



Extending Your Mainframe For More Business Value

New Data Workloads On System Z

New Data Workloads

Data is central to our operations and many of our projects

We have a few issues and challenges



**Service Oriented Finance
CIO**

New Data Workloads

Our core processing systems use DB2 for z/OS in a sysplex configuration

Organic growth is increasing our MIPS usage

Oracle says they can do the job for lower cost



Service Oriented Finance
CIO

Oracle falls short compared to DB2

Lets see why the world's largest corporations rely on DB2 for z/OS.




IBM

DB2 Proven Success In The Finance Industry

- 59 of the top 60 banks on the global Fortune 500 list use DB2 for z/OS
- Why?
 - ▶ Highest Scalability – Capacity to handle large or growing workloads
 - ▶ Highest Availability – DB2 provides nearly continuous availability
 - ▶ Proven Security and Compliance – DB2 protects business data and customer privacy
 - ▶ Lowest overall TCO for incremental growth

DB2 For z/OS Has Near Linear Scalability

- IBM benchmarked the workloads of Bank of China and another large Asian bank to demonstrate workload capacity
- Near linear scaling was achieved through a range of MIPS

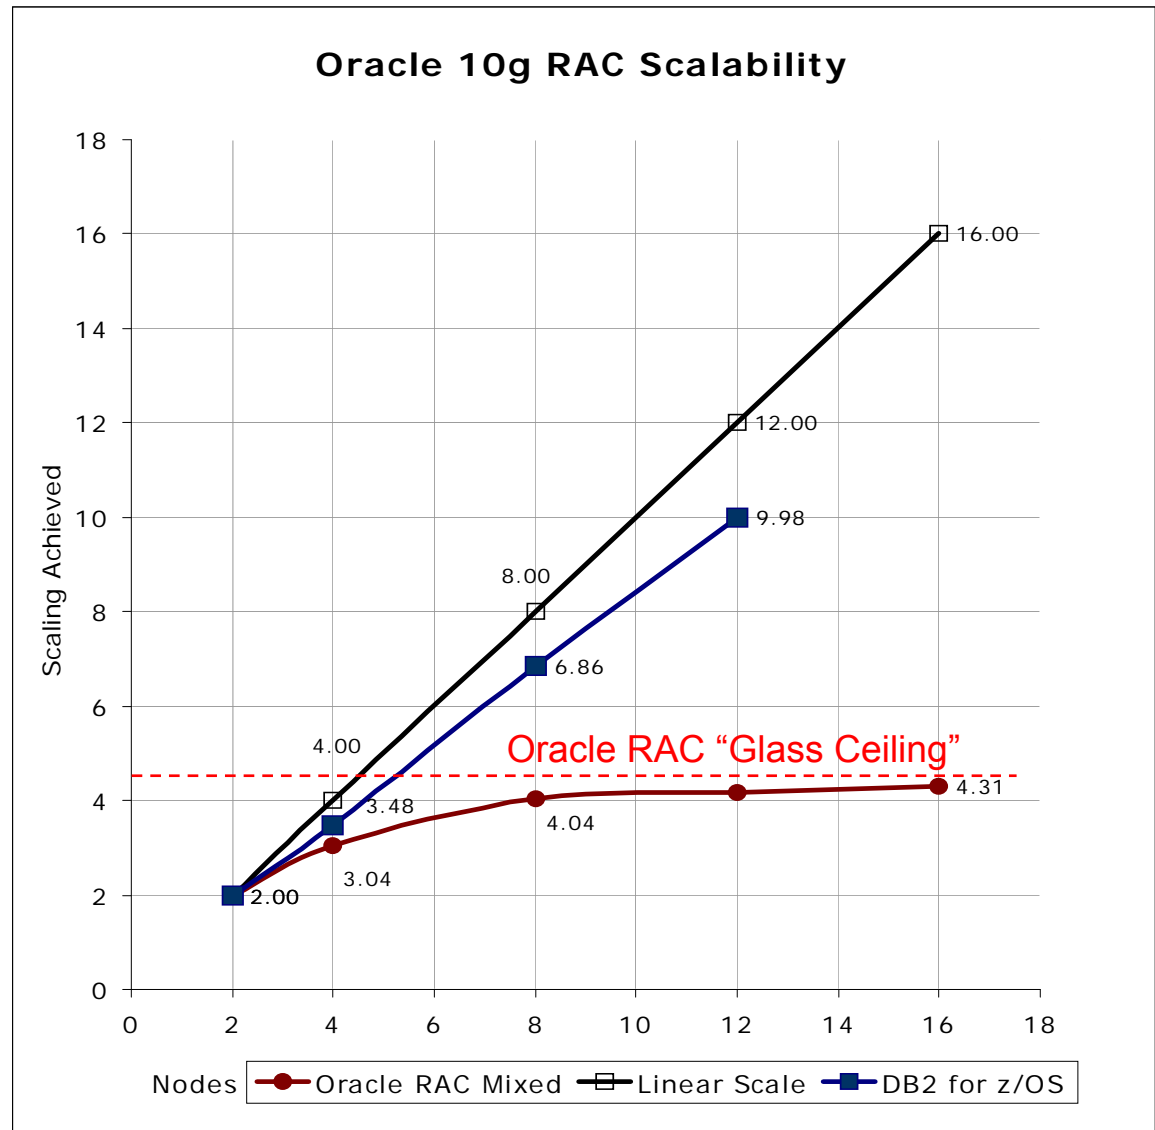


Oracle RAC Scale Out Is Limited

- 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

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


"Enterprise Data Base Clustering Solutions" ITG, October 2003



Why Does Oracle RAC Have These Problems?

DB2 for z/OS


Centralized Sysplex Design



High speed centralized
lock manager in
coupling facility

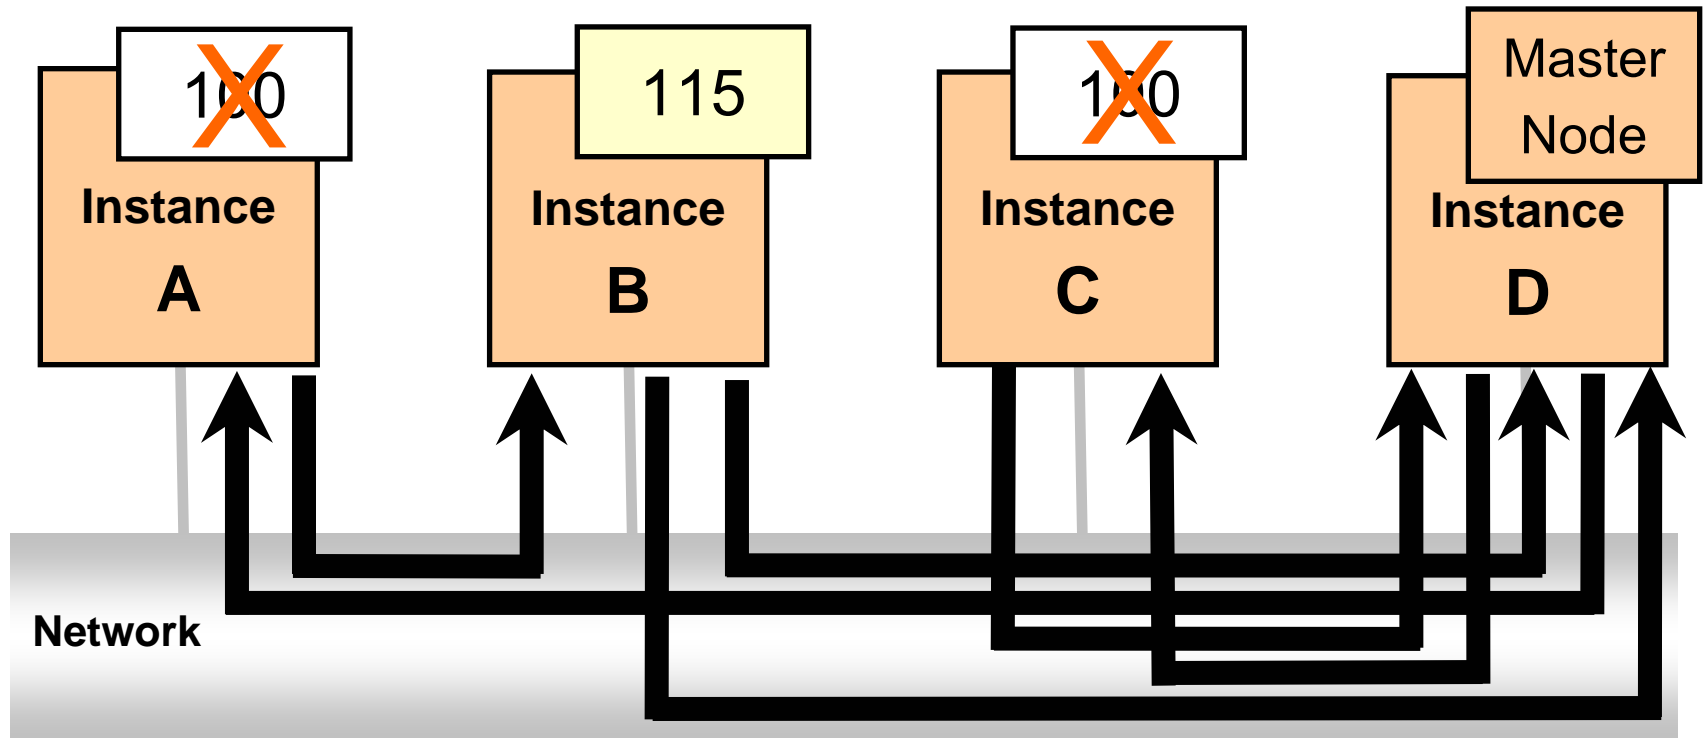
Oracle RAC

Distributed Lock and Data Design



Distributed lock
management with
high messaging **overhead**

Oracle RAC: Lock Management Overhead



Lock Assume

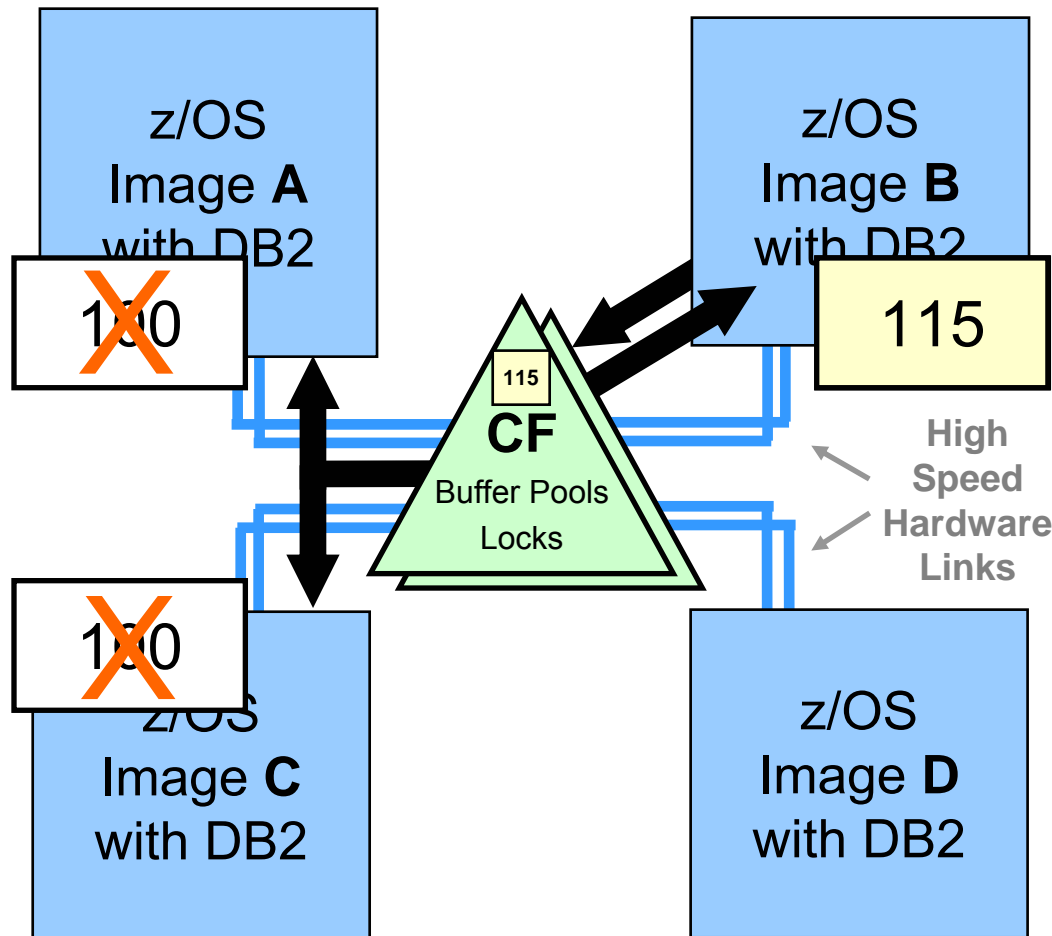
7. B Updates local copy

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.

Centralized Coupling Facility Permits Efficient Lock And Cache Management In DB2



A and C have data in local bufferpool without locks

1. B registers page to CF and obtains write lock
2. B Updates data
3. B Commits update


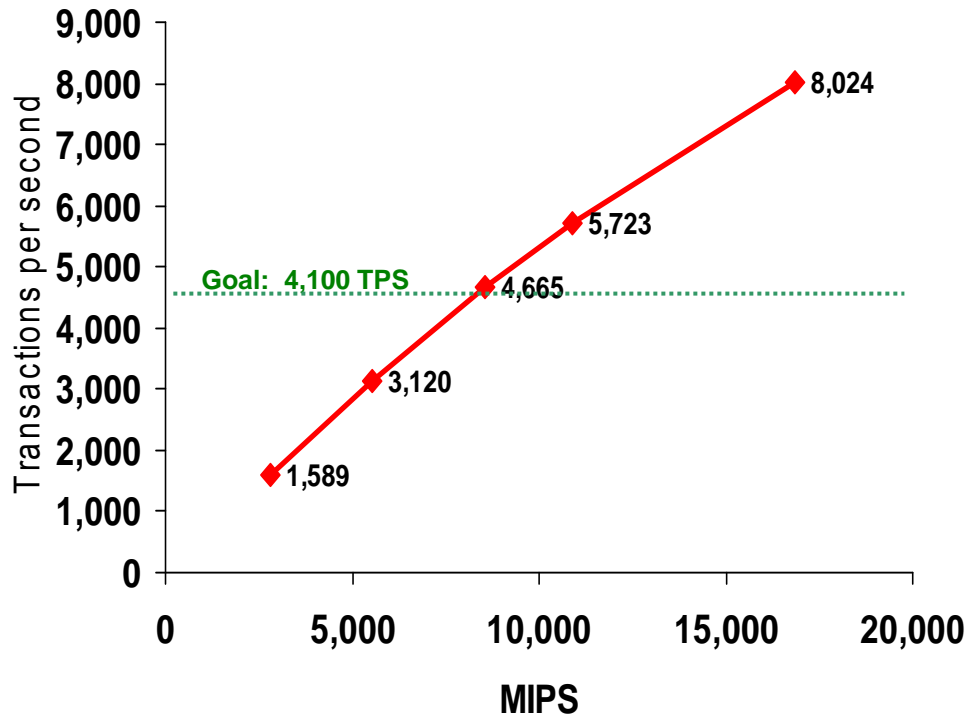
B Caches update in group buffer pool

CF invalidates all cached copies without interrupting processors

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

I/O Bandwidth Is Also Important For Scalability

Bank of China System z Benchmark required big I/O bandwidth capacity



Big scale up, requires big I/O bandwidth capacity

Data Security And Compliance: DB2 For z/OS Has A Proven Track Record

DB2 for z/OS Security

- 10 security related patches in the last 10 years
- Proven RACF and Multi Level Security
- End-to-end encryption via hardware assist
- Optim Test Data Management
 - ▶ Ensures anonymous access to data necessary for testing
- Optim Archiving Expert
 - ▶ Allows customers to easily archive and access data
- DB2 Audit Management Expert
 - ▶ Supports compliance requirements
 - ▶ Consul for enterprise wide audit

Oracle's Security Exposures

- **Oracle.com – April 2008**
41 security patches, including **17** for database
- **eWeek.com – January 2008**
26 security patches, including **9** for database
- **eWeek.com – October 2007**
51 security patches, including **27** for database
- **eWeek.com – July 2007**
45 security patches, including **17** for database

In the last year Oracle has issued 163 security patches, **70** for the database

New Data Workloads

OK, I see that Oracle RAC may not be able to handle our core transaction processing requirements

But what about our other new projects?



**Service Oriented Finance
CIO**

Mainframe Extension Solution – New Data Workloads

We need a data base server for our new **SAP** applications

Our credit report project needs to store **XML** data



Service Oriented Finance
CIO

DB2 9 for z/OS is the best choice for these projects too!



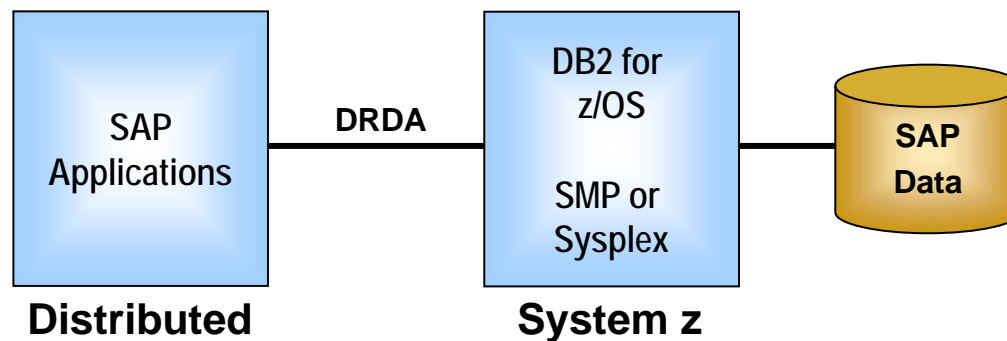
IBM

DB2 For z/OS Is Designed To Work Better With SAP

■ Partnership with SAP

- ▶ 35 years of IBM partnership with SAP, 12,000 joint customers
 - 14 years of DB2 advancements driven by SAP
 - Joint development team
 - technology roadmaps with IBM
- ▶ DB2 for z/OS 9: approximately 40 features requested by SAP
- ▶ Eligible for zIIP and new workload price incentives
- ▶ No unique features in SAP exploit Oracle

■ SAP data operations benefit from the inherent qualities of the mainframe platform



DB2 for z/OS Optimizations for SAP

■ Ease-of-Use

- ▶ Easy to clone DB2 instances, such as test environment

■ Less DBA skills and activities required

- ▶ Large Object Management, Automated Space Management
- ▶ DB2 Recovery Expert for automatic recovery and backup
- ▶ Real-time Statistics Utility provides automatic scheduling information, integration into Workload Management and Resource Limit Facility
- ▶ BACKUP and RESTORE system enhancements

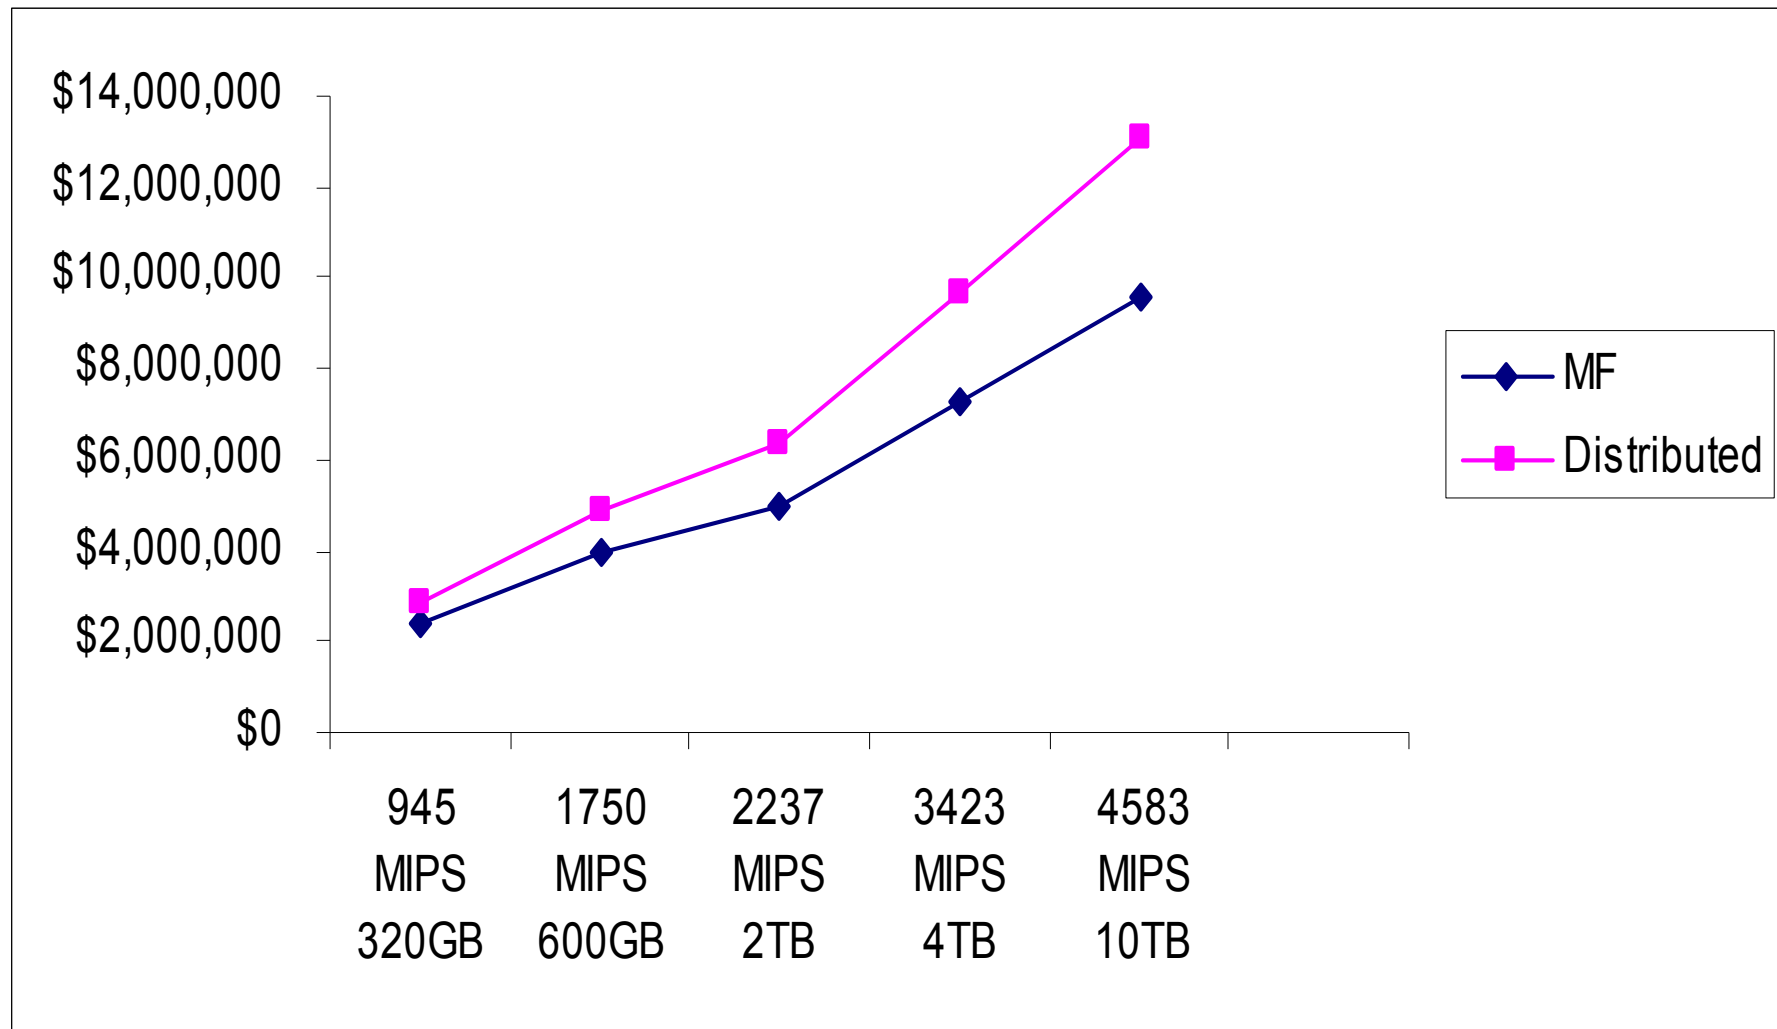
■ SAP exploitation of DB2 9 new features

- ▶ Partition by growth, Merge data and Fast load
- ▶ Simplified connectivity and seamless failover
- ▶ SAP specific Auditing by SAP transaction, report et al
- ▶ Index rename and compression, BIGINT data types

■ High Performance

- ▶ SAP Business Warehouse performance gains through Dynamic Index ANDing

Data Server with Disaster Recovery – Mainframe Costs Are Lower Regardless of Data Server Size



IBM Teams With SAP To Further Lower The Cost Of DB2 For SAP Customers

- OEM agreement allows SAP to sell DB2, DB2 Utilities and DB2 Connect for restricted use

North American Retailer Example



Assume 298
incremental
MSU's dedicated
to DB2 for SAP

	Prior to OEM Agreement	With OEM Agreement
3 Year Costs	\$1,596,997	\$692,561
Savings of over \$900K and 57% for Data Serving on System z!		


But What About The SAP *Applications*?

- Typical configuration
 - ▶ SAP data base server on z/OS
 - ▶ SAP applications on distributed servers

- Better configuration
 - ▶ SAP data base server on z/OS
 - ▶ SAP applications on **zLinux**
 - ▶ Benefit from qualities of mainframe service
 - ▶ Run on lower cost IFL processors
 - ▶ Benefit from co-location of data base and applications
 - ▶ Systematic disaster recovery

Customer Case Study: European Retailer Saves Money By Running SAP Applications On zLinux

- Cost study to replace existing SAP application on Solaris servers. Costs include all incremental Hardware and Software



- Keeping Applications and Database on System z results in
 - ▶ Better workload management and virtualization
 - ▶ Co-location benefit of SAP applications and data bases on same System z

Baldor Electric Company Consolidates Global SAP Systems Onto IBM Mainframe



Solution

- Consolidate 35 global SAP systems to one System z Server
- Portal-based applications extend customer access to inventory systems
- Used zIIPs and IFLs to reduce costs

Results

“The migration of our SAP application servers to Linux on zSeries produced an immediate increase in **performance**, has made it easier to **manage** and maintain our systems, and significantly trimmed the **total cost** of IT”

“Downtime costs us more than **\$100,000** an hour. Availability is king for Baldor, and the IBM zSeries gives us what we need.”

Mark Shackelford,
Director of Information Systems, Baldor

Baldor met customer needs and achieved company growth without a rise in IT costs

XML Solves Business Problems Today


- SOA

- ▶ Web Services messages are XML

- Business-to-Business Integration

- ▶ Platform-independent transport mechanism.

Transaction orders may be defined in XML



- Forms and Document Processing

- ▶ Government and legal industry require digital signature

Tax forms require signature & change year to year

- ▶ Documents often contain sub-documents

Literary materials contain books, chapters, and sub-chapters



XML Is Driving Many Industry Standards Today

Banking

IFX, OFX, SWIFT, SPARCS,
MISMO +++

Healthcare

HL7, DICOM, SNOMED,
LOINC, SCRIPT +++

Insurance

ACORD
XML for P&C, Life +++

Financial Markets


FIX Protocol, FIXML, MDDL,
RIXML, FpML +++

Cross Industry

PDES/STEPml
SMPI Standards
RFID, DOD XML+++

Life Sciences

MIAME, MAGE,
LSID, HL7, DICOM,
CDIS, LAB, ADaM +++



Chemical & Petroleum

Chemical eStandards
CyberSecurity
PDX Standard+++

Retail

IXRetail, UCCNET, EAN-UCC
ePC Network +++

Electronics

PIPs, RNIF, Business Directory,
Open Access Standards +++

Telecommunications

eTOM, NGOSS, etc.
Parlay Specification +++

Energy & Utilities

IEC Working Group 14
Multiple Standards
CIM, Multispeak

Service Oriented Finance Needs To Store XML Data

We need to support the MISMO standard to do credit checks. It uses XML.



**Service Oriented Finance
CIO**

DB2 9 pureXML can do this.

Let's see how...



IBM

XML – The Difference Is Fundamental

- Relational is a data model
 - Relations (tables)
 - Attributes (columns)
 - Set based w/some sequences
 - Strict schema

SSN	CreditReportID	CreditDate
111111111	1234	Dec 12, 2007
111111111	4456	Feb 8, 2008
123456789	2314	Nov 30, 2007

SSN	LastName	FirstName	Street	City	State	Zip
111111111	Haan	Brian	1 Harry Rd	San Jose	CA	95141
123456789	Smith	Joe	555 Bailey Ave	San Jose	CA	95141

CreditReportID	CreditBureau	CreditLiability	Rating
1234	ABC Credit	Collection	649
1235	ABC Credit	Collection	687
2314	TRW Reporting	Mortgage	750

- XML is a data model
 - Hierarchical tree structure
 - Nodes (elements, attributes, comments, etc.)
 - Relationships between nodes
 - Sequence based w/ some sets
 - Flexible schema

```
<MISMOVersionID="2.3.1" ?>
<RespondingParty="ABC Credit">
  <RESPONSE_DATA>
    <CREDIT_RESPONSE_MISMOVersionID="2.3.1"
      CreditResponseID="CRResp0001"
      CreditRatingCodeType="Equifax">

      <CREDIT_BUREAU _Name="ABC Credit" _StreetAddress="..
      ...
    </CREDIT_BUREAU>
    <BORROWER BorrowerID="B1" _FirstName="Joe" _LastName="Smith"
    ....
  </BORROWER>
  <CREDIT_LIABILITY CreditLiabilityID="CrL12923"
  ....
    <_CURRENT_RATING _Code="9" _Type="Collection/>
  ....
</RESPONSE_DATA>
```


Service Oriented Finance Needs To Store XML Data

```

<MISMOVersionID="2.3.1" ?>
<RespondingParty="ABC Credit">
<RESPONSE_DATA>
  <CREDIT_RESPONSE_MISMOVersionID="2.3.1"
    CreditResponseID="CRResp0001"
    CreditRatingCodeType="Equifax">

  <CREDIT_BUREAU_Name="ABC Credit" _StreetAddress="..
    ...
  </CREDIT_BUREAU>
  <BORROWER BorrowerID="B1" _FirstName="Joe" _LastName="Smith"
    ....
  </BORROWER>
  <CREDIT_LIABILITY CreditLiabilityID="CrL12923"
    ....
    <_CURRENT_RATING _Code="9" _Type="Collection/>
  ....
</RESPONSE_DATA>
  
```

BEST!
Performance
Sparse Data
Schema
Evolution

File Name

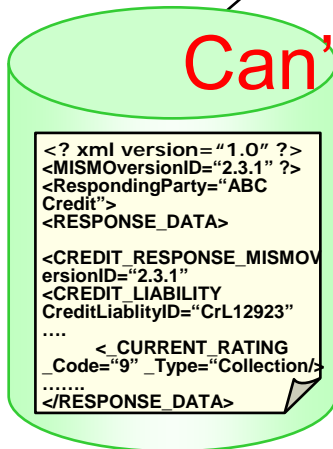
SQL XML

SQL XML

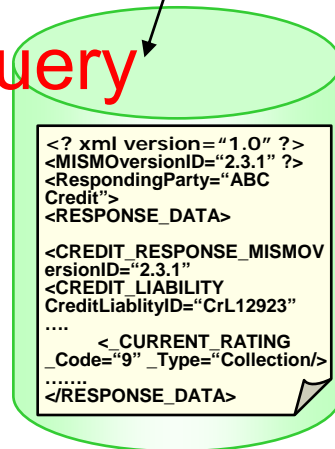
pureXML

Can't Query


Overhead



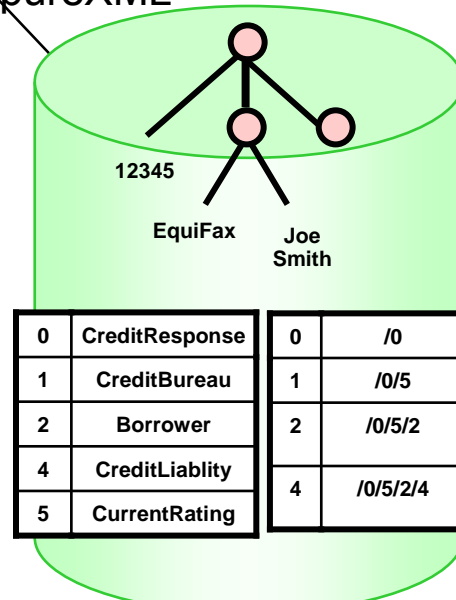
File System



**BLOB
(Oracle)**



**Shredded
(Oracle)**



**Native Support
(DB2)**

DB2 9 Native XML Storage

- A “Hybrid” data base environment combining the relational and XML hierarchical data models
 - ▶ Adds a new “XML” data type
- A new storage mechanism to efficiently manage XML data
 - ▶ “Native“ means that XML documents are stored on data base pages as parsed tree structures to reflect XML’s hierarchical structure
- This avoids conversions between XML and relational structures, and the corresponding limitations
 - ▶ Input and retrieval are faster, performance is better, and querying is better and faster
 - ▶ With BLOBs and shredding, every operation (parsing, etc.) is expensive and there is a potential loss of data
 - ▶ The XML document might be too complex to shred

DEMO: Service Oriented Finance Credit Report Processing

- Data base contains two credit reports for Brian Haan
- Schema of one report is old version
- Schema of the other report is up-level version
- New schema contains a new element (high risk loans)
- Same query can access both

