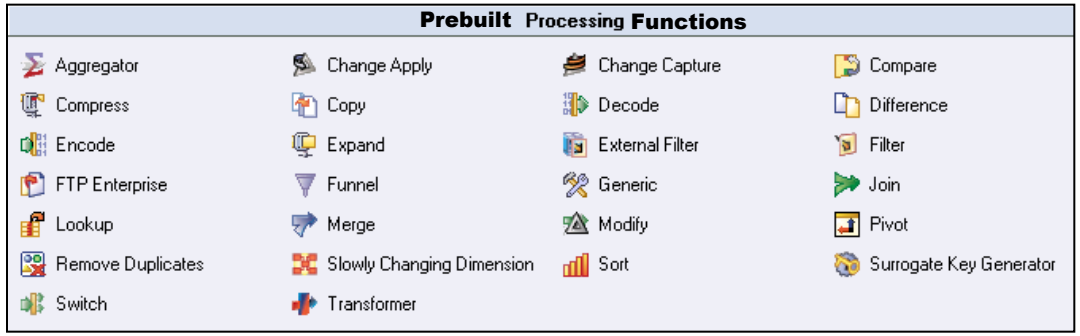

IBM Leads In Data Integration

- Only InfoSphere Information Server delivers unified metadata across all tools for collaboration and reuse
  - Oracle Warehouse Builder and Oracle Data Integrator are two separate products that are not yet unified
- Model-driven design with FastTrack and DataStage speeds development
  - Oracle Warehouse Builder has no tools to help manage source to target mappings
- InfoSphere Information Server works in heterogeneous environments
  - InfoSphere gathers, processes, and cleanses more data from more sources than Oracle

"FastTrack enables our analysts to capture more complete business requirements. The ability to translate this information directly into DataStage jobs with up to 70 percent of the code completed will significantly shorten our development lifecycle."

- Roderich Hofmann, project manager, WAVE, IT-Solutions provider of Bank Austria and member of UniCredit Group

Data Warehouse On System z Trumps Exadata Powered By Inefficient Oracle RAC

- z196 optimizes mixed query performance across OLTP and Warehouse
  - Exadata nodes designed for Warehousing; CPU resources wasted for OLTP
  - Works well for single queries but parallel query performance sags
- New analytics query accelerator makes complex queries much faster
- Application and database co-located on z196 sharing resources
  - Exadata over-provisions for peak, and supports growth through limited choice
- Exploits System z Parallel Sysplex for availability and scale
  - RAC's distributed design difficult to scale, freezes during outages
- System assist processors and I/O sub-system improve scale
  - Exadata I/O handled by x86 processors and architecture, so its performance falls dramatically as data exceeds smart cache size
- Hardware-based compression lowers costs of MIPs and storage
- System z central to disaster recovery and backup strategies
  - Locked-down Exadata does not fit easily into datacenter operations
- System z Warehouse costs less than Exadata

IBM Cognos Is An Integrated Platform For Smart Analytics

- Implemented in Java, runs on WebSphere
- 100% browser based access
  - Server side business intelligence
  - Users can access new intelligence from anywhere
- Easiest for IT to deploy and manage
  - Scales up and out across heterogeneous hardware and operating systems
  - Unified security
  - Unified administration
- Consistent user interface across tooling
  - Greater user satisfaction and increased business agility with lower IT costs
- Common meta data model
  - Author new intelligence assets once, consume anywhere
  - Common view enables open data strategy
  - Supports Unicode and multilingual features without recreating reports

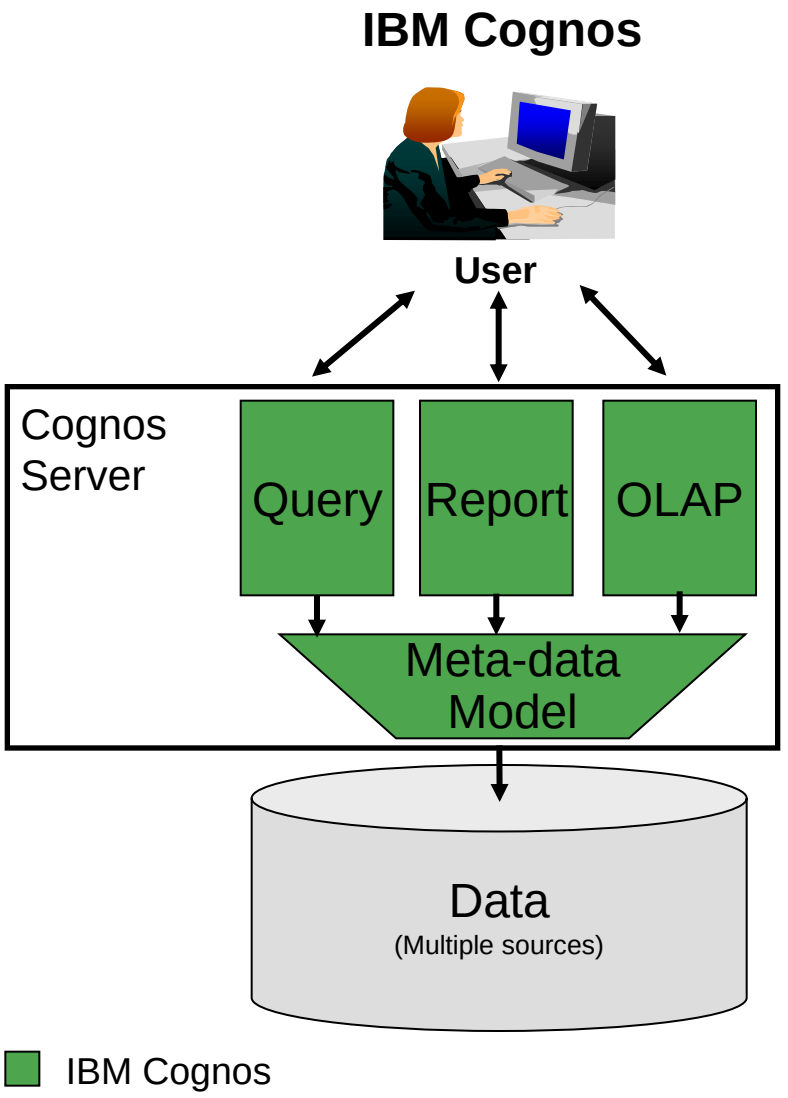

Reuse Trusted New Intelligence Assets Across The Cognos Platform

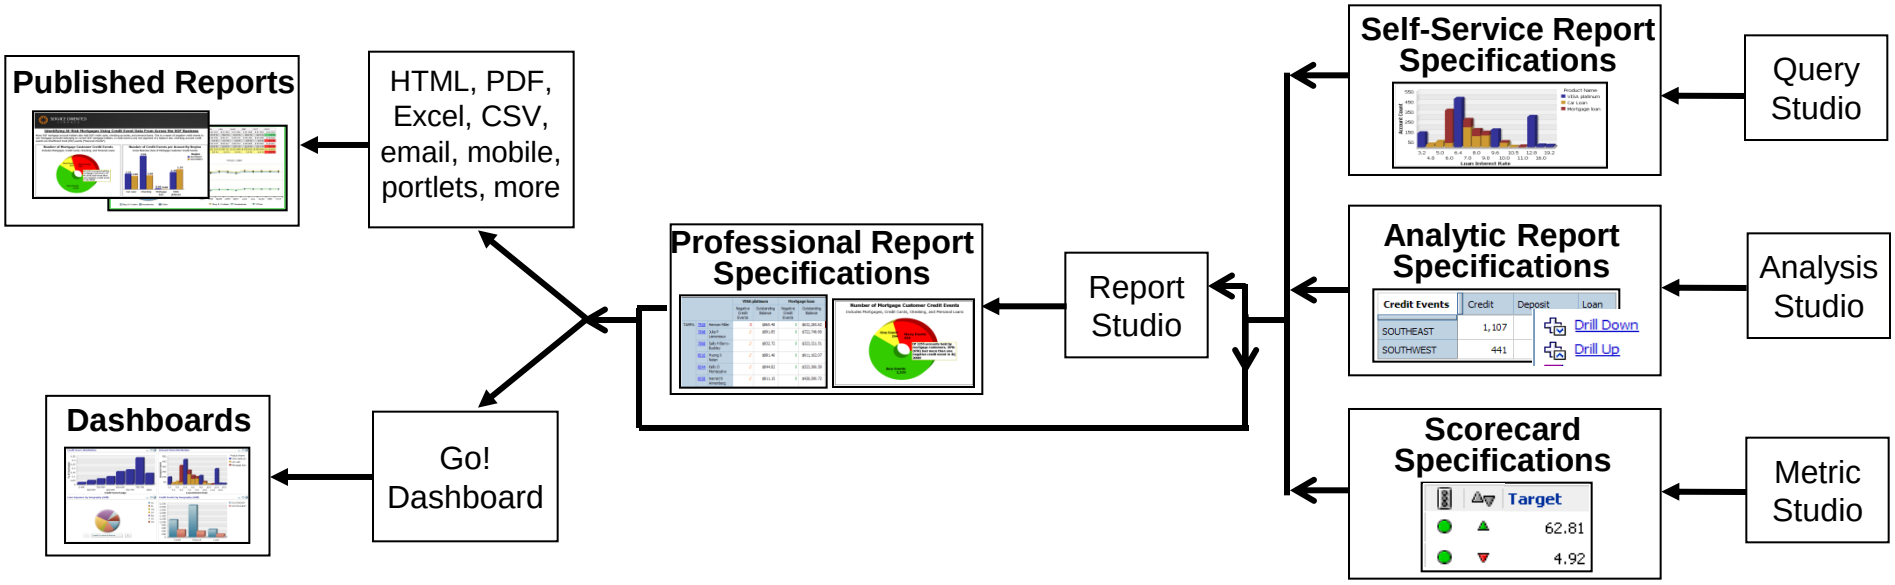

- All new intelligence assets share a common metadata model and common report specification
- Author Once – Consume Anywhere
- Ensures consistent information and enables reuse across platform functions

- Oracle has multiple metadata models depending on source type
- Oracle has multiple different report formats
- Oracle cannot reuse assets between tools

DEMO: Use Cognos To Identify New Business Insights From The Data Warehouse

1. Show report generated in Cognos Report Studio in PDF format
2. Report identifies high-risk mortgages by looking at negative credit events in other customer accounts (CC, Checking, etc…)
3. Report uses both structured and unstructured data (link to mortgage data stored in FileNet)
4. Use Go! Dashboard to monitor the business operations

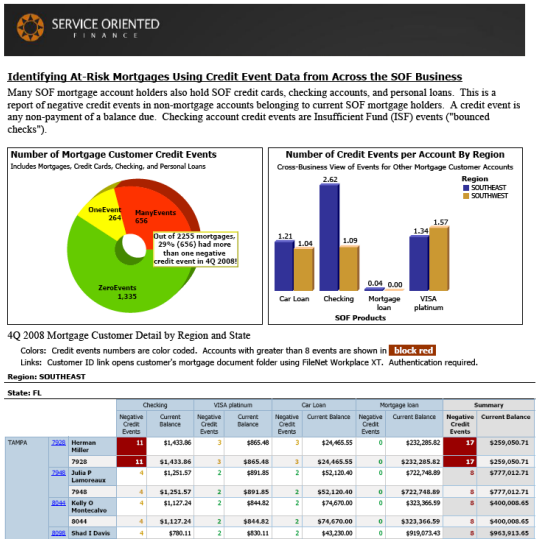

SERVICE ORIENTED FINANCE

Identifying At-Risk Mortgages Using Credit Event Data from Across the SOF Business

Many SOF mortgage account holders also hold SOF credit cards, checking accounts, and personal loans. This is a report of negative credit events in non-mortgage accounts belonging to current SOF mortgage holders. A credit event is any non-payment of a balance due. Checking account credit events are Insufficient Fund (ISF) events ("bounced checks").

Number of Mortgage Customer Credit Events

Includes Mortgages, Credit Cards, Checking, and Personal Loans

Number of Credit Events per Account By Region

Cross-Business View of Events for Other Mortgage Customer Accounts

4Q 2008 Mortgage Customer Detail by Region and State

Colors: Credit events numbers are color coded. Accounts with greater than 8 events are shown in block red

Links: Customer ID link opens customer's mortgage document folder using FileNet Workplace XT. Authentication required.

Region: SOUTHEAST

State: FL

			Checking		VISA platinum		Car Loan		Mortgage loan		Summary	
			Negative Credit Events	Current Balance	Negative Credit Events	Current Balance	Negative Credit Events	Current Balance	Negative Credit Events	Current Balance	Negative Credit Events	Current Balance
TAMPA	7928	Herman Miller	11	$1,433.86	3	$865.48	3	$24,465.55	0	$232,285.82	17	$259,050.71
		7928	11	$1,433.86	3	$865.48	3	$24,465.55	0	$232,285.82	17	$259,050.71
	7948	Julia P Lamoreaux	4	$1,251.57	2	$891.85	2	$52,120.40	0	$722,748.89	8	$777,012.71
		7948	4	$1,251.57	2	$891.85	2	$52,120.40	0	$722,748.89	8	$777,012.71
	8044	Kelly O Montecalvo	4	$1,127.24	2	$844.82	2	$74,670.00	0	$323,366.59	8	$400,008.65
		8044	4	$1,127.24	2	$844.82	2	$74,670.00	0	$323,366.59	8	$400,008.65
	8098	Shad I Davis	4	$780.11	2	$830.11	2	$43,230.00	0	$919,073.43	8	$963,913.65

- At risk customers are identified

IBM Smart Analytics System 9600: Delivers Enterprise-wide Analytics On z196

- **Extend the qualities of service, inherent in the z/OS environment to ensure the availability and security of data.**

- **Hardware/OS**
  - z/OS 1.12
  - IBM zEnterprise technology
  - IBM System Storage DS8700 Intelligent Disk controller
    - Large controller cache and 3 Tier disk offering

- **Unique Software**
  - DB2 for z/OS
  - Cognos included in base offering (Linux on System z)
  - InfoSphere Warehouse (Linux on System z)

- **Optional Components**
  - Solid State drives, integrated within DS8700
    - Easy Tier to identify and migrate “hot data” to SSD
  - IBM Smart Analytics Optimizer

SPSS Enables Customers To Predict Future Events And Drive Better Business Outcomes

Capture

Data Collection delivers an accurate view of customer attitudes and opinions

Predict

Predictive capabilities bring repeatability to ongoing decision making, and drive confidence in your results and decisions

Act

Unique deployment technologies and methodologies maximize the impact of analytics in your operation

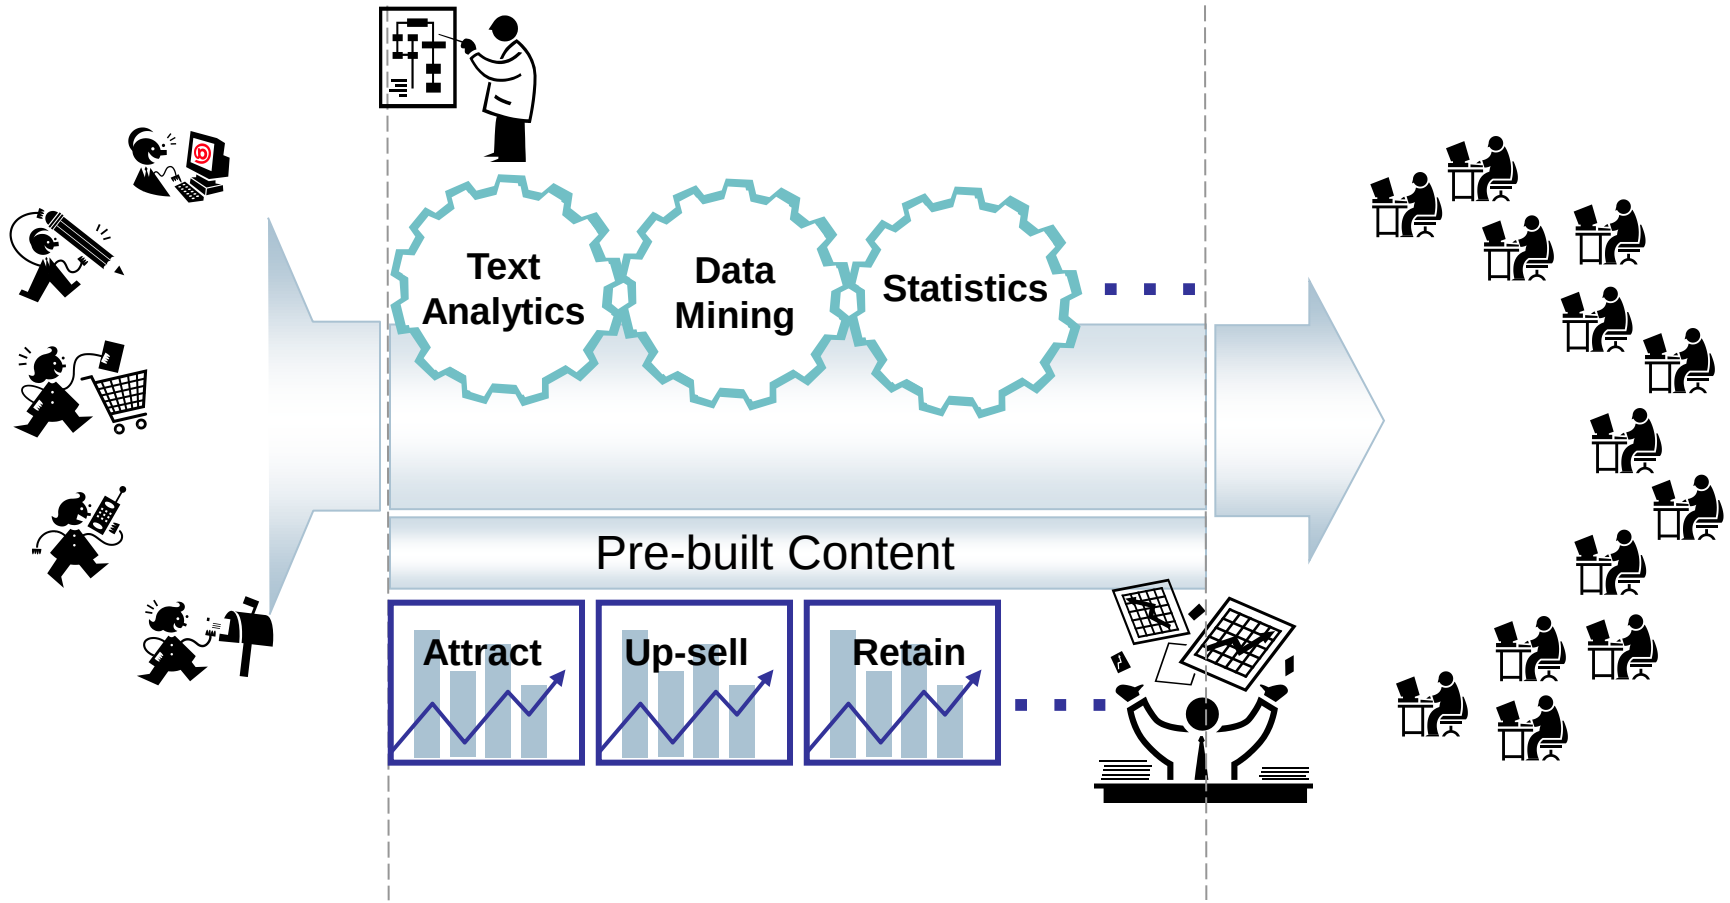

DEMO: Discover Rules For Identifying Risky Customers Using SPSS Statistics

1. Load data from Data Warehouse on DB2 for z/OS into SPSS Statistics
2. Pre-process the data to create new attributes for quantifying negative credit events across different product lines and create a risk flag for mortgage
3. Run Comparison of Means and Decision Tree to discover rules for characterizing risky customers

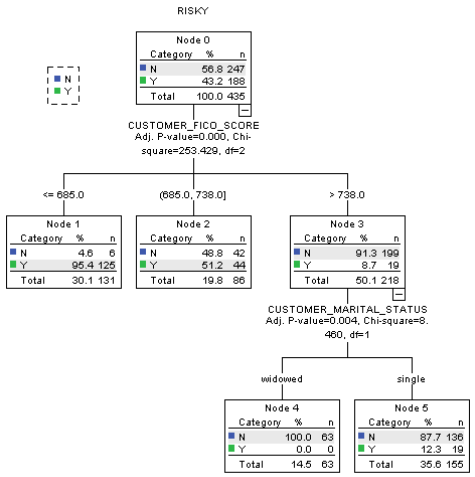

Mortgage
Line of Business VP

- Credit Limits identified for characterizing risky customers
- Use these credit limits for automated loan approval process

Consolidate Complete Business Intelligence Solution On zEnterprise

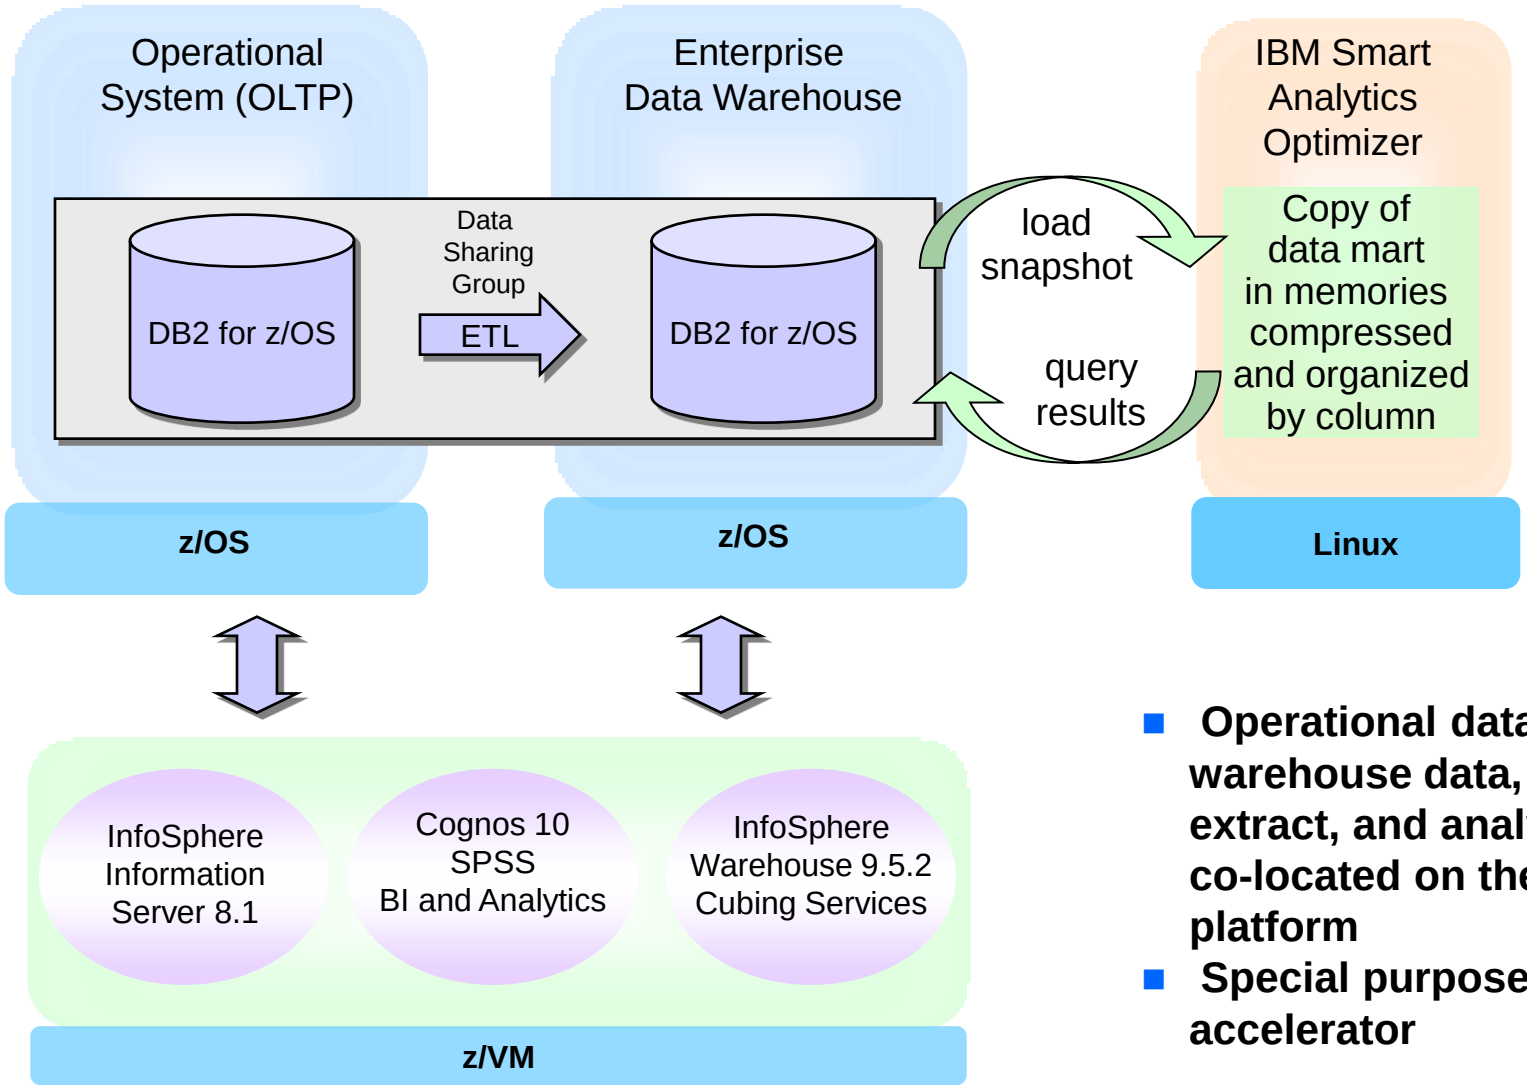

- **Operational data, warehouse data, data extract, and analysis all co-located on the same platform**
- **Special purpose query accelerator**

zEnterprise Is An Excellent Base For Your Data Warehouse And Business Analytics

- Operational and warehouse data co-located on z196
- Exploits System z Parallel Sysplex for availability and scale
- Cognos supports a common metadata model and report specification and provides 100% browser based access
- SPSS predictive analytics provides actionable insights versus hindsight
- Systematic disaster recovery and backup strategies
- Qualities of Service
- Competitive Price