



# **Extending Your Mainframe For More Business Value**

Add New Workload –  
Data Servers On System z

# 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!**



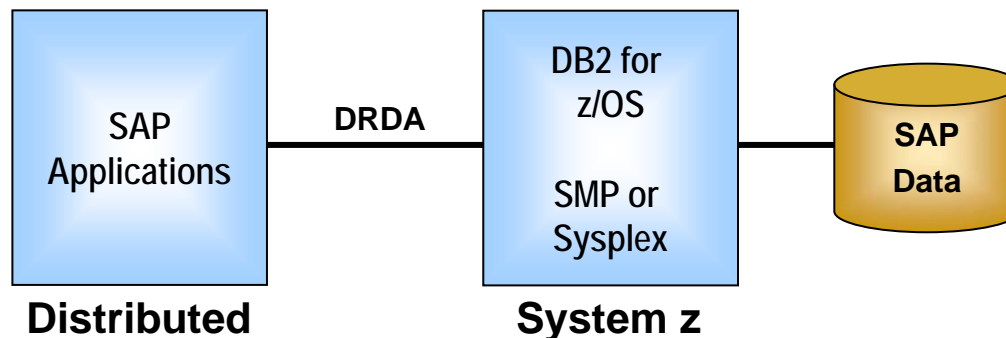
**IBM**

# 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

# Case Study: Consolidate SAP Data Server on Mainframe

## Existing Mainframe




Existing z10:  
2 GP 1,720 MIPS  
DB2 and utilities  
With 20TB storage

## Existing Disaster Recovery Site



Existing:  
1 GP processor for hot  
disaster switch-over  
1 "dark" DR  
processors  
With 20TB storage

## Add 1 LPAR for New SAP Data Server w 1TB Storage



Prod  
2,300 MIPS  
additional  
workload

Incremental:  
2 GP 1,380 MIPS (60%)  
1 zIIP 920 MIPS (40%)  
1 GB memory


## And Add Disaster Recovery w 1TB Storage



3 year  
cost of  
acquisition  
\$5.07 M

Capacity Backup:  
1 GP  
1 zIIP

## Or Add HP Integrity Superdome 9140N Servers w 1.55 TB storage



Prod  
28 Chip  
56 Core  
  
201,977\*  
Performance Units

## And Add 1 Server for Disaster Recovery, Development & QA w 1.55 TB storage



Prod  
28 Chip  
56 Core  
  
201,977\*  
Performance Units

3 year  
cost of  
acquisition  
\$10.77 M

\*Production Performance Units required = 2,300 x 87 = 200,100

# 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

- Solaris cost study to replace existing SAP application
  - ▶ 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

# 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



# 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 Solves Business Problems Today

## ■ SOA

- ▶ Platform independence
- ▶ Web Service messages are XML



## ■ Business-to-Business Integration

- ▶ Any to any communication
- ▶ Information exchange may be defined in XML

## ■ Forms and Document Processing

- ▶ Forms may have sparse data and change format often
- ▶ Documents often have a hierarchical structure
- ▶ Forms and documents may be represented in XML



# 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+++

# XML – The Difference Is Fundamental

- Relational is a data model
  - Relations (tables)
  - Attributes (columns)
  - Set based w/some sequences
  - Strict schema
- ▶ Normalized data
  - Data replication limited to KEY fields

SSN	CreditRespons
111111111	1234
111111111	3456
123456789	2314

CreditRespons	CreditBureau	Rating
1234	ABC Credit	649
3456	ABC Credit	687
2314	TRW Reporting	750


SSN	LastName	FirstName
111111111	Smith	Joe
123456789	Haan	Brian

- XML is a data model
  - Hierarchical tree structure
  - Nodes (elements, attributes, comments, etc.)
  - Relationships between nodes
  - Sequence based w/ some sets
  - Flexible schema
- ▶ Data NOT normalized
  - Data can having repeating elements

```
<CREDIT_RESPONSE>
  </CreditResponseID="1234">

  <ResponseData>
    <CREDIT_BUREAU> ABC Credit </CREDIT_BUREAU>
    <BORROWER> Joe Smith </BORROWER>
    <CREDIT_LIABILITY>
      </CURRENT_RATING=" 649">
    </CREDIT_LIABILITY>

  </ResonseData>
</CREDIT_RESPONSE>
```



# 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

# Service Oriented Finance Needs To Store XML Data

```

<CREDIT_RESPONSE>
  </CreditResponseID="1234">

  <ResponseData>
    <CREDIT_BUREAU> ABC Credit </CREDIT_BUREAU>
      <BORROWER> Joe Smith </BORROWER>
    <CREDIT_LIABILITY>
      </CURRENT_RATING=" 649">
    </CREDIT_LIABILITY>

  </ResonseData>
</CREDIT_RESPONSE>
  
```

**BEST!**  
 Performance  
 Sparse Data  
 Schema  
 Evolution  
 pureXML

File Name

SQL XML

SQL XML

**Can't Query**

**Overhead**

```

<CREDIT_RESPONSE>
</CreditResponseID="1234">
<ResponseData>
  <CREDIT_BUREAU>ABC Credit.
</CREDIT_BUREAU>
<BORROWER>Joe Smith
</BORROWER>
  <CREDIT_LIABILITY ....
  </CURRENT_RATING=" 649">
  .....
</ResponseData>
</CREDIT_RESPONSE>
  
```

```

<CREDIT_RESPONSE>
</CreditResponseID="1234">
<ResponseData>
  <CREDIT_BUREAU>ABC Credit.
</CREDIT_BUREAU>
  <BORROWER>Joe Smith
</BORROWER>
  <CREDIT_LIABILITY ....
  </CURRENT_RATING=" 649">
  .....
</ResponseData>
</CREDIT_RESPONSE>
  
```

1234	ABC Credit	Joe Smith	649

0	CREDIT_RESPONSE
1	CreditResponseID
2	ResponseData
4	CREDIT_BUREAU
5	BORROWER
6	CREDIT_LIABILITY
7	CURRENT_RATING

/0/1 1234 /0/2/4/ ABC Credit /0/2/5 Joe Smith /0/2/6/7 649

**File System**

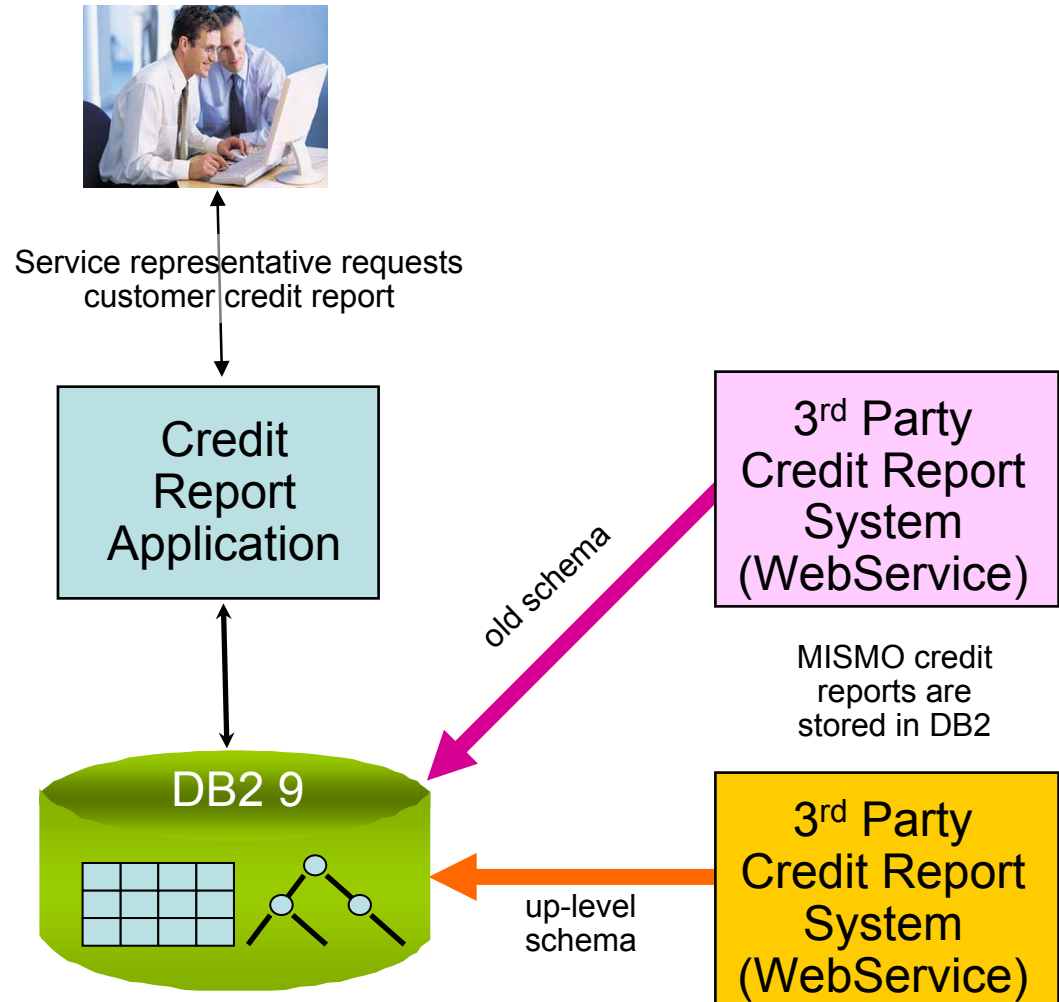
**BLOB  
(Oracle)**

**Shredded  
(Oracle)**

**Native Support  
(DB2)**

# 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



# Remember This Extension Solution? Add an SAP Or XML Database For Less

*Existing Mainframe*




Existing z10:  
2 GP 1,720 MIPS  
DB2 and utilities  
With 20TB storage

*Existing Disaster Recovery Site*



Existing:  
1 GP processor for hot disaster switch-over  
1 "dark" DR processors  
With 20TB storage


*Add 1 LPAR for New SAP Data Server w 1TB Storage*



2,300 MIPS additional workload

Incremental:  
2 GP 1,380 MIPS (60%)  
1 zIIP 920 MIPS (40%)  
1 GB memory

*And Add Disaster Recovery w 1TB Storage*



3 year cost of acquisition  
**\$5.07 M**

Capacity Backup:  
1 GP  
1 zIIP

*Or Add HP Integrity Superdome 9140N Servers w 1.55 TB storage*



201,977\*  
Performance Units

*And Add 1 Server for Disaster Recovery, Development & QA w 1.55 TB storage*



201,977\*  
Performance Units

3 year cost of acquisition  
**\$10.77 M**

\*Production Performance Units required = 2,300 x 87 = 200,100



