

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

02 - New Data Workloads on System z v1.7

3

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

02 - New Data Workloads on System z v1.7

4

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



02 - New Data Workloads on System z v1.7


5

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




02 - New Data Workloads on System z v1.7

6


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

02 - New Data Workloads on System z v1.7

8

Oracle RAC: Lock and Cache Management Overhead



Lock Assume

7. B Updates local copy

Inter-node connections: 6


Communications overhead limits Oracle RAC scalability

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

02 - New Data Workloads on System z v1.7

9

Mainframe Centralized Coupling Facility Permits Efficient Lock and Cache Management in DB2



A and B have read locks with local copies

1. B registers page to CF and obtains write lock
2. B Updates local copy
3. B Caches update in group buffer pool
4. CF invalidates all cached copies without interrupting processors


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

02 - New Data Workloads on System z v1.7

10

I/O Bandwidth Is Also Important For Scalability

Bank of China Benchmark required huge I/O bandwidth capacity



Huge scale up, requires huge I/O bandwidth capacity


02 - New Data Workloads on System z v1.7

11

Oracle RAC I/O Bandwidth Is Limited By The Underlying Server Platform

An attempt at 8+ Gbytes per second of I/O

16 - 4 Processor Servers with 8 Fiber Ports
 8 Fiber Switches
 64 - Fiber Arrays (1 port each)
 960 - 72G 15K Disk Drives



ORACLE®


Source: Oracle Open World, Sept 2005 - Session 1973

02 - New Data Workloads on System z v1.7

12

DB2 for z/OS Availability

- Fractional Improvements Result in Millions in Savings
- Financial Impact of Downtime Per Hour for financial industry is \$1.145M
- Financial Services Company Example:
 - ▶ \$300B assets, 2500+ branches, 15M customers
 - ▶ Retail banking, loans, mortgages, wealth management, credit cards
 - ▶ CRM System – branches, financial advisors, call centers, internet
 - ▶ Number of users – 20,000+



	Unix/Oracle	zSeries/DB2
Availability %	99.825%	99.975%
Annual outage	15h 20m	2h 11m
Cost of Downtime	\$17.6M	\$2.5M

\$15.1 Million dollar difference!

Sources: Picking up the value of PKI: Leveraging z/OS for Improving Manageability, Reliability, and Total Cost of Ownership of PKI and Digital Certificates by Jerald Murphy: 2007

02 - New Data Workloads on System z v1.7

13

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

- **eWeek.com – 10/16/2007**
51 security patches, including **27** for data base
 - **eWeek.com – 07/17/2007**
45 security patches, including **17** for data base
 - **eWeek.com – 04/17/2007**
36 security patches, including **13** for data base
 - **C/NET – 01/17/2007**
"Oracle plugs 51 security flaws" including **26** for data base
- In the last 15 months Oracle has issued more than 146 security patches

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

We're also thinking about how we can leverage **spatial** data in our business



Service Oriented Finance
CIO

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




IBM

02 - New Data Workloads on System z v1.7

16

DB2 for z/OS is Designed to Work Better With SAP

- Partnership with SAP
 - ▶ 35 years of IBM partnership with SAP, 12000 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



02 - New Data Workloads on System z v1.7

17

DB2 for z/OS Optimizations for SAP


Examples:

- Ease-of-Use
 - ▶ Easy to clone DB2 instances, such as test environment
 - ▶ Customized SAP Tuner and Administration Console
- 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-specific enhancements to DB2 Query Optimizer
 - ▶ Enhancements for SAP Business Intelligence query performance
 - ▶ Enhancements for SAP OLTP products
- High Performance
 - ▶ SAP Business Warehouse performance gains through Dynamic Index ANDing

02 - New Data Workloads on System z v1.7

18

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



02 - New Data Workloads on System z v1.7

22

IBM Teams with SAP to 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!		

02 - New Data Workloads on System z v1.7

23

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

02 - New Data Workloads on System z v1.7

24

Customer Case Study: European Retailer Saves Money by Running SAP Applications on zLinux

- Cost study to replace existing SAP application on Solaris servers
 - ▶ CASE 1: Applications and data bases on distributed
 - 5 year TCO €15.0M
 - ▶ CASE 2: Applications on distributed, data base on z/OS
 - 5 year TCO €12.6M
 - ▶ CASE 3: Applications on zLinux, data base on z/OS
 - 5 year **TCO €11.1M**
 - Better workload management and virtualization
 - Co-location benefit of SAP applications and data bases on same System z
- All cases incremental cost of additional Hardware and Software

02 - New Data Workloads on System z v1.7

25

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

02 - New Data Workloads on System z v1.7


26

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*



02 - New Data Workloads on System z v1.7

27

XML is Driving Many Industry Standards Today

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


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

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

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

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

Insurance
ACORD
XML for P&C, Life +++



Telecommunications
eTOM, NGOSS, etc.
Parlay Specification +++

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

Automotive
ebXML,
other B2B Stds.

Energy & Utilities
IEC Working Group 14
Multiple Standards
CIM, Multispeak

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

Chemical & Petroleum
Chemical eStandards
CyberSecurity
PDX Standard+++

02 - New Data Workloads on System z v1.7

28

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...

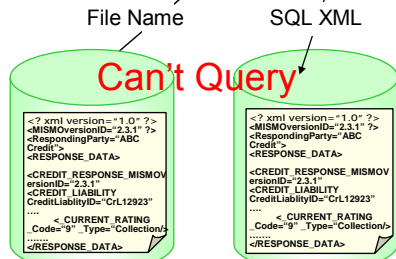


IBM

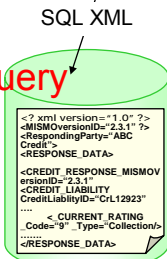
Service Oriented Finance Needs to Store XML Data

```
<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 System



BLOB (Oracle)



Shredded (Oracle)



Native Support (DB2)

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="CRRsp0001"
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>
    
```

02 - New Data Workloads on System z v1.7

31

DB2 9 Native XML Storage

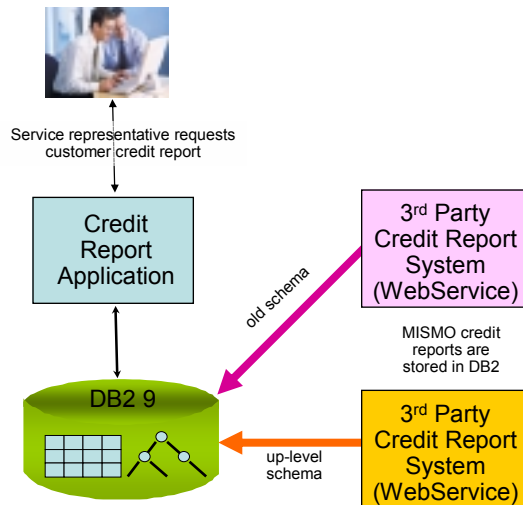
- A “Hybrid” data base environment combining the relational and XML hierarchical data models
 - ▶ Adds a new “XML” data type
- A new storage mechanism to efficiently manage XML data
 - ▶ “Native” means that XML documents are stored on data base pages as parsed tree structures to reflect XML’s hierarchical 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

02 - New Data Workloads on System z v1.7

32

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



02 - New Data Workloads on System z v1.7

34

Financial Services Need Location Data

Business areas that need location data:

- Workers Compensation
- Risk Prediction
- On-Demand Claims
- Service Optimization



02 - New Data Workloads on System z v1.7

35

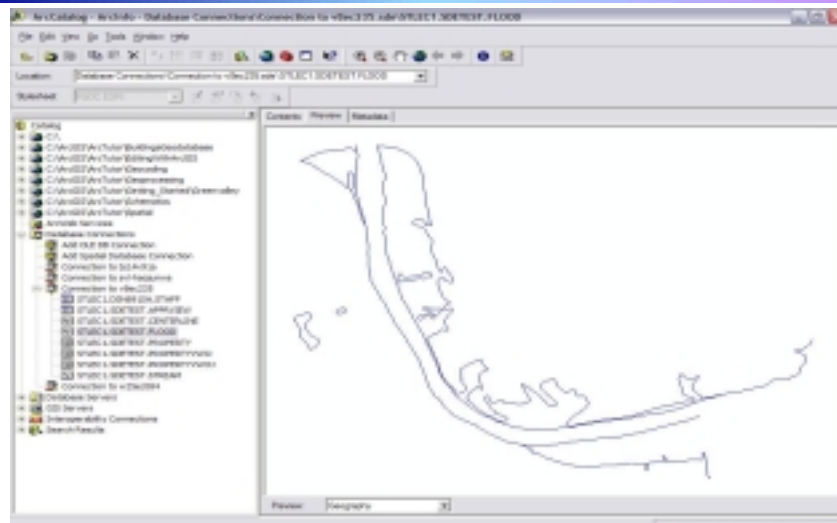
Example: Risk Prediction Needs Location Data

- Every insured RISK has a location
 - ▶ House, office, vehicle, warehouse, person, goods, etc
- This location can change
 - ▶ Vehicle, property, policy, restructuring, merger, acquisition, etc
- Every PERIL influencing the risk can be geolocated
- The perils are influenced by geography
 - ▶ Urban development, population, demographics, climate, postal units, flood, fire, crime, earthquake, tsunamis, landslip, etc
- The locations impacted change over time

02 - New Data Workloads on System z v1.7

36

Show Flood Plain

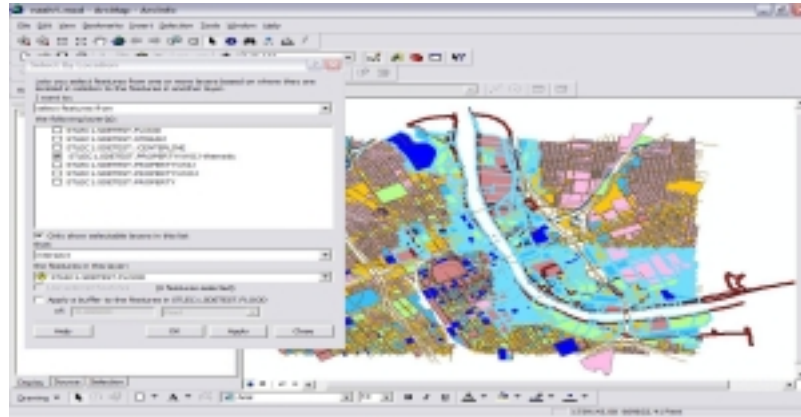


Screen shows ArcCatalog with the preview of flood data set

02 - New Data Workloads on System z v1.7

37

Now Intersect with Our Customers in Flood Zone



Screen shows the query result for "categorized properties intersect flood"

DB2 Simplifies Processing With Spatial Queries

Sample Spatial Query

```
SELECT branchid ,count(*), avg(income)
FROM branches b, customers c
WHERE ST_Distance(b.loc, c.loc)<100
GROUP BY branchid;
```

BRANCHES

BRANCHID	NAME	ADDR	LOC	ZONE

"tell me the average income, and number of all customers who live within 100 miles of each branch"

CUSTOMERS

CID	NAME	INCOME	ADDR	LOC

