





System z Enables Solutions For A Smarter Planet

Handling The Information Explosion

Data Volumes Are Exploding ...

- Information is doubling every 18 months
 - ▶ Structured data growing at 32%
 - ▶ Unstructured data growing at 63%
 - ▶ Replicated data growing at 49%
- IDC predicts by 2011, digital data will be ten times its size in 2006
- We now create more data than we can store
 - ▶ By 2011, half of the data created will not have a permanent home



Sources: IDC worldwide enterprise disk in Exabytes from "Changing Enterprise Data Profile", December 2007 and "The Diverse and Exploding Digital Universe", March 2008

Massive Amounts Of Data Present Challenges To Datacenters

- Lots of data means lots of devices
 - ▶ Growing costs for hardware and management
- Data comes in different forms
 - ▶ High value online
 - ▶ Older, infrequently referenced
 - ▶ Archived
- Requirements to move data
 - ▶ Jobs that process data need high I/O bandwidth
 - ▶ Synchronization requirements

Too Much Data!

Business growth and new regulations on data retention have our data growing out of control!




**Service Oriented Finance
CIO**

You need a smarter information management strategy.




IBM

A Smarter Information Management Strategy




Building A Scalable, Cost Effective Storage Environment Is The First Step



System z Storage Management Is Designed To Handle Massive Amounts Of Data

- System z Data Facility Storage Management Subsystem (DFSMS)
- System z dedicated I/O hardware offloads I/O processing cycles
- IBM System Storage DS8000 and IBM Virtualization Engine TS7700 virtualize storage and deliver massive capacity
- System z integrates these capabilities to deliver optimized storage




System z DFSMS Storage Management

- Provides System z file system and access methods
 - ▶ E.g. BSAM, QSAM, VSAM, z/OS Unix file system ...
 - ▶ Extendable while running

- Storage management features
 - ▶ Automates management of datasets, catalogs, objects and z/OS UNIX files and logical volumes
 - ▶ Move, copy, backup, recovery and automatic space management
 - ▶ Manage removable media
 - ▶ Manages movement of data in storage hierarchy
 - ▶ Concurrent access of VSAM data


Hierarchical Storage Management (HSM) Autonomically Migrates Data For Archival



- References to data typically diminish over time
- Hierarchical storage management moves older data to slower devices
- DFSMShsm provides automated hierarchical storage management for System z
 - ▶ But distributed servers require a separate product like Tivoli Storage Manager for Space Management


System z Also Has Dedicated I/O Hardware To Enhance Performance

- Offload I/O operations to dedicated hardware
- Up to 11 System Assist Processors (SAP) coordinate I/O requests
- Up to 336 RISC processors handle I/O operations
- I/O Offload saves general purpose CPU cycles
- Maximum I/O Bandwidth of 288 GB/sec without impact to workload capacity
- HP Superdome uses general processors for I/O – no dedicated processors
 - ▶ Sustained I/O bandwidth less than half, while impacting workload



Virtualization Of I/O Enables Redundant I/O Paths

- I/O Virtualization provided by Logical Channel Subsystem
 - ▶ Up to 1024 logical channel paths
- Virtualization enables optimal Physical I/O path to be used
 - ▶ Dynamic path selection
 - ▶ Load balances I/O traffic
- Transparent Failover
 - ▶ SAP recovers I/O operations in progress and switches to alternate path



Solid State Disk Drives Are Here To Revolutionize Storage

- Semiconductor storage delivered in DS8000 storage subsystems
 - ▶ Random access solid state storage – no moving parts
 - ▶ Electronically erasable medium
- 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 two orders of magnitude higher than traditional spinning disk
 - ▶ Reduce the “batch window”
 - ▶ DFSMS automatically controls allocation of new datasets to SSD drives
- Cost reductions
 - ▶ 75% reduction in space
 - ▶ 80+% reduction in power and cooling
 - ▶ Reduce RAM requirements

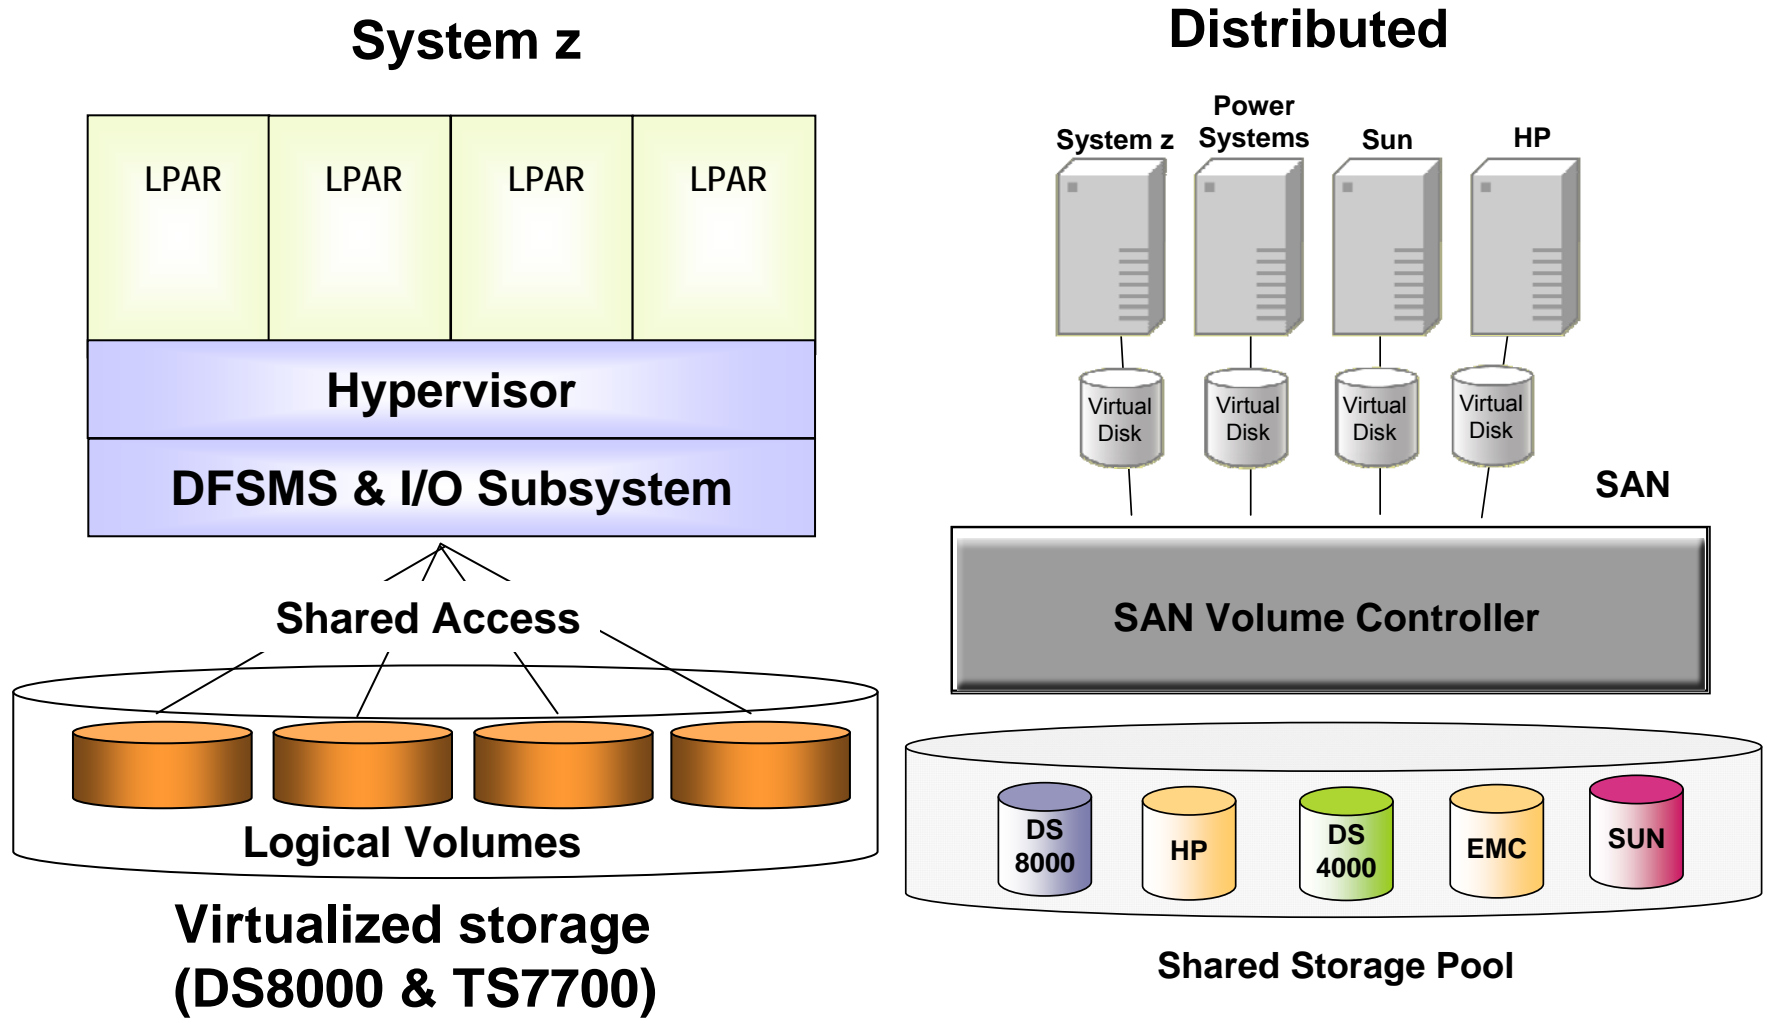


IBM DS8000 And TS7700 Provide High Capacity Storage For System z

- DS8000 supports a mix of disk drive types up to 461 TB
 - ▶ Maximum of 1024 disk drives
 - ▶ Solid State Disk drives (146GB)
 - ▶ 450 GB Fiber Channel Hard Disk Drives (450 GB)
- Up to 4.9 million I/O Operations per second
- Stripe data across multiple RAID arrays
 - ▶ Minimize disk “hot spots”
- Data mirroring for business resilience
 - ▶ Synchronous copies up to 300 km apart
 - ▶ Asynchronous copies over virtually unlimited distances
- Supports System z Extended Address Volume
 - ▶ Up to 223 GB per volume
- Supports Dynamic Volume Expansion
 - ▶ Increase volume size while running
- TS7700 provides virtual tape solution
 - ▶ Up to 70TB disk cache and 11PB capacity with TS3500 Tape Library



Storage Virtualization Is Built Into System z... Distributed Solutions Need Additional Products



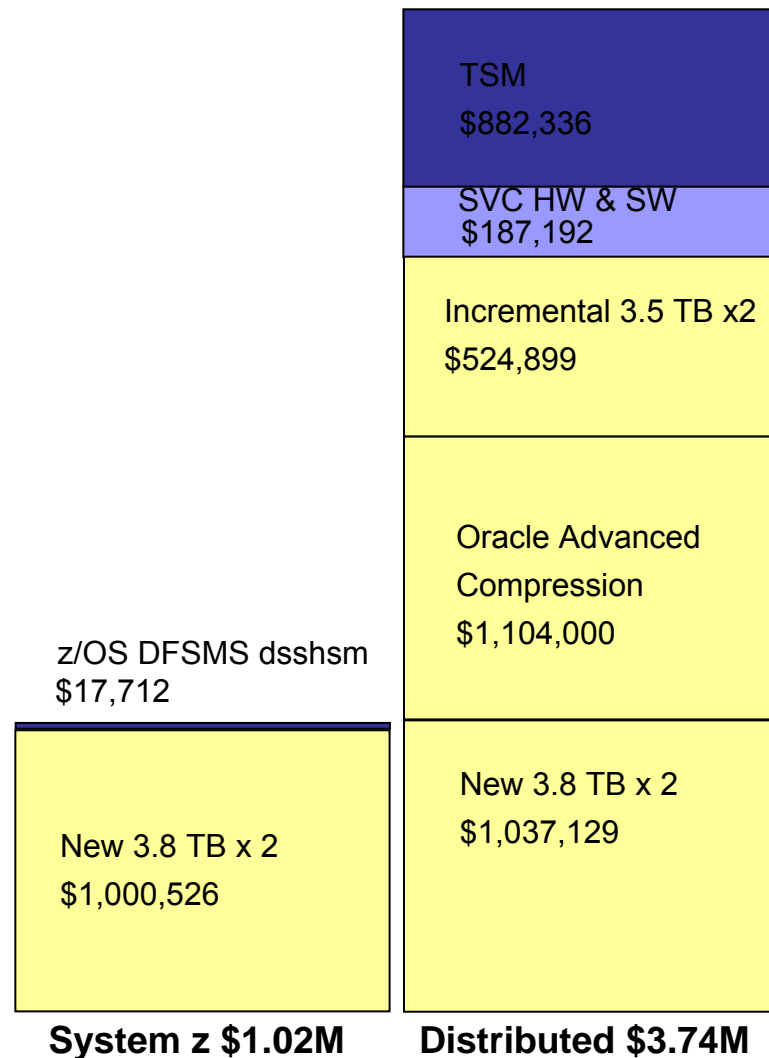
PAV supports parallel access of logical volumes within the same system and MA supports I/O parallelism across different systems

DB2 Hardware Compression For System z Further Reduces Storage Costs

- Data Warehouses (TPC-H benchmark):
 - ▶ 62% (DB2) vs. 27% (Oracle)
- Save over **TWICE** as much on disk space over Oracle
- DB2's compression also saves on memory and I/O used
 - ▶ You'll need less buffer cache than with Oracle
 - ▶ You'll also do less I/O than with Oracle
 - ▶ You'll also need substantially less backup storage space
- Flexible DB2 compression algorithm applies to more database tables
 - ▶ Oracle algorithm limitations limits it's effectiveness

System z And DB2 Reduce The Cost Of Storage By 73% For A New 10 TB Database


- For new storage capacity, 3.8TB x 2 (Primary+Secondary),
 - ▶ DS8100 for System z and HPXP2400 for Distributed
- Data Compression (10TB Storage)
 - ▶ System z – No incremental storage required, since DB2 uses built-in hardware compression, which supports up to 62%
 - ▶ Distributed – Incremental 3.5TB x 2 capacity since Oracle Advanced compression supports up to 27%
- Storage Management (HSM) and Virtualization (Data Sharing)
 - ▶ System z – DFSMS
 - ▶ Distributed – San Volume Controller (SVC) for Virtualization and Tivoli Storage Management (TSM) for HSM



Philippine Airlines Selects System z And IBM System Storage To Support Exponential Growth

- As PAL prepared to launch new routes to both domestic and international destinations, it realized it needed to upgrade its current information infrastructure
- PAL required better performance and superior throughput of the storage systems to run more efficiently. PAL also required an off-site fallback storage for business continuity and disaster recovery
- PAL replaced three different multi-vendor disk systems with an IBM storage solution consisting of the IBM System Storage Turbo DS8300

DB2 For z/OS Is The Smarter Choice For Structured Data



System z With DB2 Scales Further Than Oracle On HP Superdome Best Banking Benchmark

Asian Bank

- ▶ IBM System z9 and DB2
- ▶ TCS BaNCS (Cobol)
- ▶ 15,353 Transactions/second
- ▶ 50 Million Accounts
- ▶ IBM benchmark for customer


Bank of China **

- ▶ IBM System z9 and DB2
- ▶ TCS BaNCS (Cobol)
- ▶ 8024*** Transactions/second
- ▶ 380 Million Accounts
- ▶ IBM benchmark for customer

HP/Temenos *

- ▶ HP Itanium and Oracle
- ▶ Temenos T24 (Java)
- ▶ 2,153 Transactions/second
- ▶ 13 Million Accounts
- ▶ Largest banking benchmark performance claimed by HP

System z and BaNCS Online Banking Benchmarks



* SOURCE: TEMENOS BENCHMARKS; <http://h71028.www7.hp.com/enterprise/downloads/TemenosBenchmark.pdf>


** SOURCE: <http://www.enterprisenetworksandservers.com/monthly/art.php?2976> Source: InfoSizing FNS BaNCS Scalability on IBM System z – Report Date: September 20, 2006

*** Standard benchmark configuration reached 8024 tps, a modified prototype reached 9445 tps

DB2 Optimized For z/OS Outperforms Oracle Designed For Commodity Hardware

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**

HP Agrees! 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



DB2 pureXML Is Most Efficient

```

<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

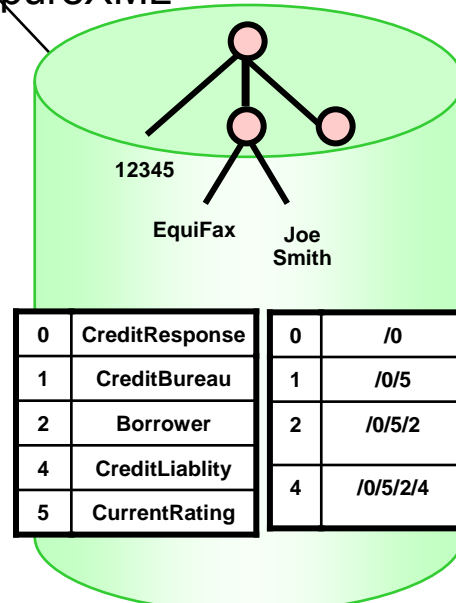

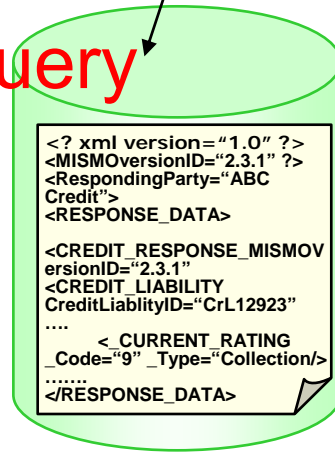
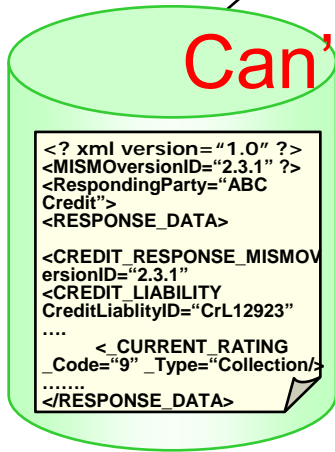
SQL XML

SQL XML

pureXML

Can't Query

Overhead



File System

BLOB
(Oracle)

Shredded
(Oracle)


Native Support
(DB2)

System z Is An Unparalleled Platform For DB2

- Coupling Facility enables unmatched scalability through centralized cache and lock management
- Hipersockets enable fast, secure communication between DB2 z/OS and zLinux applications
- Integrated with RACF and Multi Level Security
- Supports System z hardware encryption – Crypto Express 2
- DB2 Recovery Expert for automatic recovery and backup
- Offload up to 40% of workload onto zIIP processors to reduce licensing costs
- Hardware compression in addition to “Venom” deep compression can reduce storage up to 70%
- System z parallel sysplex supports rolling updates to running DB2

Oracle is Not !

Handling Unstructured Data Is A Key Step In A Smarter Information Management Strategy



Capture Paper Documents And Manage Electronic Data With FileNet

Our mortgage business is drowning in paper and folders! It's costly and slows us down



**Mortgage Business
VP**

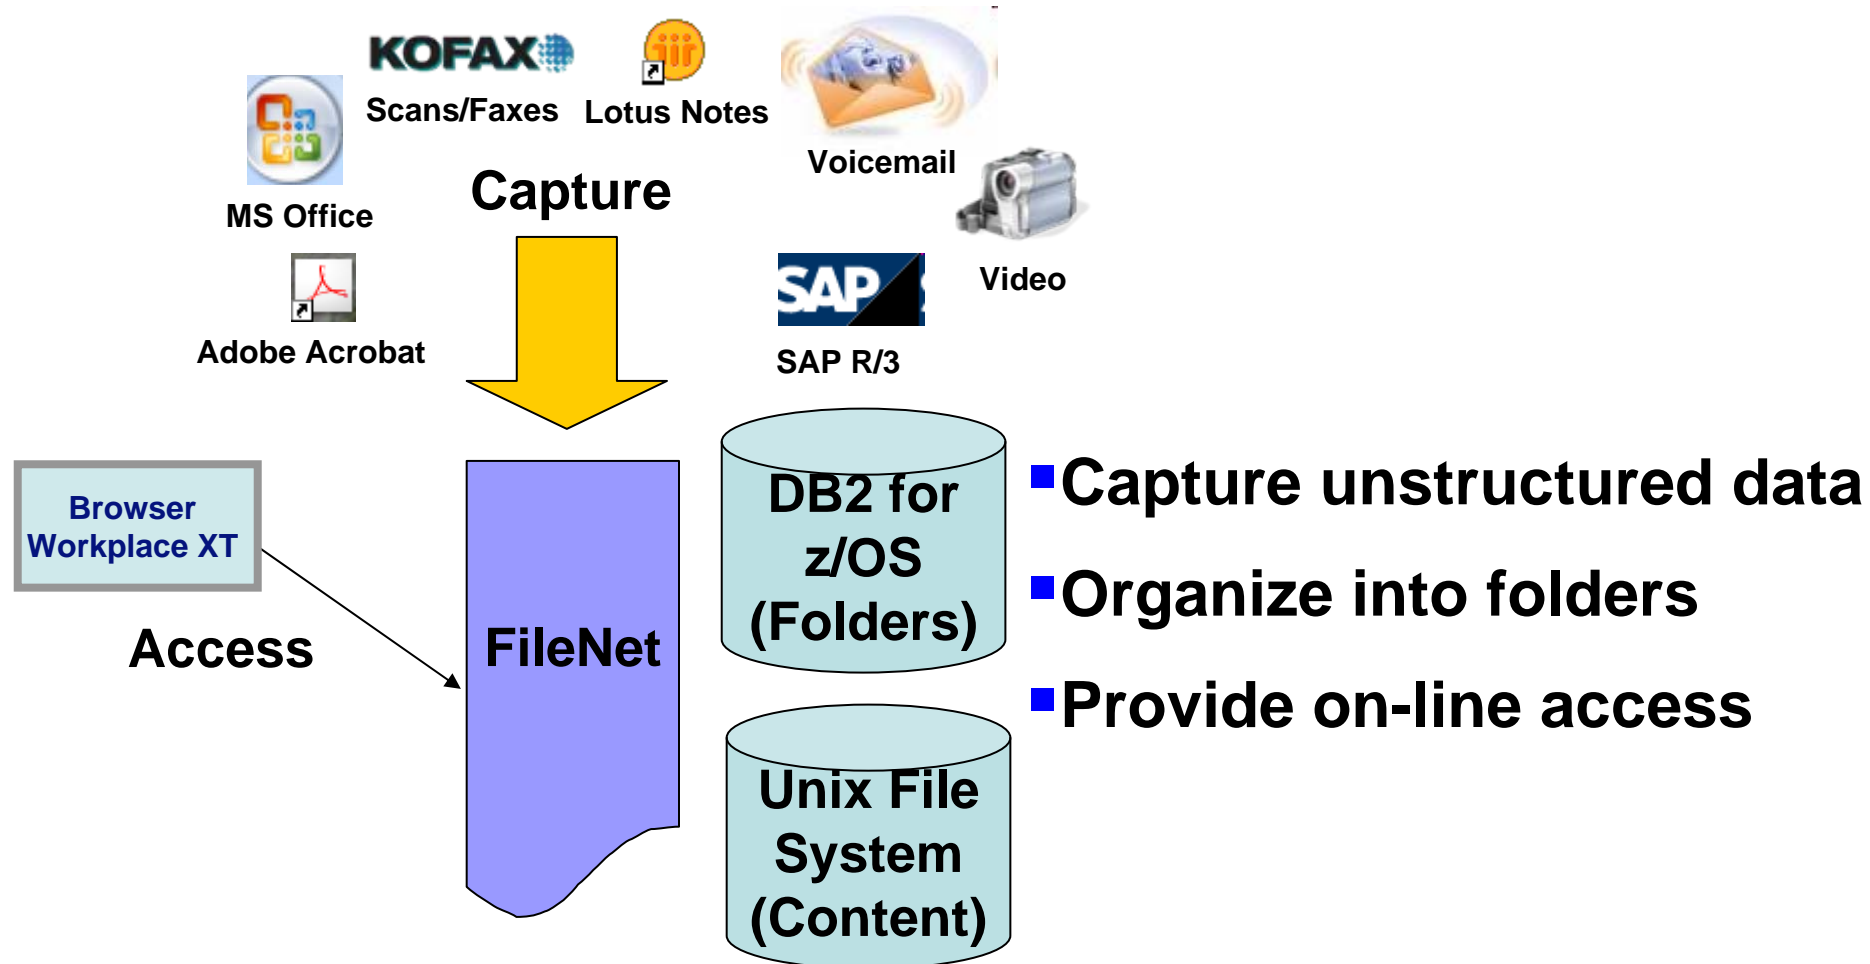


FileNet Content Manager on System z can help you efficiently store and use data in case management scenarios



IBM

FileNet Content Manager Captures A Variety Of Content Data Online



FileNet Content Manager Can Help Service Oriented Finance Go Paperless

Case Management Paradigm

- “Folder” collects all the documents for each mortgage:
 - ▶ Credit Reports
 - ▶ Proof of Identity (Driver’s License)
 - ▶ Email
 - ▶ Change of Address eForm
 - ▶ Picture and Video of house
 - ▶ Appraisal, Inspection Report
 - ▶ Federal Tax Return
- Role-based security
- Library Services (Check-In/Check-Out)
 - ▶ Versioning and Tracking for compliance

Mortgage documents can be accessed on-line using a case management approach.

Let’s review some of the capabilities IBM’S FileNet P8 Platform provides



IBM

FileNet Enterprise Content Management Solution For Mortgage Document Handling



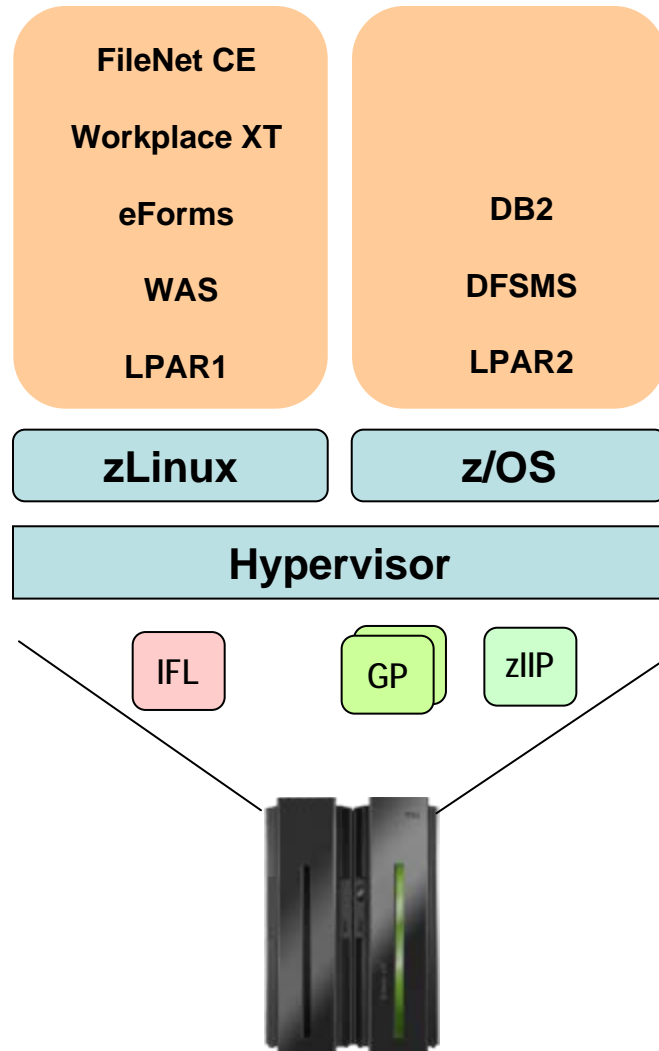
Processor (Lisa)



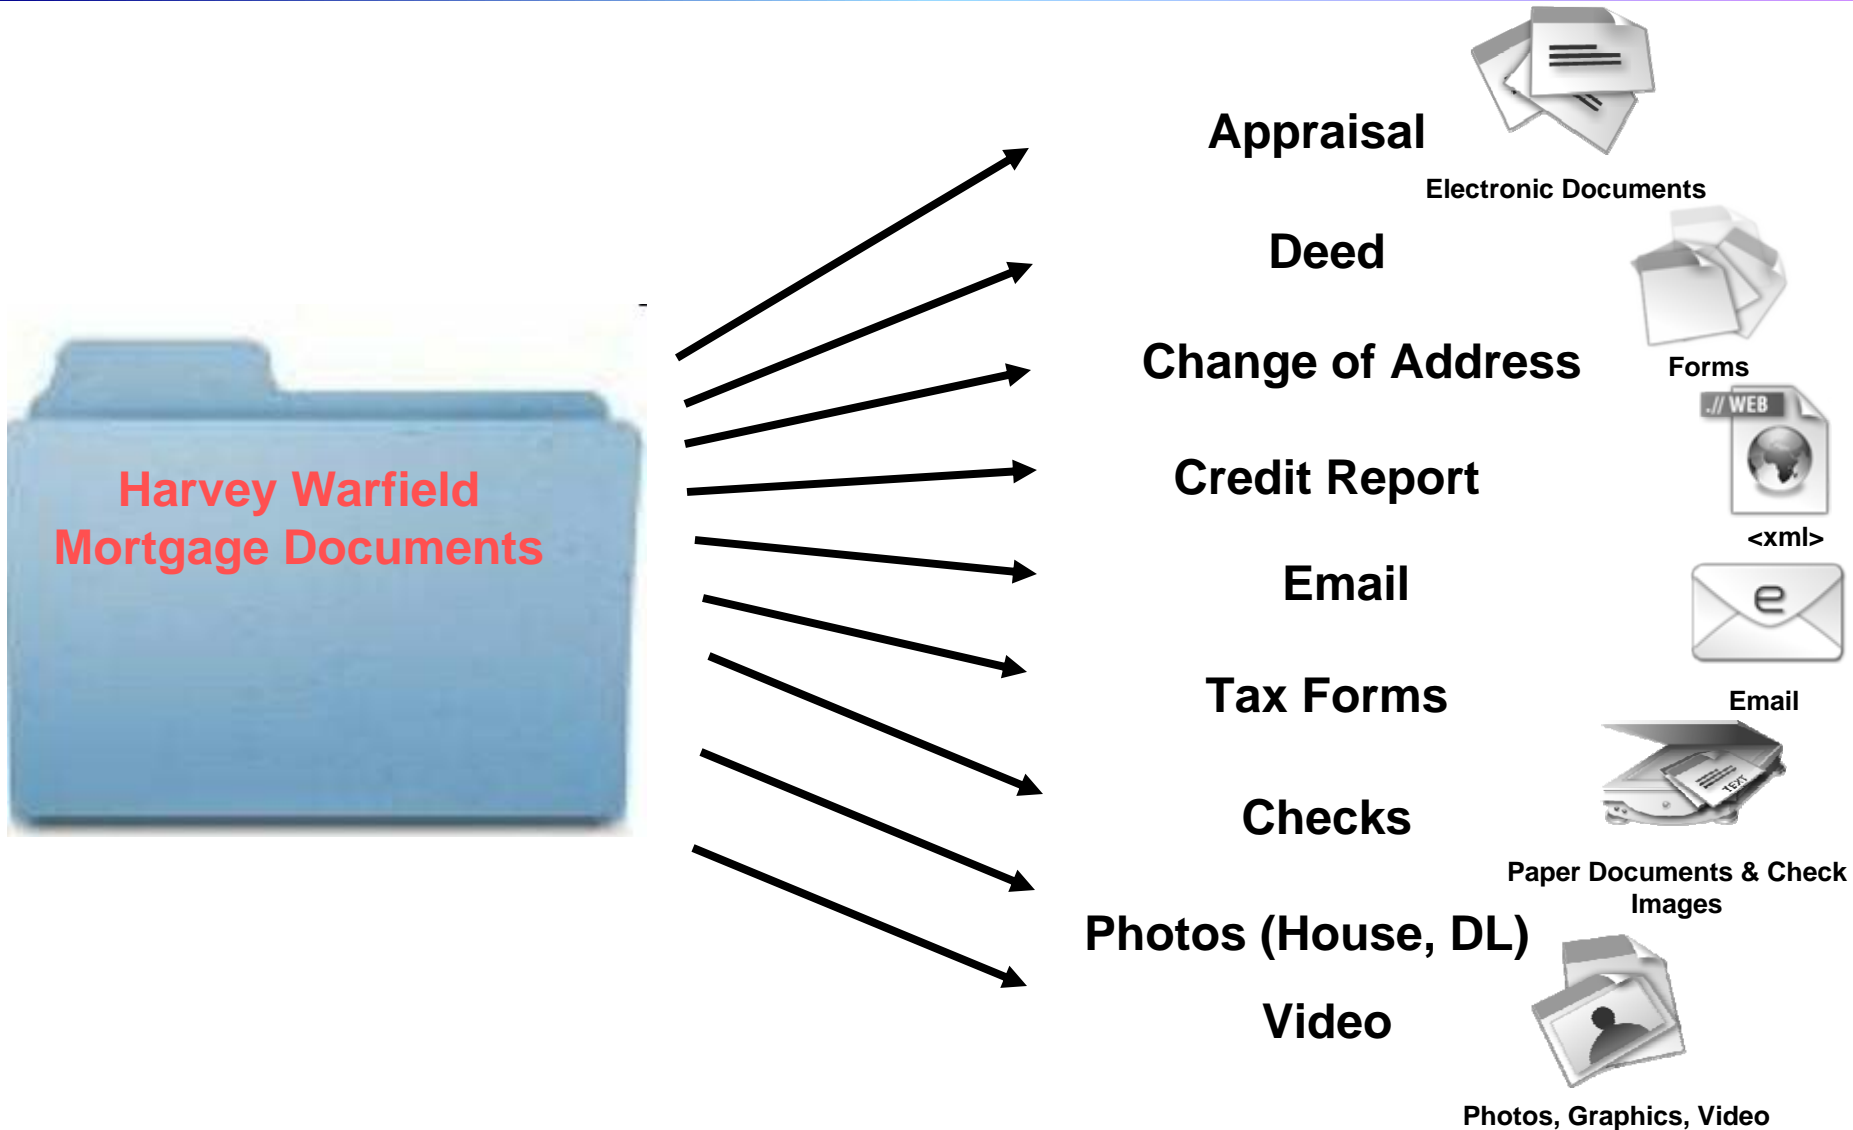
Compliance Officer (Homer)



Supervisor (Maggie)



DEMO: IBM FileNet Workplace XT, eForms And IBM FileNet Content Engine



Why Deploy FileNet On System z?

- Lower cost as an incremental workload
- Take advantage of System z storage management capabilities
 - ▶ Capability to store massive amounts of data
 - ▶ Virtualized storage is included
 - ▶ Hierarchical storage management is included
 - ▶ Dedicated I/O subsystem offloads I/O
- Linux for System z quality of service
 - ▶ Reliability and serviceability

Case Study: Deploy FileNet Content Manager On System z With Disaster Recovery (1000 Users)

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 processor
With 20Tb storage

Add 1 LPAR for FileNet Content Manager w 3.8 TB incr. storage




2,184 MIPS
additional
workload on z/OS
and 920 MIPS on
zLinux

Incremental:

2 GP 1,310 MIPS (60%) DB2, Utilities & DFSMS
1 zIIP 874 MIPS (40%) DB2
1 IFL 920 MIPS FileNet Content Manager & WAS
Add 10 GB memory

And add Disaster Recovery w 3.8 TB Storage




Capacity Backup:
2 GP
1 zIIP
1 IFL

*3 year
cost of
acquisition
\$5.8M*

Or add HP Integrity Superdome sx2k 9140N Server w 3.8 TB incr. storage

Prod



40 Chip
80 Core

272,902*
Performance Units

Documentum, Oracle,
SVC, TSM

And add Disaster Recovery w 3.8 TB storage

Prod



40 Chip
80 Core

HP DR solution is used in
software and hardware

*3 year
cost of
acquisition
\$12.6M*

*Performance Units required = (2,184+920) MIPS x 87 = 270,048

New Intelligence Is The Next Step For A Smarter Information Strategy

