



The Modern Mainframe – At the Heart of Your Business

Information on Demand



Data is at the Heart of Every Enterprise

We are using data in more ways than we've ever used it before. What is IBM doing in the area of database technology?



**On Demand Insurance
CIO**

The world's largest corporations rely on DB2 for z/OS for their most important database needs



IBM

Superior Design of DB2 for z/OS

DB2 has the most advanced technology for databases including large-scale clustered databases. Let's start by looking at our clustering technology

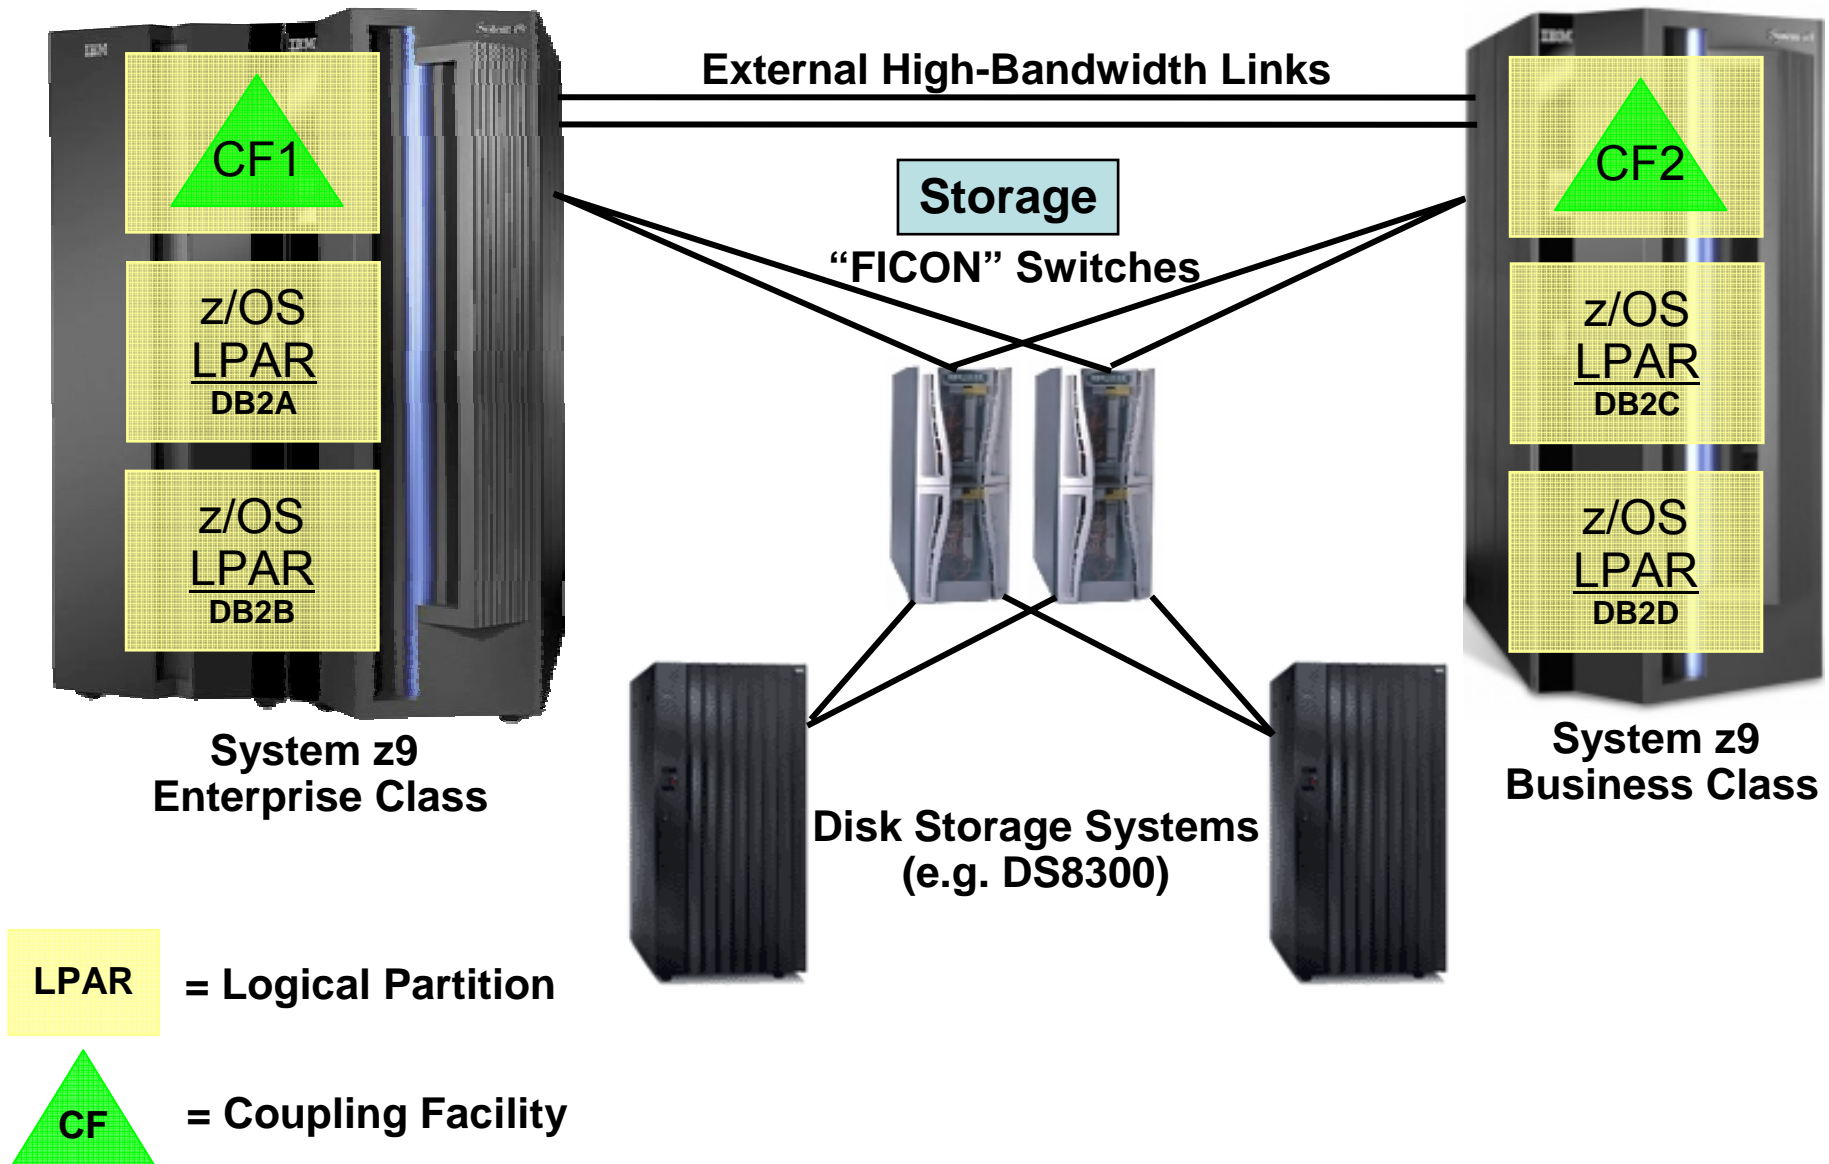


IBM

- DB2 for z/OS has better
 - ▶ Scalability
 - ▶ Recovery
 - ▶ Availability
 - ▶ Security, Encryption & Compression

DB2 on System z Sysplex Design


Example: 4 DB2 Instances in a Cluster



Comparison of Data Sharing Architectures

DB2 for z/OS


Centralized Sysplex Design



High speed centralized
lock manager in
coupling facility

Oracle RAC

Distributed Lock and Data Design

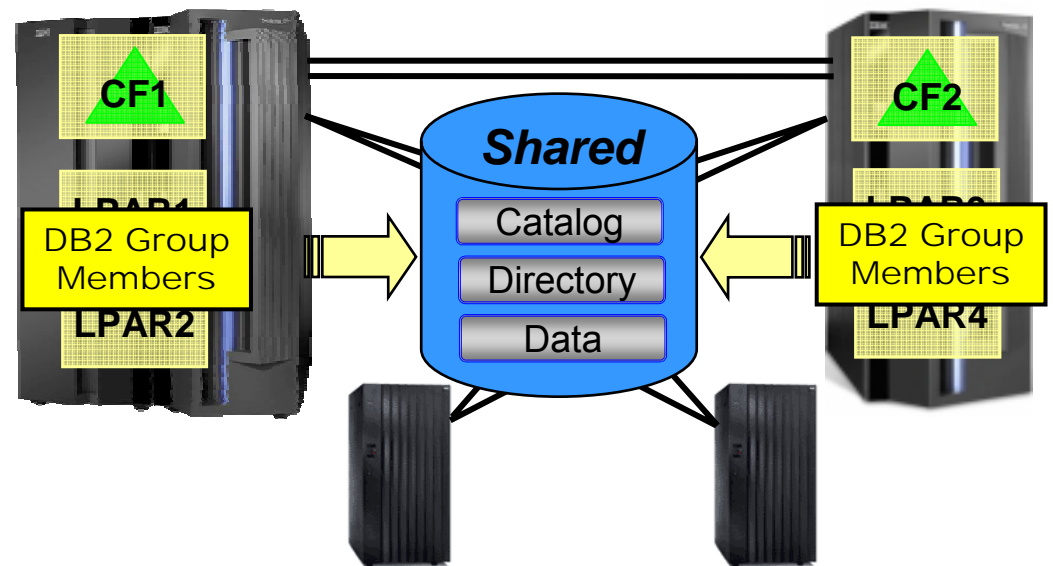


Distributed lock
management with
high messaging **overhead**

DB2 for z/OS Data Sharing Groups

■ DB2 Data Sharing Group

- ▶ Up to 32 DB2 subsystems in the same Parallel Sysplex that all simultaneously access data on shared disks and within Coupling Facilities
- ▶ All group members share the same DB2 catalog and directory
- ▶ Coupling Facilities (CF) provide high-speed caching and lock processing
- ▶ Multiple data sharing groups can be defined within the same Parallel Sysplex



■ *Sysplex query parallelism*

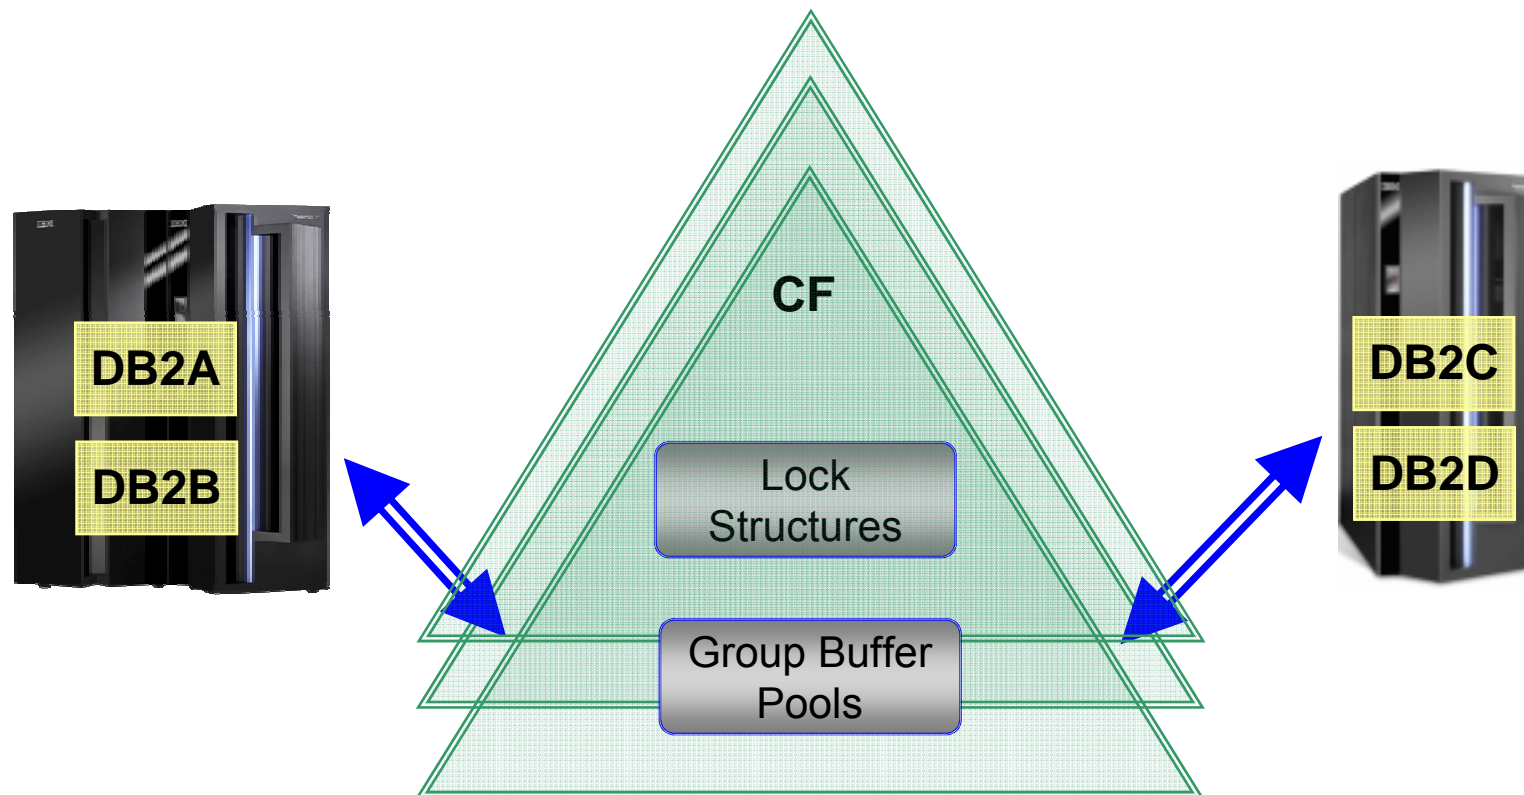
- ▶ DB2 can use all the processing power of the data sharing group to process a single query.

How DB2 for z/OS Data Sharing Works

- To control access to shared data, DB2 uses the locking capability provided by the Coupling Facility (CF)
- DB2 caches data in a storage area inside the CF called a Group Buffer Pool
- Each DB2 member also caches data in local buffer pools
- When a page of data is changed by a DB2 member, the CF invalidates any copy of the page contained in the local buffer pools of other members.
- Subsequent requests for that same data by other DB2 members can retrieve it from the CF's Group Buffer Pool.

How DB2 Uses the System z Coupling Facility

“Duplexing” asynchronously copies cache and lock data between multiple Coupling Facilities

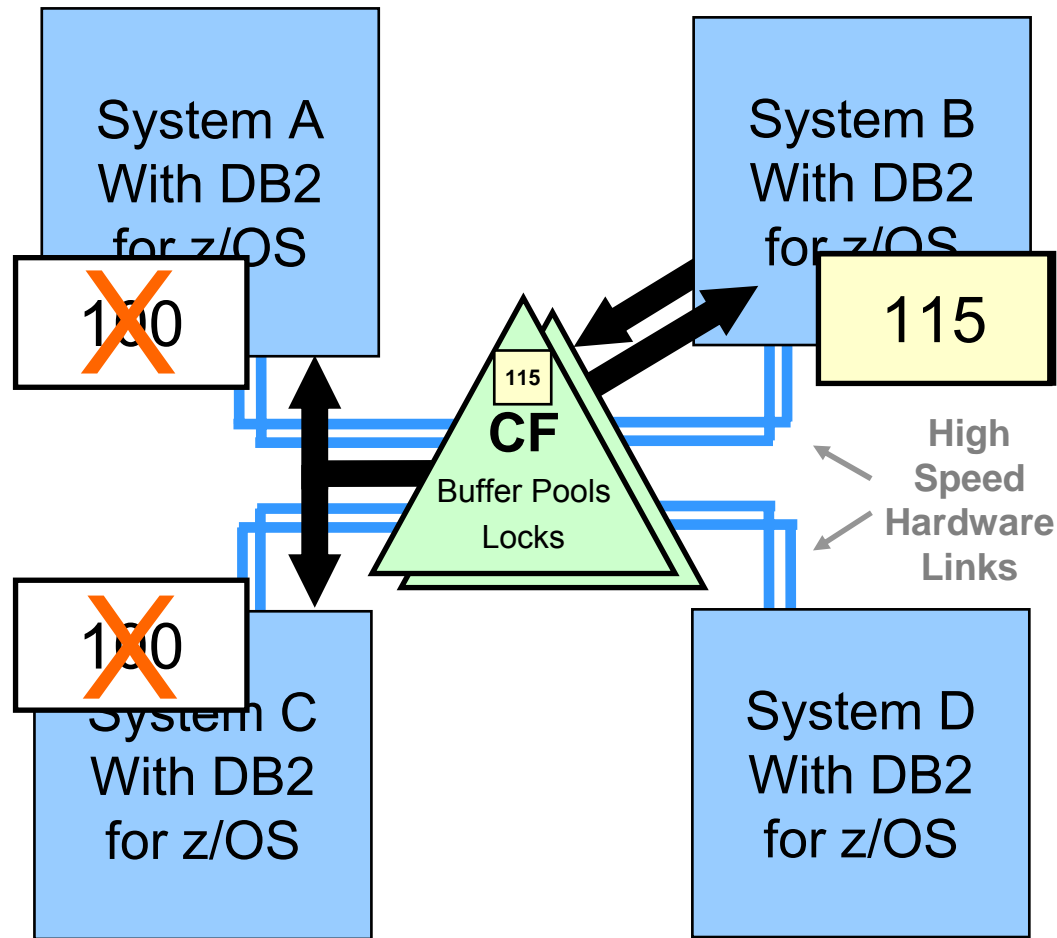


4-Way DB2 "Data Sharing Group"

The DB2 Sysplex Design Scales Linearly as the Cluster Grows

- Beyond two members, each additional member in a DB2 data sharing group adds a minor incremental performance overhead of usually between .5% and 1%
 - ▶ No direct interactions with other nodes
 - ▶ CF handles the workload of lock and buffer management
 - ▶ Result is near linear scale out
- With Oracle RAC, processing resource in each node is used to service lock and buffer requests from other nodes
 - ▶ The processing overhead grows as nodes are added
 - ▶ More overhead means less transaction throughput
 - ▶ Result is limited scale out

Centralized Coupling Facility Permits Efficient Lock and Cache Management in DB2 for z/OS



A, B, and C have read locks with local copies

1. B Obtains write lock
2. B Updates local copy
3. B Caches update in buffer pool
4. CF invalidates all cached copies without interrupting processors

Cache and locks are maintained with no inter-node disturbance!

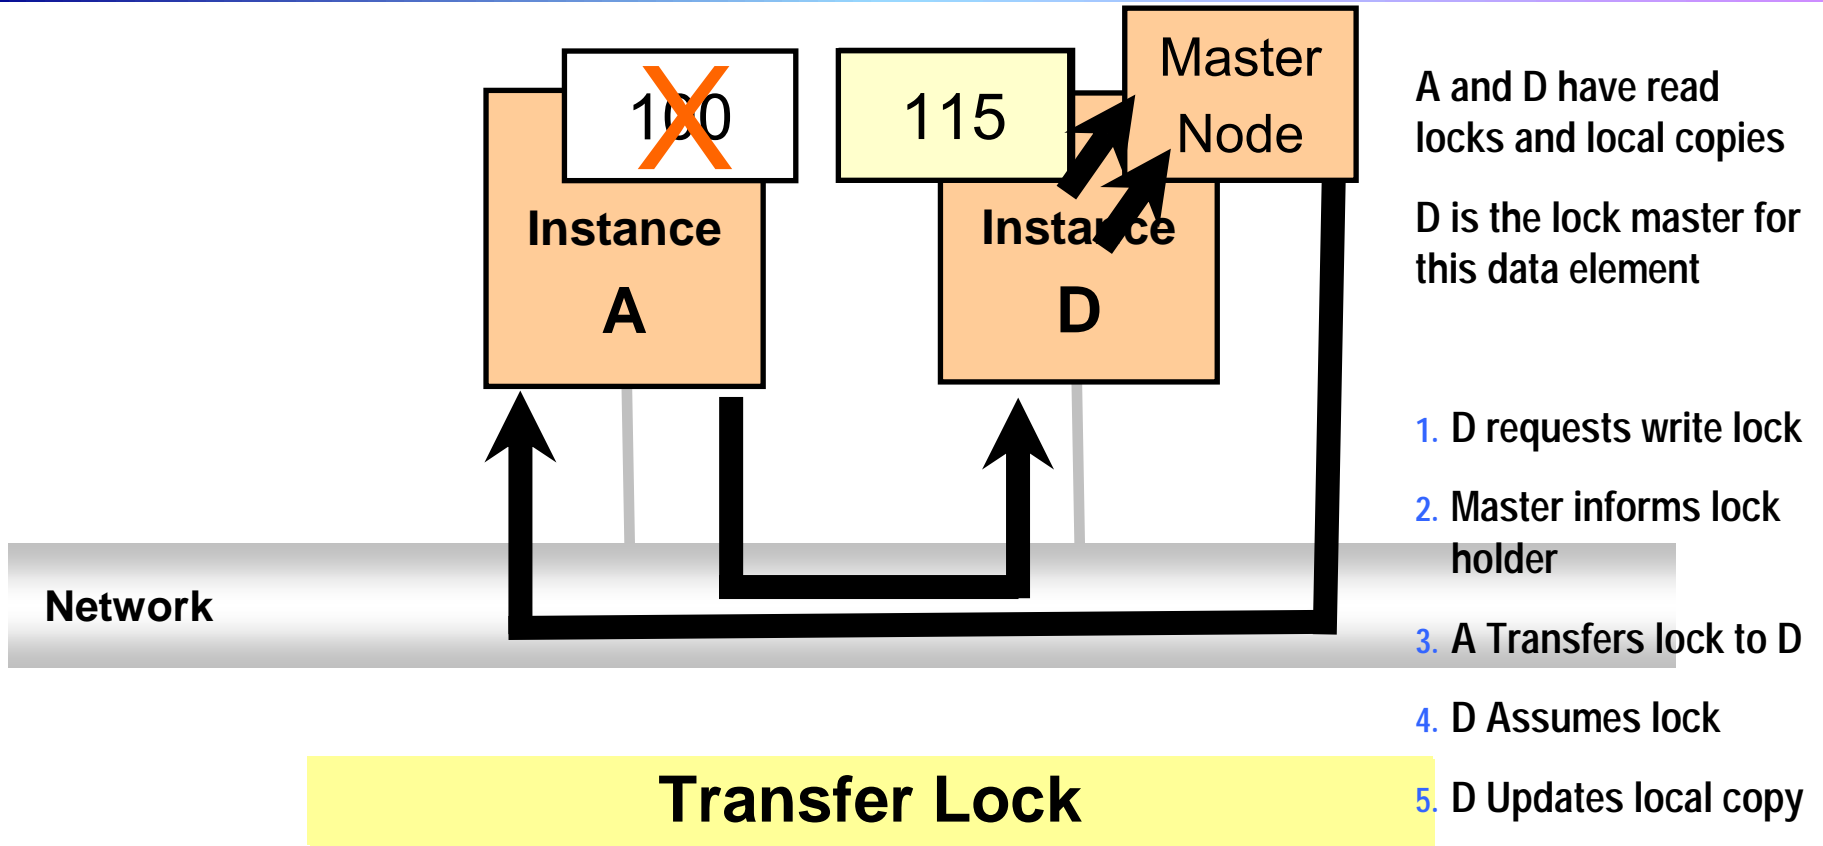
Why is Oracle RAC Scalability Limited?

RAC Inefficiencies increase as a cluster grows

- RAC Nodes must constantly communicate to process requests to maintain distributed cache and lock data.
- Adding additional nodes to the cluster results in increased inter-node communication and requires additional CPU and network time to complete the same operations.
- RAC distributed lock management overhead increases faster than the added capacity of more nodes.

Let's look at some examples...

Oracle RAC: Network Based Distributed Lock Management Involves High Overhead




Inter-node connections: 2

In a cluster with 2 nodes an update operation may need 2 network connections and two in-memory calls (not shown).

Example based on Oracle's US Patent 7,107,319 B2.

Oracle RAC: Lock Management Overhead Increases as the Cluster Grows Larger



Lock Assume


7. B Updates local copy

Inter-node connections: 6

In a cluster with 4 nodes, an update operation may need 6 network connections and two in-memory calls (not shown).

Example based on Oracle's US Patent 7,107,319 B2.

Oracle RAC Overhead Wastes Processing Power in Each Node



Oracle RAC source:
"Scale-up versus
scale-out using Oracle
10g with HP
StorageWorks",
Hewlett-Packard, 2005


DB2 for z/OS source:
"Enterprise Data Base
Clustering Solutions"
ITG, October 2003

Result: DB2 for z/OS Scales Out, Oracle RAC is Limited

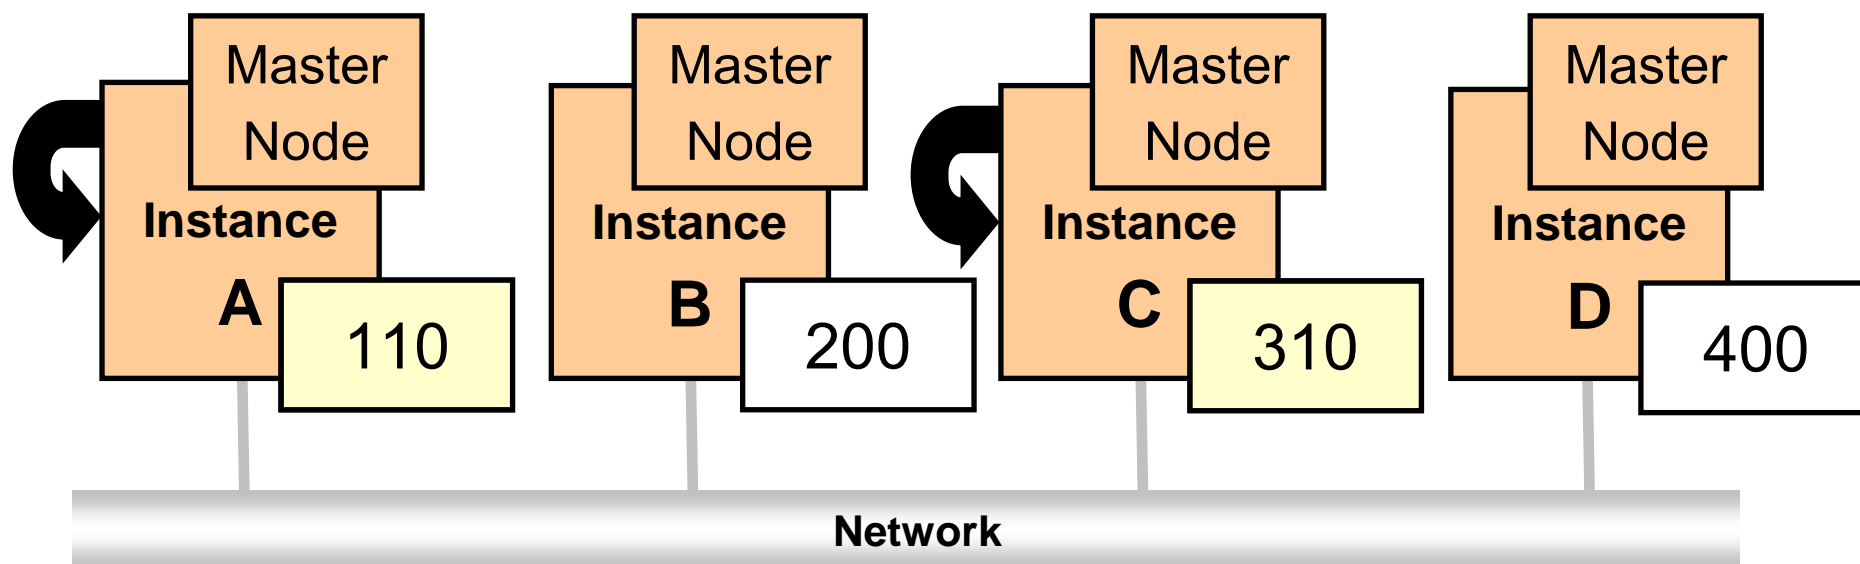
- DB2 for z/OS provides near linear scalability with relatively little overhead as nodes are added
- With Oracle RAC, overhead increases rapidly as additional nodes are added and performance degrades after only 4 to 6 nodes

Oracle RAC source: "Scale-up versus scale-out using Oracle 10g with HP StorageWorks", Hewlett-Packard, 2005

DB2 for z/OS source: "Enterprise Data Base Clustering Solutions" ITG, October 2003




Partitioning is Necessary to Effectively Use Oracle RAC Above 4 to 6 Nodes



- Partitioning associates the data, buffer, and lock manager with a particular server node
- All work requests affecting the associated data partition must be executed by the owning node
- Data partitioning requires a redesign and reimplementing of the database and the application (big job!)
- Data partitioning may not always be possible depending on the workload

Near-Linear Scalability with DB2 for z/OS with No Partitioning Required

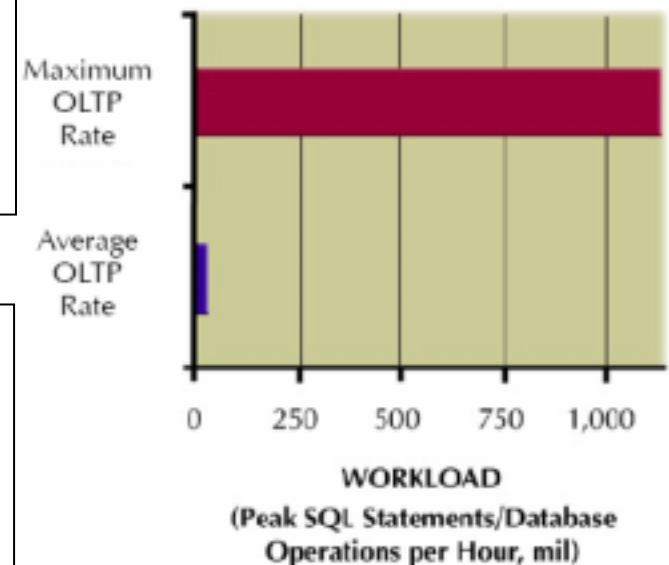
Large Asian Bank Surpasses Throughput Goals



Compare Scalability Achievements: Winter Corporation's "2005 Top Ten" Awards

- "The highest performing transaction processing system in the 2005 program, a [DB2] z/OS implementation, executed **over one billion** SQL statements in an hour. The average for operational systems was **35 million** SQL statements or database operations per hour"
 - The study lists the largest known peak workload on **Oracle RAC** to be **8.6 million** SQL statements per hour
-
- "The largest transaction processing [database] in the program, **23 TB**, was hosted on [DB2] z/OS, as in the last program"
 - The study lists the largest transaction processing database on **Oracle RAC** to be **9.6 TB**

Leading OLTP System,
DB2 for z/OS, Processed
Over One *Billion* SQL
Statements per Hour



http://www.wintercorp.com/VLDB/2005_TopTen_Survey/TopTenWinners_2005.asp
http://www.wintercorp.com/WhitePapers/WC_TopTenWP.pdf

Why Do the World's Largest Companies Rely On DB2 for z/OS?

Here are few good reasons:



A very large Asian banking customer uses DB2 for z/OS to meet and exceed all of their current and projected future transaction processing workloads!

- ▶ Manages 380 million accounts
- ▶ System z and DB2 for z/OS readily achieved this customer's projected targets of 15 million banking transactions per hour and proved additional performance and scalability capacity to more than double this throughput.



UPS, the international shipping company, runs DB2 for z/OS to support the world's largest known peak workload for any database - 1.1 Billion SQL statements per hour! *



The UK government's Land Registry runs DB2 for z/OS to maintain the world's largest known transaction processing (OLTP) database – 23.1 TB! *

* As listed in Winter Corporation's "2005 Top Ten" awards - <http://www.wintercorp.com/index.html>

Availability: DB2 for z/OS Can Eliminate Unplanned Outages


- **With DB2 for z/OS, unplanned outages due to hardware or software failure are highly unlikely**
 - ▶ The System z environment is very resilient to hardware and software failures and can recover from or adjust to most problems transparently without impacting DB2 availability or response times
 - Inherent z platform RAS
 - Sysplex redundancy
- **However, if a DB2 instance is lost or taken offline due to a hardware or software failure:**
 - ▶ Recovery of locks occurs when the failed DB2 instance automatically restarts.
 - “Restart Light” option allows for a faster recovery
 - ▶ During the recovery process, all transactions continue to run successfully except transactions that require access to data pages locked by the failed DB2 instance.
 - ▶ Coupling Facilities prevent any loss of cache or locks.

Availability: Unplanned Outages are Probable with Oracle RAC

- **Oracle RAC is dependent on less reliable hardware than System z**
 - ▶ For example, with an HP configuration, a processor or other possible hardware failures within an O/S partition results in a total partition outage
 - O/S partition loss results in the loss of an Oracle RAC cluster member
 - Hardware such as processors, RAM, and disk dedicated to this O/S partition cannot be used by other partitions, which may have a significant impact on database response times (No partition virtualization like System z PR/SM)


- **By design, RAC loses cache and lock data whenever a node fails!**
 - ▶ While RAC rebuilds the cache and locks across the remaining nodes, the entire cluster is unavailable for some period of time.
 - ▶ “Remastering” time depends on amount of transaction activity at the time of failure.

Summary – Unplanned Outages (Loss of a Node)



DB2 for z/OS

on Parallel Sysplex




A

- Unlikely due to highly reliable Z architecture
- Never loses cache or lock data
- No period of zero availability
- In-flight transactions rolled back.

Oracle 10g RAC

on distributed hardware



C

- More Likely due to less reliable mixed hardware platforms
- Always loses cache and lock data
- Variable period of zero availability
- In-flight transactions rolled back.

Availability: Oracle RAC Clusters are Vulnerable to TCP/IP Network Failures

- **RAC uses a private TCP/IP or UDP/IP network between nodes to manage shared cache and lock information.**

- **A TCP/IP network failure between cluster nodes can bring down the entire Oracle RAC cluster**
 - ▶ Depending on how the network interconnect fails, some instances may survive, but all instances in the cluster may shut down.
 - ▶ False node “evictions” can occur under heavy loads if this network becomes saturated
 - Oracle “evicts” a server from the RAC cluster by telling it to reboot itself if it doesn’t “check in” with the Oracle Cluster Synchronization Services (CSS) master within a certain amount of time.

- **Additional networking hardware can be added to provide a failover configuration for the network interconnect, at a significant additional complexity**

Availability:

Rolling Upgrades Minimize Planned Downtime

- DB2 allows rolling “on-line” upgrades and maintenance with mixed releases of DB2
 - ▶ Activity at the node needing maintenance is first quiesced, then work is routed to the remaining nodes, and activity is resumed once the node becomes active again
 - ▶ Can fall-back to prior release without outage
- The Oracle RAC solution requires downtime to upgrade the database
 - ▶ Oracle claims rolling upgrade support for minor fixes, but only if the fixes support it, and few actually do.
- Frequent Oracle patches drive high administration and operational costs.
 - ▶ Oracle has shipped over 100 patches in a single quarter.
 - “Downtime for both routine maintenance and upgrades is too high and becoming more critical due to global operations”*

Quote from AMR Research: Oracle Peer Forum Teleconference, Feb 2006
<http://www.amrresearch.com/Content/View.asp?pmillid=19165>

Availability:

DB2 Online Maintenance Minimizes *Planned* Downtime

- Database changes can be made while running, for example:
 - ▶ Add or modify a table schema without dropping and recreating
 - Table or column changes, type and length
 - Dynamic adding and rotating partitions

- Performance changes can be made with database online
 - ▶ Reorganization of the database
 - ▶ Secondary Indexes partitioning
 - ▶ Partition without an index; cluster on any index
 - ▶ Online database parameter changes

- Housekeeping operations can be done without taking databases down
 - ▶ Image copy (backup) while database is running

TD Bank Financial Group

10 Years of Continuous Availability with DB2 for z/OS!

 DB2 in TD Bank's Parallel Sysplex has had zero unplanned downtime since October, 1996. **10 years!!**

 TD Bank migrated from DB2 V7 to V8 (32bit->64bit upgrade) with zero downtime.


- zSeries is used for Customer Account Data for the applications supporting Tellers, Internet Banking and Automated Banking Machines
 - ▶ Parallel Sysplex supports 42 million business transactions a day




■ Client Environment

- ▶ System z
- ▶ z/OS
- ▶ DB2
- ▶ IMS
- ▶ WMQ
- ▶ GDPS .


Availability: Minor version upgrade



DB2 for z/OS
on Parallel Sysplex



Oracle 10g RAC




A


New fixes can be applied in a **rolling fashion without impact to the application**

C+ Not all fixes are rolling upgradeable

Availability: Major Version Upgrade




DB2 for z/OS
on Parallel Sysplex



A

New versions of DB2 for z/OS can be applied in a **rolling fashion without impact to the application**

Oracle 10g RAC



D

Does not support rolling version upgrades

DB2 for z/OS:

The Highest Availability Database Server

	Oracle RAC	DB2 for z/OS
Availability during unplanned outage	<ul style="list-style-type: none">■ Failover model■ Distributed lock management■ Node failure -> redistribute & remaster locks■ Entire database frozen■ Cluster Management SW fails -> all nodes in cluster fail	<ul style="list-style-type: none">▪ Continuous Availability model▪ Integrated with z/OS CF▪ Locks persist across nodes▪ Never entire blackout▪ Proven with most challenging workloads
Availability during Maintenance and Upgrades	<ul style="list-style-type: none">■ No mixed release capability■ No easy fallback (need full restore)■ No data availability on upgrade■ Significant security patches	<ul style="list-style-type: none">▪ Rolling upgrades including HW▪ Run mixed releases▪ Fall back without outage▪ Continuous availability▪ Planned outages minimized▪ Schema Evolution + in V8
Availability during disaster recovery	<ul style="list-style-type: none">■ Generic backup and recovery■ Not integrated■ More complex■ Added performance overhead	<ul style="list-style-type: none">▪ Integration means less complexity▪ Industry leading backup & recovery▪ XRC▪ GDPS

Security has Risen to the Top of Requirements

As DW becomes more real time we need to ensure we have tighter security.

- DB2 for z/OS has
 - ▶ One security model, platform for data
 - ▶ HW exploitation of encryption
 - ▶ Roles, Contexts
 - ▶ Simpler process for auditors & administrators
 - ▶ Simpler process with Compliance tools

Payment Card Industry Data Security Standard

- Encrypt transmission of data
- Assign unique ID to each user accessing cardholder data
- Track and monitor all access to network and cardholder data
- Global Standard set by American Express, Discover, JCB VISA, Mastercard et al...
- Mandatory compliance for all businesses dealing in Payments

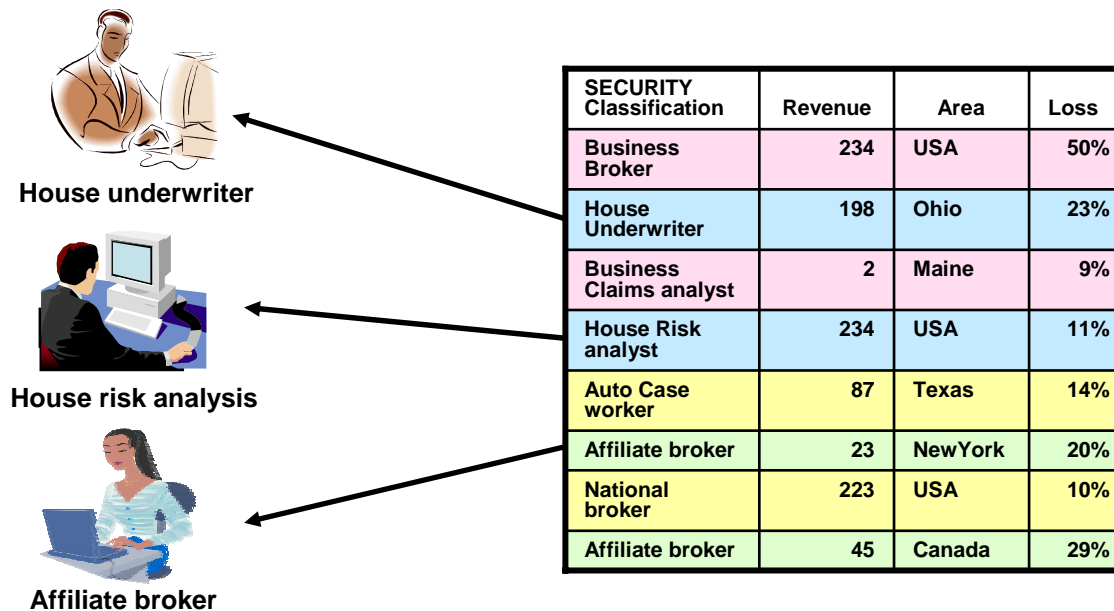
DB2 Protects All Data on Disk

- We will allow encryption for the key disk resources used by DB2:
 - ▶ Tables, LOBs, Indexes
 - ▶ Image copies
 - ▶ Logs and Archives

- Future Directions –
Extending Encryption to IBM TotalStorage
 - ▶ zOS manages keys like Tape encryption
 - ▶ Encryption CPU cycles offloaded

Multi Level Security Access (MLS) Helps Reduce Security Complexity

- Builds on traditional “discretionary” access permissions by adding additional access controls that are based on a comparison between the *classification* of the user with the *classification* of the data.
- Facilitates compliance with audit and governmental regulations, and protection of highly sensitive data down to the row level - **without writing custom code.**
- Eliminates need for duplicate I/T infrastructures, silos, and redundant software

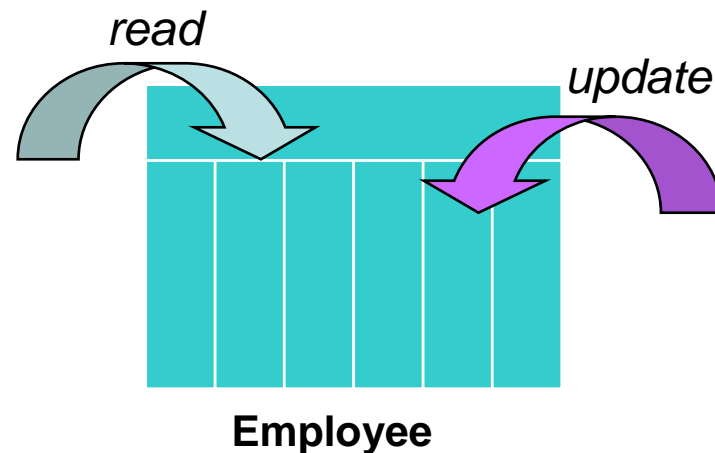


A single image of data is sharable by multiple enterprise departments with different levels of “need to know”

Compliance made easy – Audit Management Expert

- ▶ Centralized auditing and tracking...who is accessing sensitive data
- ▶ Enable Auditors to be independent of DBA's
 - Enable auditors to collect, view, analyze, and report on data via the audit repository
- ▶ Provides an auditor-friendly reporting UI....
- ▶ Many user friendly options for examining data in the repository
 - Allows detailed analysis and visualization of data collected by the DB2 auditing tool
 - Auditors can export audit data into other applications such as Excel®.
- ▶ Product provides Batch reporting...report security violations
- ▶ Can perform Log Analysis to view changed data values
- ▶ Reduce DBA tasks related to auditing, setting traces, log analysis et al

Scenario: Using DB2 AME to uncover unauthorized access attempts to read and update sensitive data



Now we'll walk through a scenario to show how DB2 Audit Management Expert can help auditors uncover unauthorized access attempts to read and update sensitive data. Let's say the sensitive data in question is an employee table that contains social security numbers and salary information. We'll show how DB2 Audit Management Expert helps the auditors spot unauthorized access attempts, both failed and successful, to this table.

Profile Name	Description	Last Modified	Rules	Active Collections
Authorizations Audit Profile	Authorizations Audit Profile	2006-08-16 11:23:47	1	0
Command Audit Profile	Commands Audit Profile	2006-08-16 11:21:35	1	0
Emp Audit Profile	AME Demo Profile	2006-08-16 11:19:45	1	1
Profile 1	Profile 1	2006-08-14 21:54:27	1	0
Utilities Audit Profile	Utilities Audit Profile	2006-08-16 11:22:44	1	0

Let's begin by looking at the audit data that is being collected. The audit data to collect is defined in a collection profile within Audit Management Expert's Administration user interface. In this scenario, the administrator has already defined a few collection profiles. Let's look at the one named "Emp Audit Profile", since this one covers the employee table.

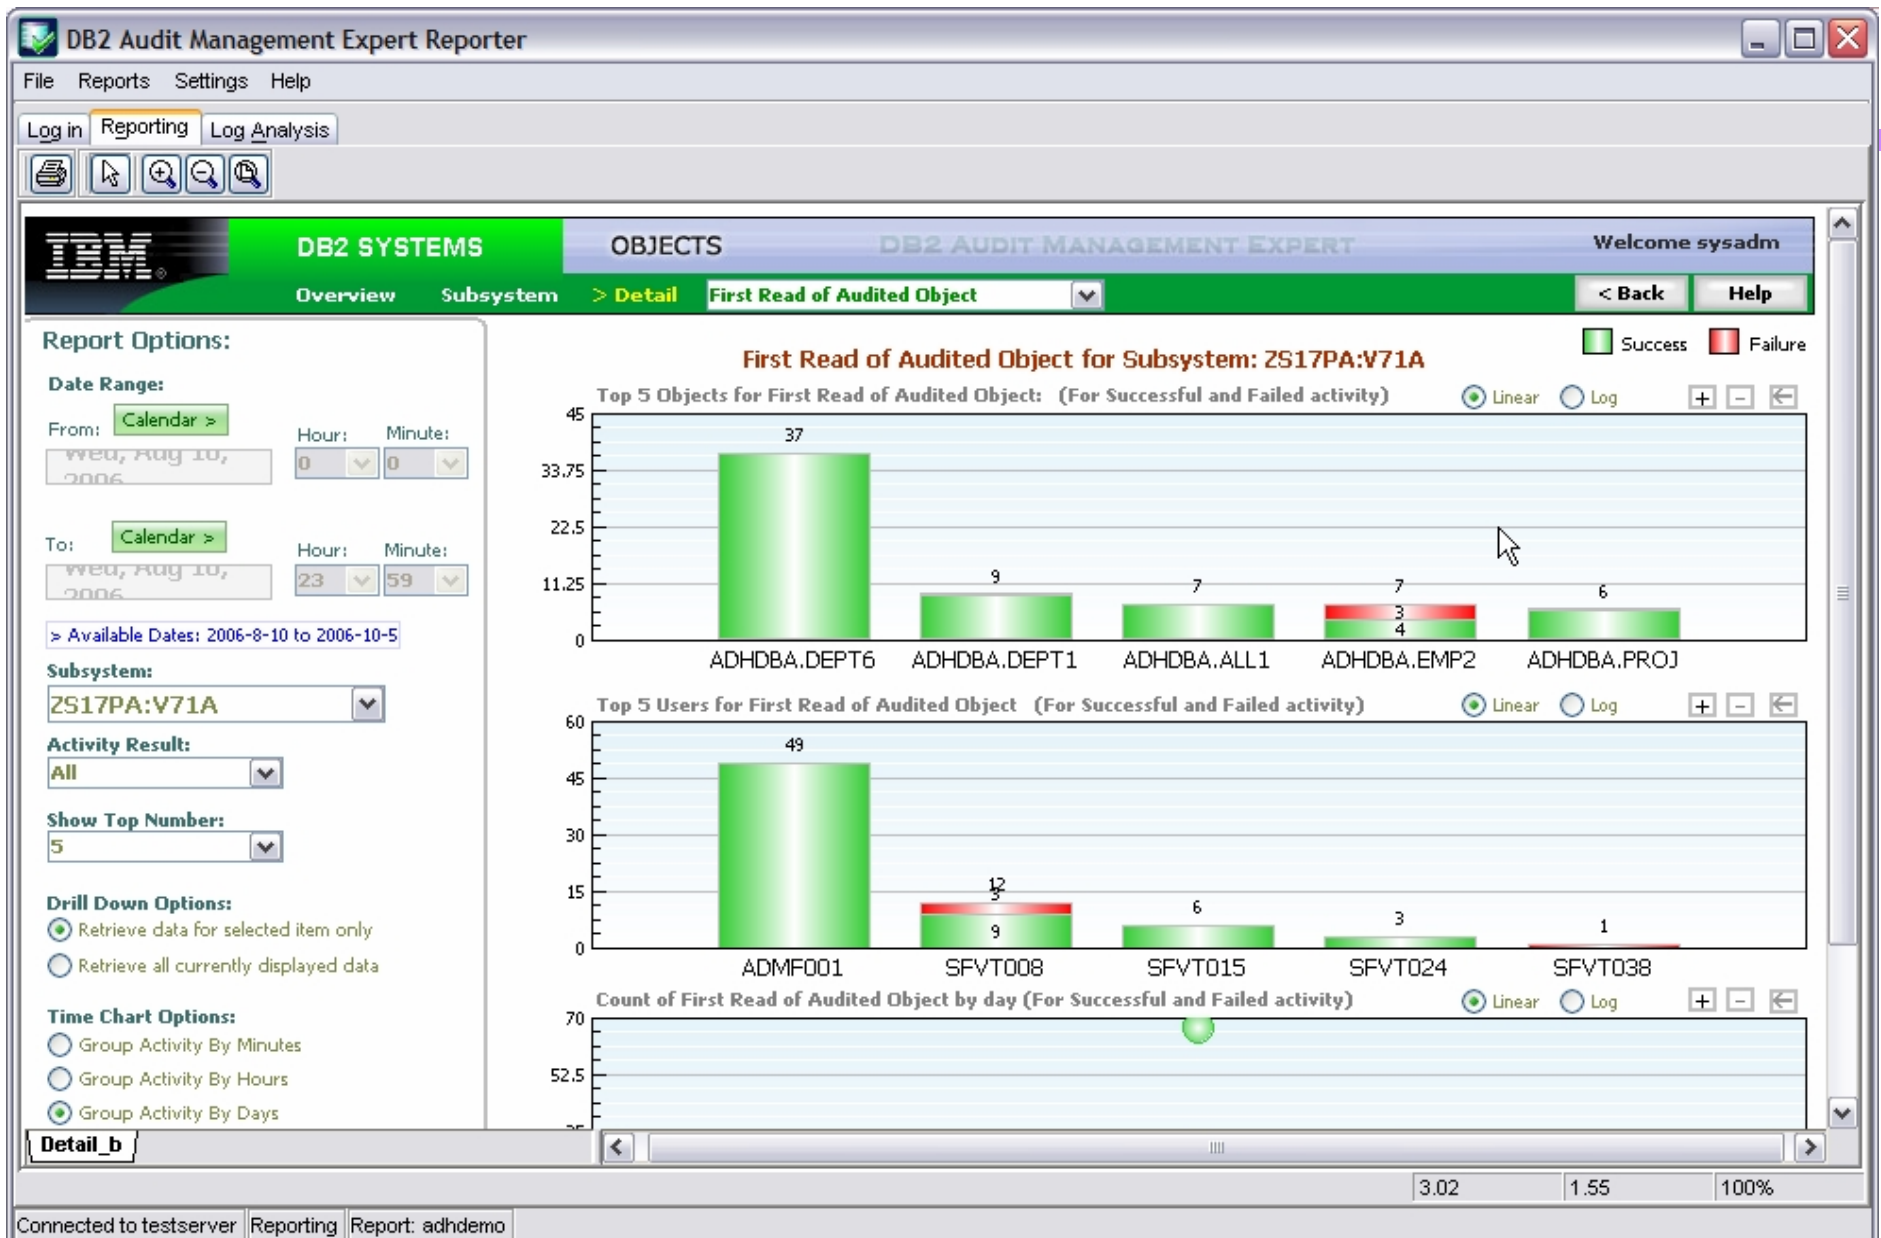
The profile summary shows the rules that determine how Audit Management Expert performs auditing.

The list of Audit targets/tables shows the objects to be audited. We can see that in this collection profile, auditing is turned on for several tables including the Employee table that contains our sensitive data.

For each audited table, we are monitoring the read (or selects) events and the change (insert, update, and delete) events for those tables.

Now let's look at Audit Management Expert's Reporting User interface to see the auditing reports that are generated after the data has been collected. The left side of the Reporting User Interface lists the various report options. You can filter by date, userID, activity type and threshold. The right side of the user Interface shows a summary/overview of all the audit activity for a given subsystem. Thresholds can be set by the user in order to flag potential warning and problem situations that would need to be further investigated.

Continuing with this scenario, we're interested in looking at read attempts on the Employee table, so we'll click on Detail, then select the First Read of the Audit Object, which provides a more detailed view of this audit activity. The first graph shows the top five objects that have the most read attempts. (The number of objects shown is a user setting.) Red represents a failed read attempt and green represents a successful read attempt. The Employee table that contains the sensitive data, Emp2, is one of these top objects. The Emp2 table had 4 successful and 3 failed read attempts. Clicking on the Emp2 table will generate a more detailed report for first read audit activity.



The detailed audit activity report shows several interesting things. We can see the time in which the read events were attempted, and the auth ID attempting to read as well as what they were attempting to read: social security numbers. We also see the same auth ID (sfvt008) had several failed read attempts followed by successful read attempts.

Data Security: DB2 for z/OS Has a Proven Track Record

DB2 for z/OS Security

- ▶ Proven System z security
- ▶ Proven RACF and MLS security
 - ▶ End to end security including applications disks, printers and network
- ▶ DB2 Test Database Generator
 - ▶ Ensures anonymous access to data necessary for testing
- ▶ DB2 Archive Expert
 - ▶ Allows customers to easily archive and access data
- ▶ DB2 Audit Management Expert
 - ▶ Allows automated record retention
- ▶ Less than 10 security related patches in the last 10 years.

Oracle's Security Exposures

- **IT Week - 01/16/2007**
"Oracle flags up 52 security flaws"
- **ComputerWorld - 10/17/2006**
"Oracle releases 101 patches in quarterly update"
- **NGS Research - 11/21/2006**
"The conclusion is clear – if security robustness and a high degree of assurance are concerns when looking to purchase database server software – given these results one should not be looking at Oracle as a serious contender."
<http://www.nextgenss.com/research/papers/comparison.pdf>
- **Gartner - 01/26/2006**
"Oracle has not yet experienced a mass security exploit, but this does not mean that one will never occur."
http://www.gartner.com/DisplayDocument?doc_cd=137477

Oracle RAC Quiz

- How many nodes are in an Oracle cluster on average?

▶ under 3

4

6 or more

- What % of Oracle database installs run RAC, 5 years after it was first available?

▶ under 3%

9.5%

19%

- How many DB2 for System z customers have migrated to Oracle RAC?

▶ zero

two

130

Oracle is struggling with clustering scalability since >14 years

DB2 for z/OS Has the Lowest TCO

- DB2 for z/OS has the lowest TCO to host a data warehouse than any other vendor or platform
- Best Compression saves storage costs
 - ▶ Hardware compression lowers CPU overhead
 - ▶ Providing 50-80% reduction in storage requirements, saving money
- DBA tools improve TCO
- zIIP specialty engine enables cost effective data consolidation

Storage Costs: Customer Feedback Shows 50-80% DB2 Compression

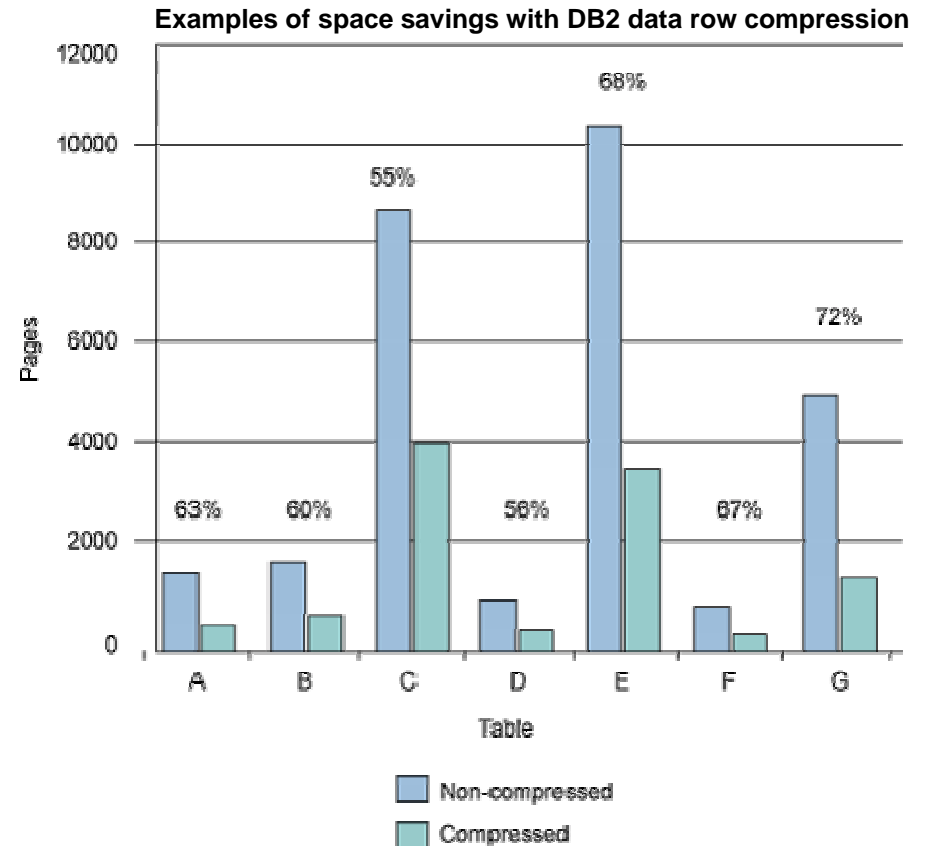
- Michael Henson *SunTrust Bank, Inc.:*

“With DB2 9, we’re seeing compression rates up to 83% on the data warehouse. The projected cost savings are more than \$2M initially with ongoing savings of \$500,000 a year.”

- For another bank customer data set, a 179.9GB table using 32KB pages was reduced to only 42.5GB: a savings of 76.4%

- Sean McCown of *InfoWorld:*

“The new compression method translates into direct storage savings at an average ratio of 45 to 75 percent, depending on the data type. To test this, I created a 40GB table (approximately 500 million rows) with mixed numeric and text data, exported it to a text file, and then imported it into a compressed table format. The size of the compressed table was about 17.75GB -- a savings of 56 percent.”




<http://www-128.ibm.com/developerworks/db2/library/techarticle/dm-0605ahuja/>

Storage Costs:


Key DB2 for z/OS advantages over competition



Oracle	DB2
Software Implementation	Hardware Implementation
<ul style="list-style-type: none">■ At <i>page</i> level■ Based upon repeating <i>column values</i>■ Only supports <i>bulk load</i>	<ul style="list-style-type: none">■ At <i>table</i> level■ Based upon repeating <i>patterns</i>■ Supports <i>load, update, insert, reorg</i>



DB2 Has a Better Performance Impact



DB2 Has Better Compression Ratios

SQL server does NOT support compression

DB2 – Better Compression Ratio Than Oracle

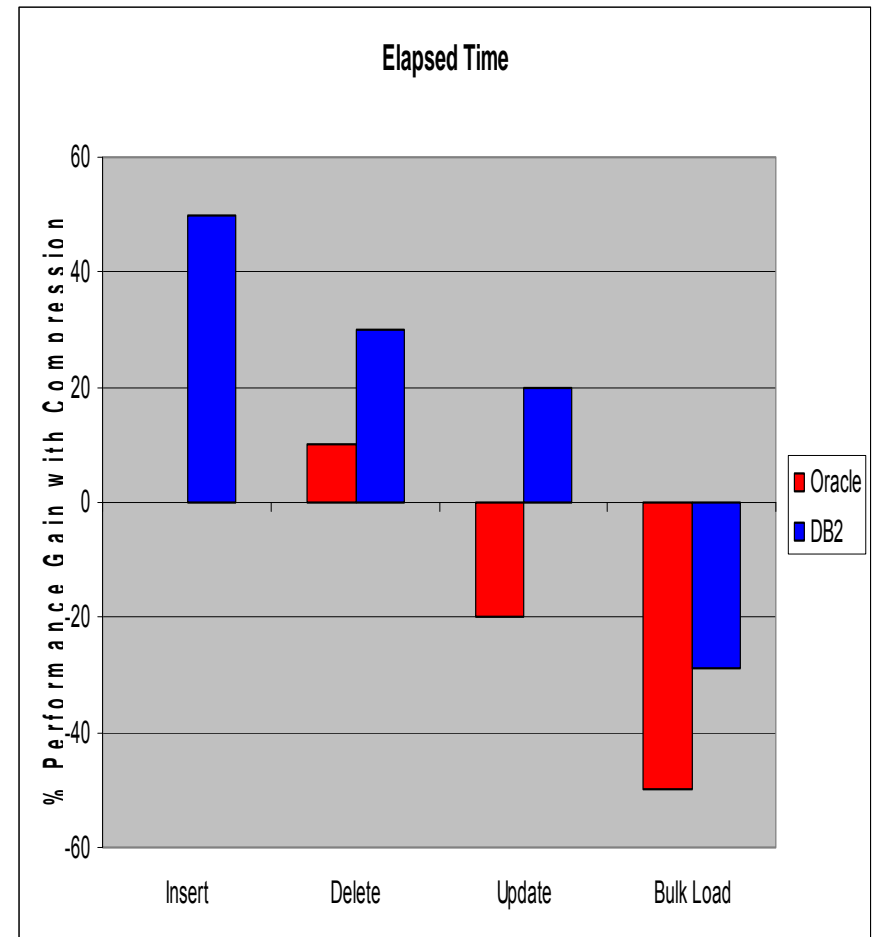
- TPC-H is a well known data warehouse benchmark
 - ▶ Each vendor uses the same tables and same data
 - ▶ Oracle published their compression rates for TPC-H tables at the VLDB conference in 2003
 - ▶ IBM ran the same tests on the same tables
- Test results

Table	Compression Ratio	
	Oracle	DB2
LINEITEM	38%	58% (1.5x better)
ORDERS	18%	60% (3x better)
Entire Database	29%	59% (2x better)

Storage Costs:

DB2 for z/OS Has a Better Performance Impact

- Benefits of hardware compression:
 - ▶ Faster than software compression
 - ▶ Executes in dedicated microprocessors, thus relieving the main processor for other tasks
- Tests from older versions of DB2 for z/OS (v3 and v5) already show superior performance compared to Oracle's current 10g R2 release.
 - ▶ DB2 v3 was released in 1993 and v5 in 1997, using processors older than the current G6
- DB2's hardware compression keeps getting faster with every new generation of hardware
 - ▶ G6 was faster than its precursor
 - ▶ zSeries improved performance significantly, compress or decompress is up to 5 times faster than G6
 - ▶ z9, both BC and EC, are even faster than zSeries



Storage Costs:

DB2 Provides More Storage Savings than Oracle

- DB2 for z/OS lowers TCO by reducing storage use
 - ▶ Typical tables compress in the range of 40% - 80%.
 - ▶ Overall database storage savings average around 30% - 45%.
 - ▶ For certain large database installations 50–80% savings were reported.

- Example storage savings with DB2 vs. Oracle

	Oracle	DB2 v9 for LUW *
Reference price for 3TB of disk	EMC disk on HP Superdome: \$52K	IBM System Storage DS8000: \$60K
Expanded price for 200TB of disk	\$3.5M (\$17.5K per TB)	\$4M (\$20K per TB)
Overall database compression ratio (using TPC-H benchmark results)	30%	60%
Compressed storage cost per TB	$\$17.5K * 70\% = \$12.25K \text{ per TB}$	$\$20K * 40\% = \$8K \text{ per TB}$
Cost of storing 200 TB data compressed	$\$12.3K * 200 = \mathbf{\$2.45M}$	$\$8K * 200 = \mathbf{\$1.6M}$
With compression, storage for DB2 costs <u>35% less</u> than for Oracle		

* DB2 for z/OS achieves similar compression ratios to those of DB2 v9 for LUW

TCO Improvements – DBA tools

- Autonomic Policy-based SQL query management/monitoring:
- Optimization Service Center (Web-based DBA admin – no 3270 screens)
- DBA tool suite for tuning/managing SQL queries (Query Performance Warehouse
 - ▶ Execution history of queries
 - ▶ Identification of query patterns
 - ▶ Identification of usage patterns for tables/indexes
- Resource Limit Facility to save CPU from runaway queries

IBM Tools Offer a Lower Cost Alternative

VENDOR	PRODUCT	IBM PRODUCT	IBM ADVANTAGE
Computer Associates	RC/Migrator	DB2 Administration Tool	Equivalent capability of RC/Migrator and RC/Query
BMC	Mainview	Omegamon Performance Expert & monitor	Industry standard Batch reports DB2 Connect monitoring Web Based Interface
NONE	NONE	Audit Expert	Give auditors access to compliance data and takes DBA audit overhead away

New zIIP Processor from IBM

IBM has recently announced a specialty processor called zIIP for data base workloads



IBM

New zIIP Processor Dramatically Lowers Cost

- IBM System z9 Integrated Information Processor (zIIP)
- New specialty engine for the System z9 mainframe designed to:
 - ▶ Lower the cost of ownership for eligible data base workloads
 - ▶ \$95,000 - \$125,000 one time charge for zIIP processor
 - ▶ No IBM software charges on the zIIP
- DB2 for z/OS V8 is the first IBM exploiter of the zIIP

What Workloads Can Be Run on a zIIP?

- DRDA over a TCP/IP connection
 - ▶ May benefit ERP, CRM, Business Intelligence and other enterprise application solutions

- Parallel queries
 - ▶ May benefit data warehouse solutions


- Internal DB2 utility functions used for index maintenance
 - ▶ Benefit to utility operations

- IBM has tools to help customers estimate their off load potential

XML Solves Business Problems Today...Faster

- Business-to-Business Integration
 - ▶ Platform-independent transport mechanism.

Mortgage applications may be defined in XML




- Forms and Document Processing
 - ▶ Government and legal industry require digital signature

Loan applications forms require signature & change year to year


- ▶ Documents often contain sub-documents

Literary materials contain books, chapters, and sub-chapters



- Business Insight
 - ▶ Universal representation from multiple sources

Risk adjustor reviews risk data from multiple service bureau consideration of original format




- SOA
 - ▶ Web Services messages are XML

XML in DB2

We can now join XML DB2 data with data from heterogeneous data sources?



SQL Person... "I see a world class RDBMS that also supports XML"



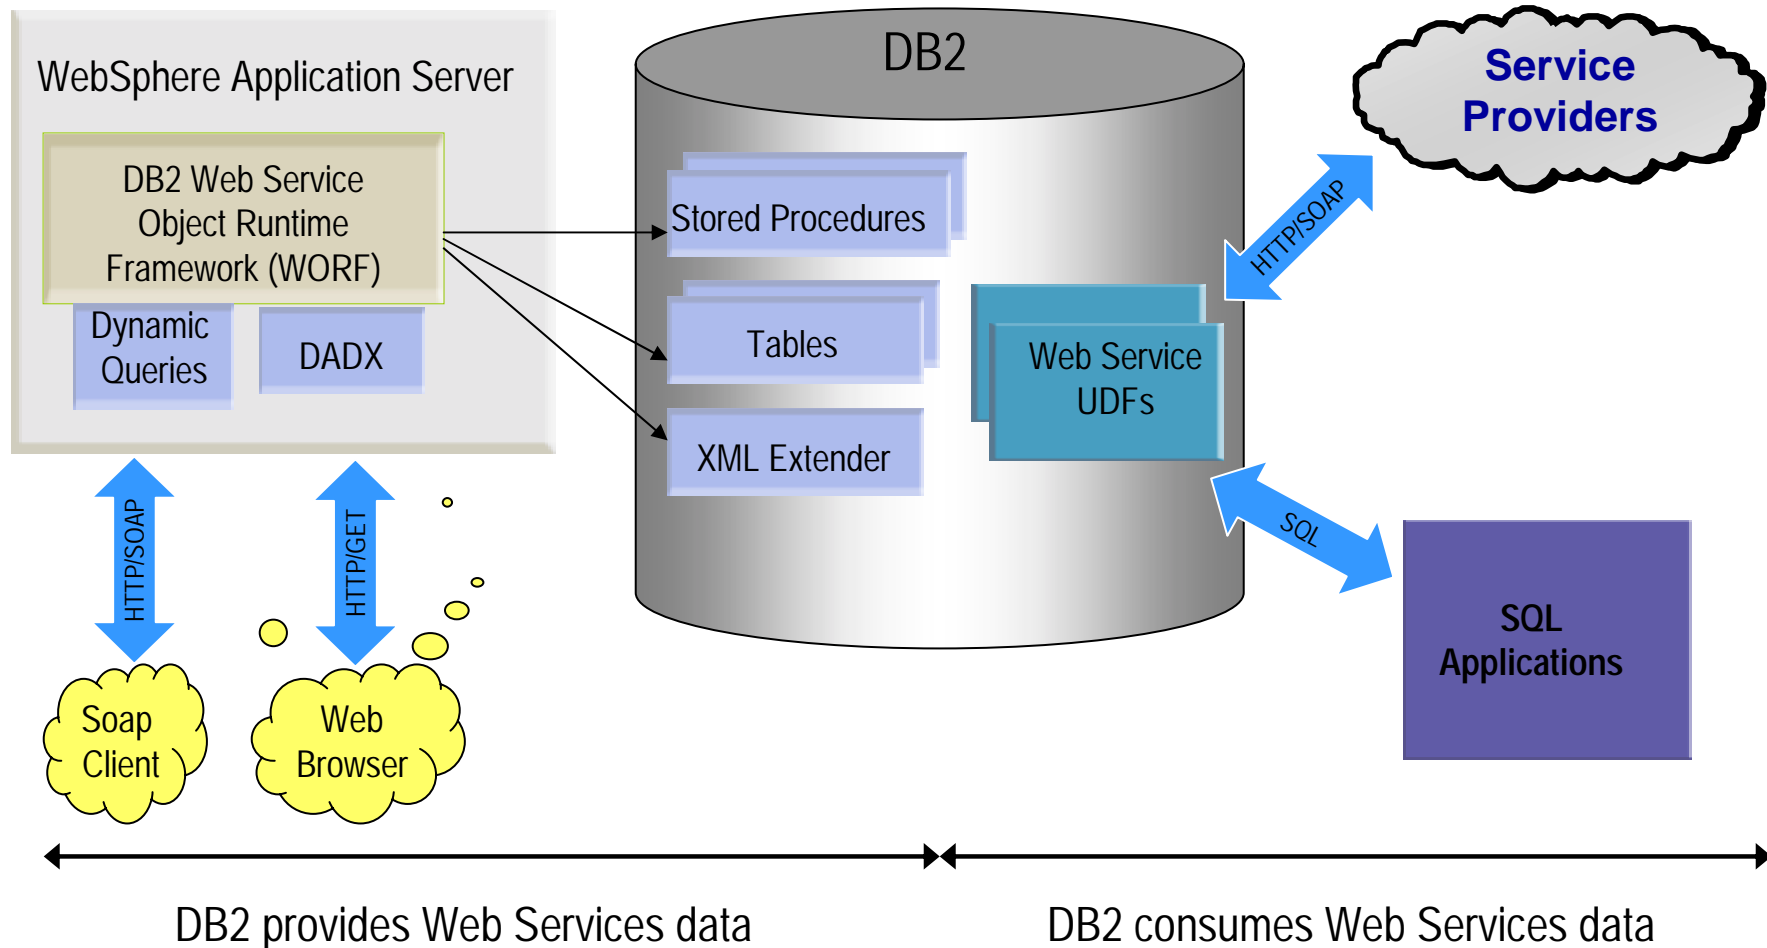
XML Person... "I see a world class XML repository that also supports SQL"

XML integrated in all facets of DB2!

New XML applications benefit from:

- Ability to seamlessly leverage relational investment
- Proven Infrastructure that provides enterprise-class capabilities

DB2 Web Services Supports XML



Proto-type Results Using DB2 9 Based XML/SOA

*Now in Beta
On System z*

	Relational	DB2 9
Development Time	CLOB: 8 hrs Shred: 2 hrs	30 min.
Add field to schema	1 week	5 min.
Relative lines of I/O code	100	35
Queries	24 - 36 hrs	20 sec - 10 min
Query non-shredded XML	1 week	½ day

1-3 M documents inserted per hour on z9

On Demand Insurance

Our sales promotions should be more effective.

We need to do a better job of identifying promotion opportunities.



**On Demand Insurance
CEO**

We have the information, but it's spread across a lot of systems and not easy to analyze.




**On Demand Insurance
CIO**

Data Warehousing and Business Intelligence

- Powerful products to help manage Data
- Extract and load data
 - ▶ Perform ETL operations using DataStage and DataStage TX
- Cleanse Data
 - ▶ WebSphere QualityStage can help cleanse data by identifying duplicate information
- Customized Insight for BI
 - ▶ With Alphablox and Cube Views, business managers can have customized insight into business performance


WebSphere DataStage

- Execute and monitor jobs to extract, transform, and load data
- Graphically design, develop, deploy, and reuse transformations quickly and easily
- Parallel Processing delivers Unmatched Function, Performance, and Scalability
- Deploy jobs in real-time, batch mode, or as services using SOAP, JMS or EJB bindings
- Take advantage of native mainframe processing as needed





Loading Your Data Warehouse with WebSphere DataStage

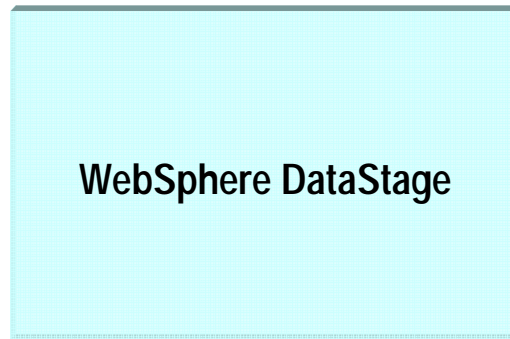
Multiple Databases




Multiple Files



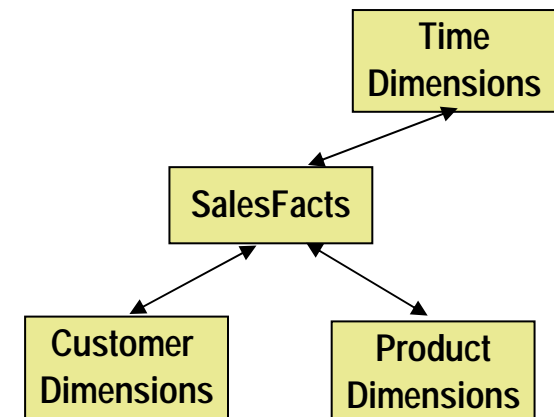

Extract



Transform



Load



Star schema tables

Ensuring the Quality of Your Data

We find a lot of duplications in our databases

We don't want these duplications to show up in our data warehouse



**On Demand Insurance
CIO**

WebSphere QualityStage can help cleanse your data by identifying duplicate information



IBM

Avoiding Duplicate Data: Two Methods to Decide a Match

Are these two records a match?

WILLIAM J	HOLDEN	128	MAIN	ST	02111	12/8/62
WILLAIM JOHN	HOLDEN	128	MAINE	AVE	02110	12/8/62

B B A A B D B A = BBAABDBA

+5 +2 +20 +3 +4 -1 +7 +9 = +49

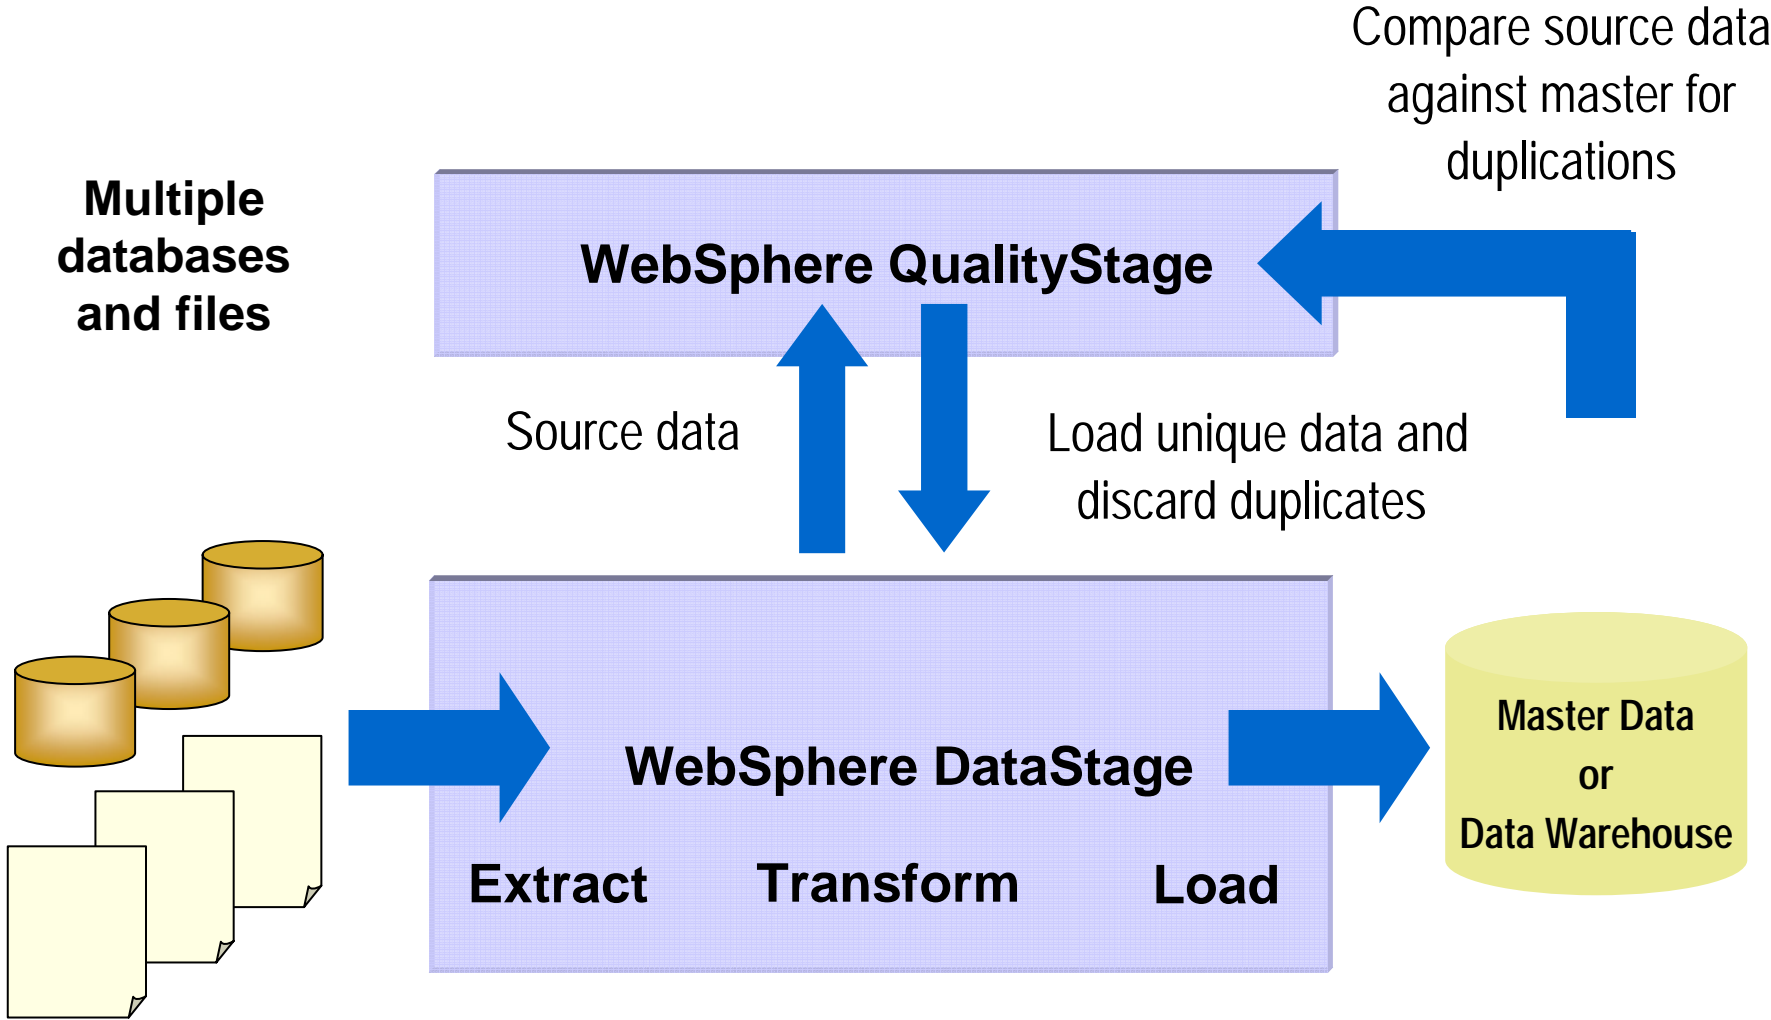
Deterministic Decisions Tables:

- ▶ Fields are compared
- ▶ Letter grade assigned
- ▶ Combined letter grades are compared to a vendor delivered file
- ▶ Result: Match; Fail; Suspect

Probabilistic Record Linkage:

- ▶ Fields are evaluated for degree-of-match
- ▶ Weight assigned: represents the “**information content**” by value
- ▶ Weights are summed to derived a total score
- ▶ Result: Statistical probability of a match

QualityStage Typical Scenario



Real Time Business Intelligence

Now we can start answering those business questions



**On Demand Insurance
CIO**

With Alphablox and Cube Views, business managers can have customized insight into business performance



IBM

Alphablox – Quick, Easy, and Customized Analytical Applications

Customized analytical applications

- Business performance and Key Performance Indicator dashboards
- Self-service reporting and analysis
- Financial reporting and analysis
- Operational analysis and planning

Heterogeneous information sources

- OLTP, OLAP, data warehouses, EIS, others

Pre-built Blox capabilities

- Interactive and real time
- Web-based spreadsheets
- Guided analysis and alerts
- Personalization and customization
- Sharing and collaboration

Web browser interface

- No client to download or install

- Dashboards
- Scorecards
- Portal KPIs
- Reporting
- Real time BI



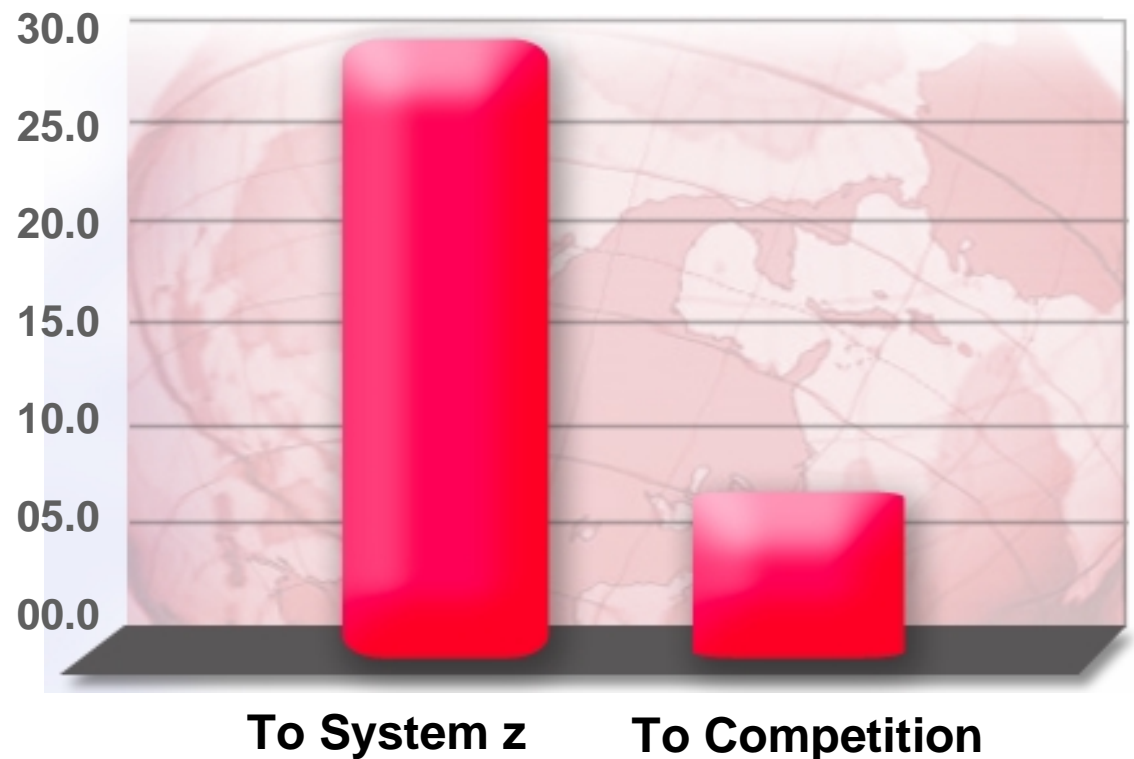
The World Relies on DB2 for System z

- Over 10,000 Licenses World Wide
- Owns 96% of Relational System z Market Place
- Over 3,000 TBs of Production Data
- Over 8 Million Clients
- Over 700 ISV Applications and Packages
- Many, Many More Customer References

More Workload Moving to the Mainframe

DBMS Migration Plans

- Nearly 30% intend to transition DBMS to z/OS®
- Large percentage moving from Microsoft® Windows®
- “This [trend] reflects the robust reputation of the mainframe hardware”
- Amount of data doubles every 12 months



Gartner: CIO Study March, 2006

The World Relies on DB2 for System z – You Can Too

- Over 10,000 Licenses World Wide
- Owns 98% of Relational System z Market Place
- Over 3,000 TBs of Production Data
- Over 8 Million Clients
- Over 700 ISV Applications and Packages
- Many, Many More Customer References

