

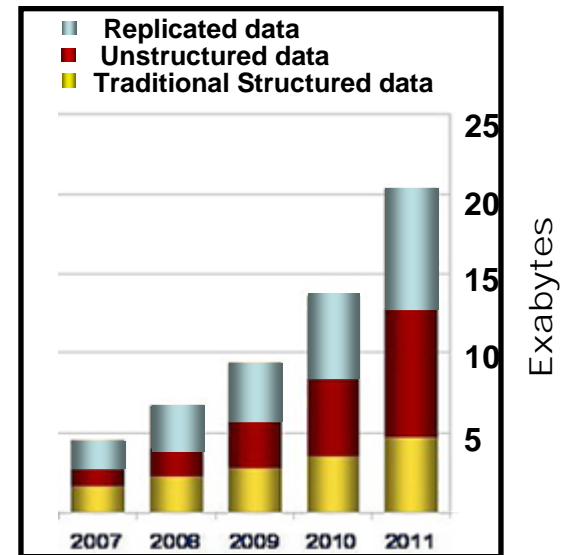


System z Enables Solutions For A Smarter Planet

Handling The Information Explosion

Data Data Everywhere

- Information is doubling every 18 months
 - ▶ The number of bits in the digital universe is already larger than the number of stars in the universe
 - ▶ Structured data growing at 32%
 - ▶ Unstructured data growing at 63%
 - ▶ Replicated data growing at 49%
- IDC predicts by 2010, storage of unstructured data will exceed that of structured



Sources: IDC worldwide enterprise disk in Exabytes from "Changing Enterprise Data Profile", December 2007

In This Brutal Economic Environment, Survival Implies A Smarter Approach To Information Management

We are drowning in paper and multiple copies of electronic content. It's costly and slows us down



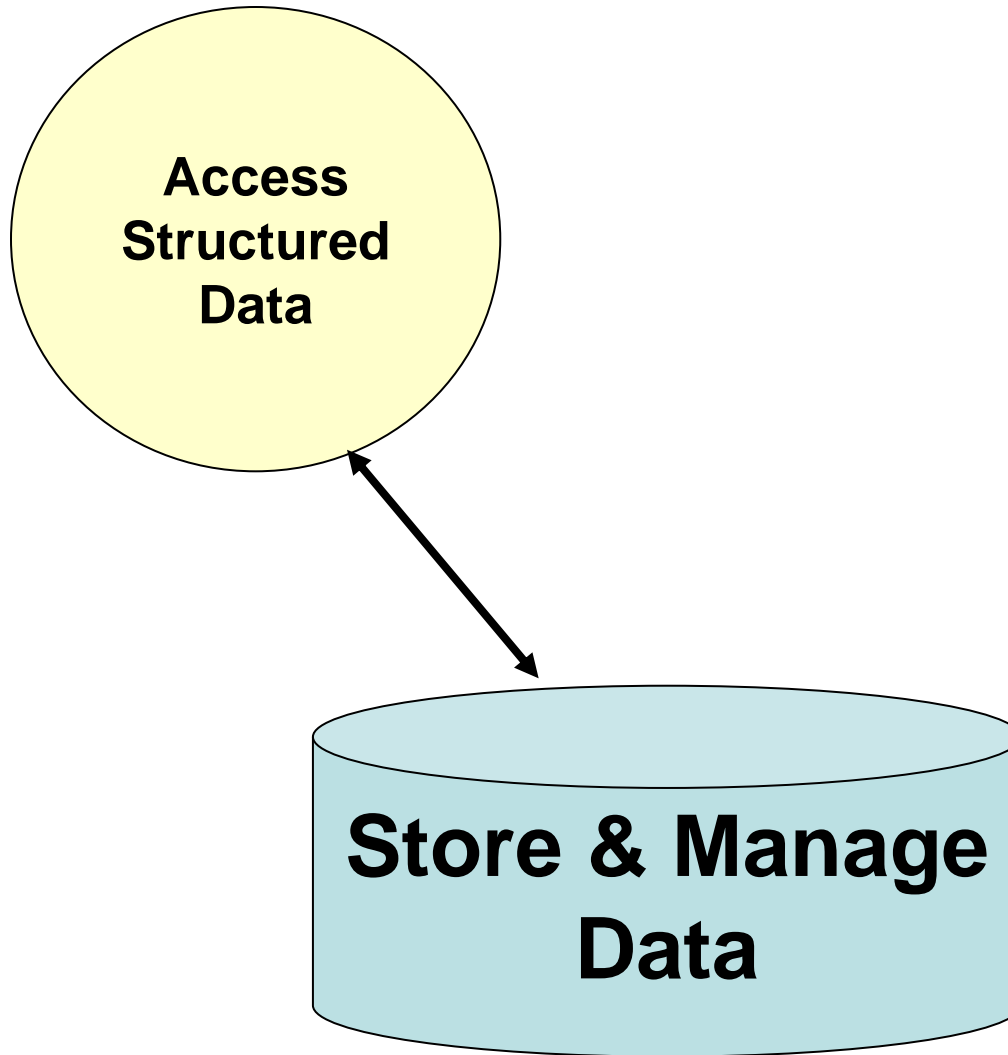
**Service Oriented Finance
COO**

You need a smarter information management strategy




IBM


Service Oriented Finance Stores Their Structured Data On Operational Systems That Run On System z



A Smarter Information Management Strategy



Content Management Is The First Step



Content Management Requirements

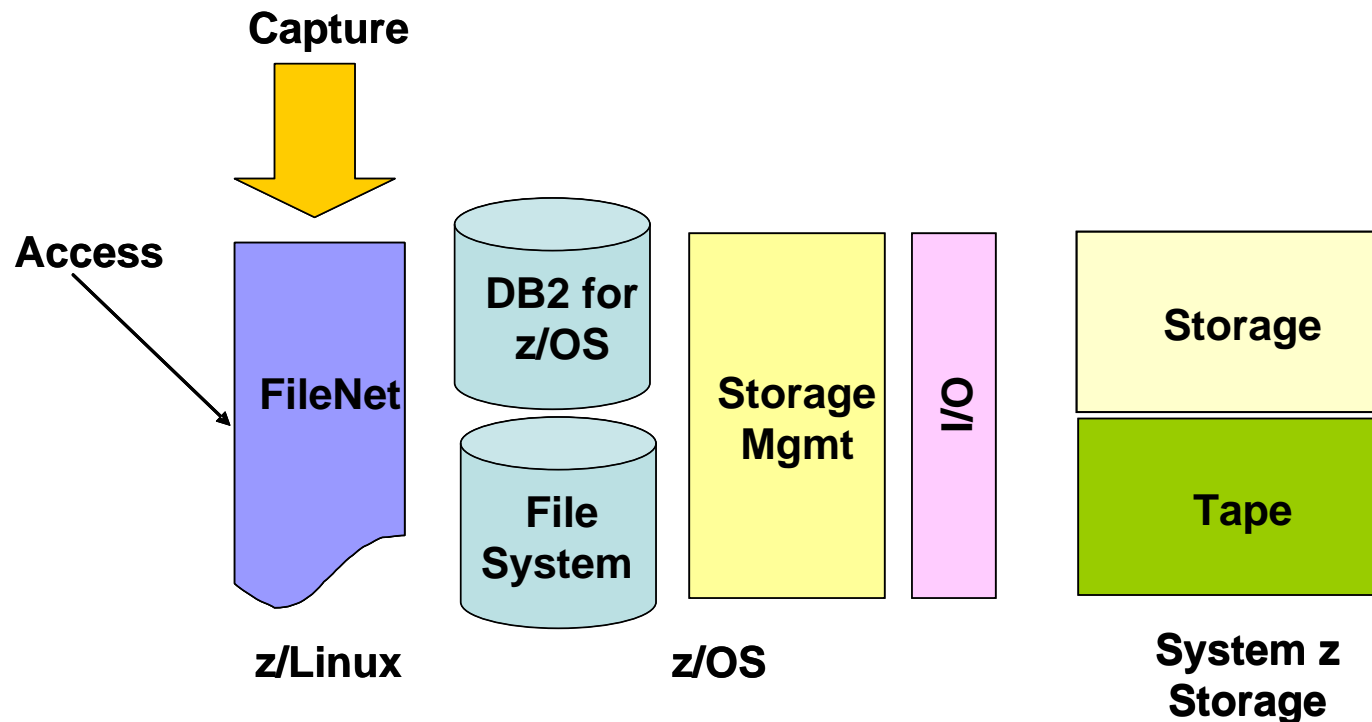
- Capture unstructured data

FileNet Content Manager

- Provide on-line access

- Store and archive massive amounts of data

System z




Content Management Requirements

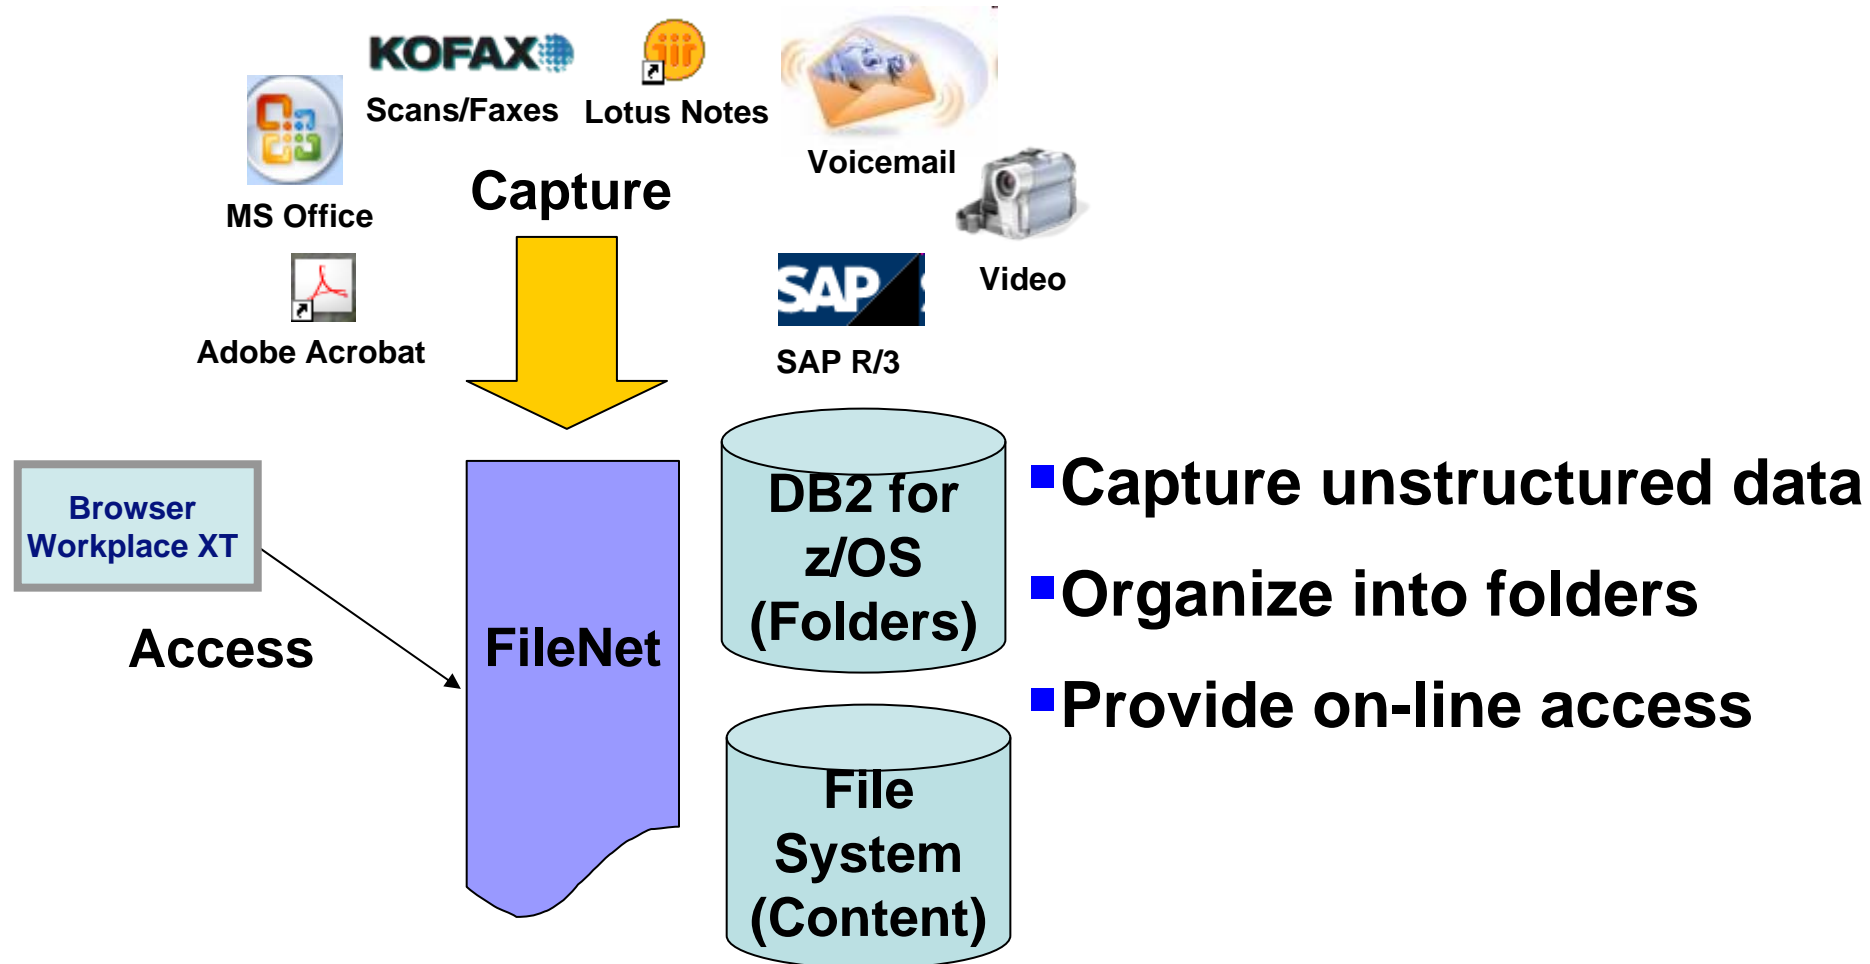
- Capture unstructured data
- Provide on-line access
- Store and archive massive amounts of data

FileNet Content Manager

System z



FileNet Content Manager Captures A Variety Of Content Classes And Provides Web Access



Mortgage Document Handling At Service Oriented Finance Is Currently Paper Based

To comply with regulations we need a better way to manage our mortgage documents



**Mortgage Business
VP**



**Service Oriented Finance
CIO**

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

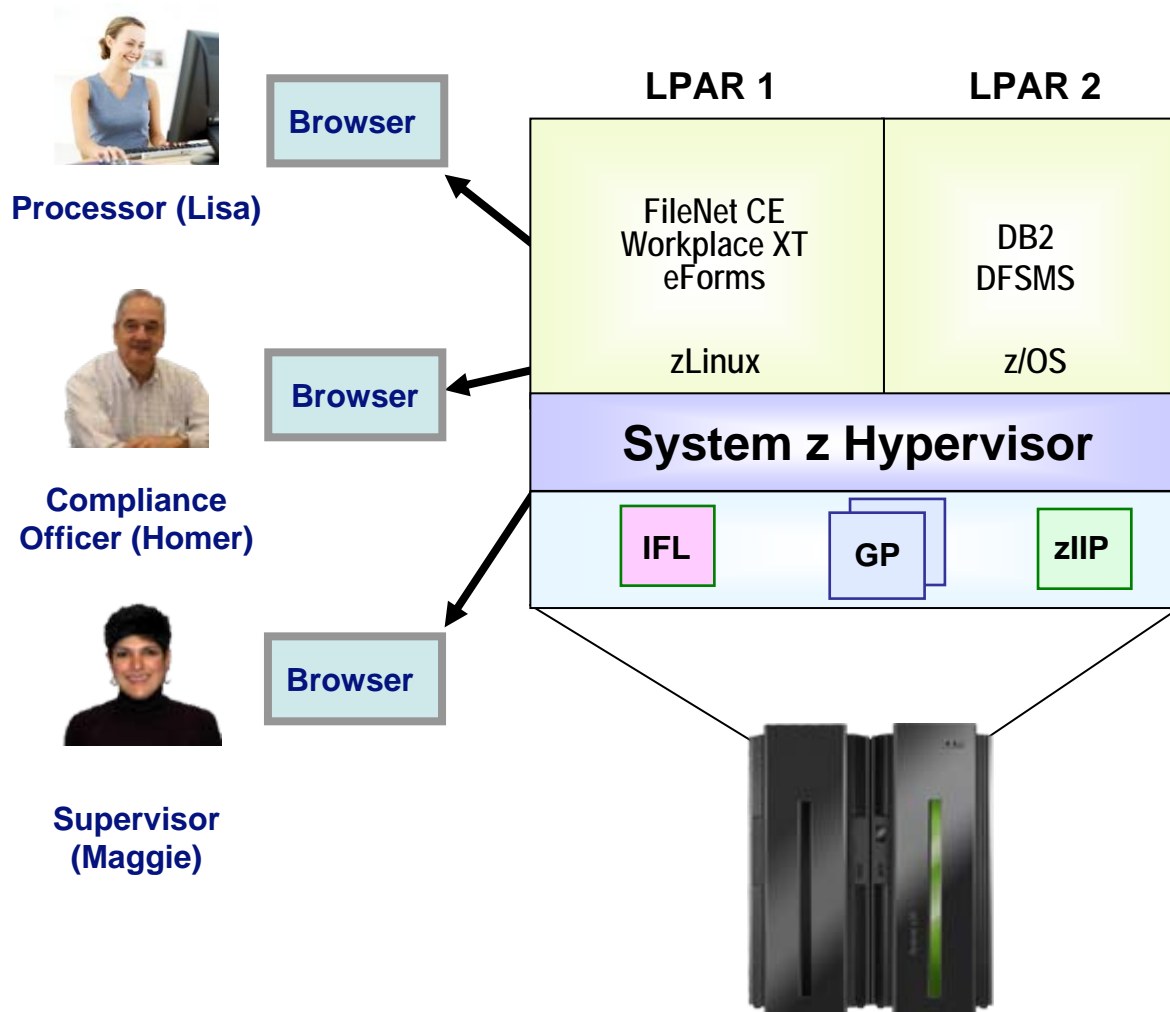
Mortgage documents can be accessed on-line using a case management approach.

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

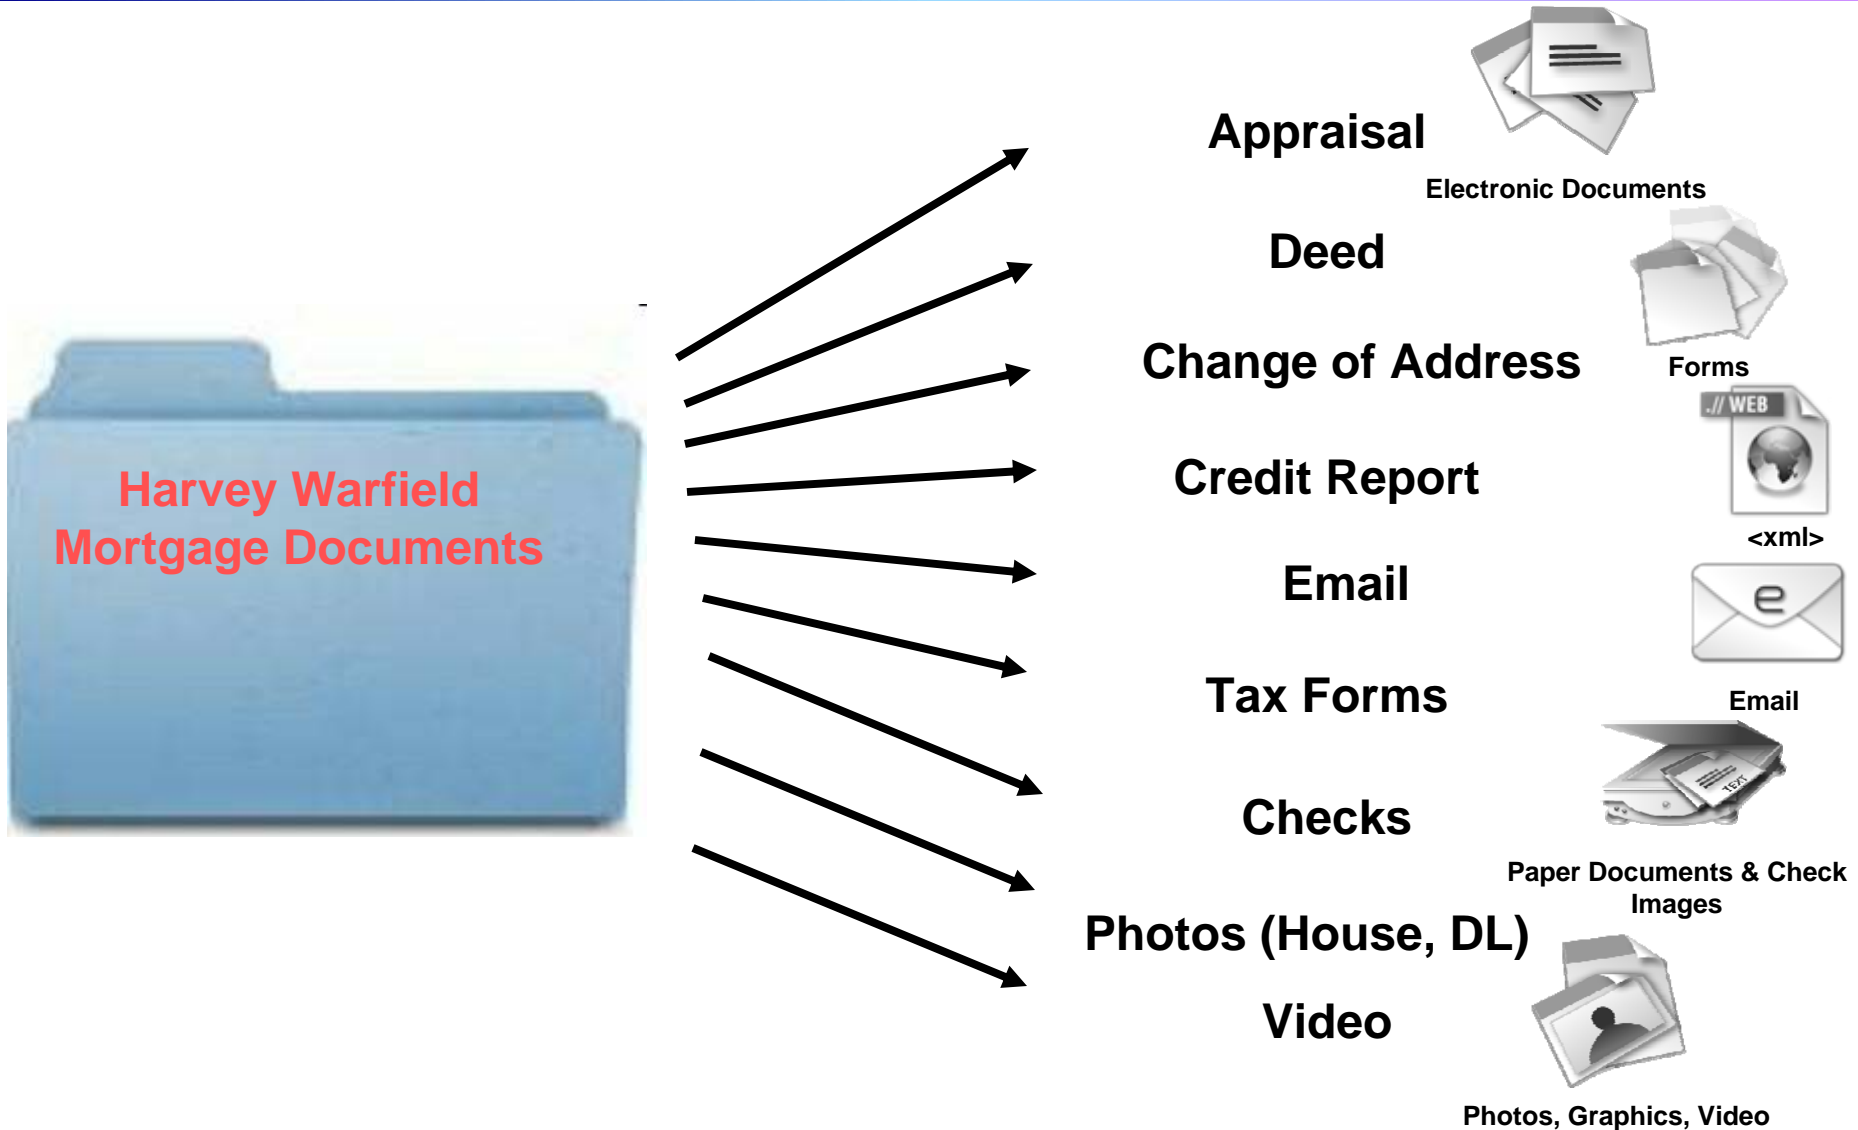


IBM

FileNet Enterprise Content Management Solution For Mortgage Document Handling



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



Why Deploy FileNet On System z?

- 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
- Lower cost as an incremental workload
- 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 LPAR for FileNet Content Manager w 3.8 TB Storage




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

Incremental:

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

And add Disaster Recovery w 3.8 TB Storage




3 year
cost of
acquisition
\$5.8M

Capacity Backup:
2 GP
1 zIIP
1 IFL

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

Prod




40 Chip
80 Core

Documentum &
Oracle

272,902*
Performance Units

And add Disaster Recovery w 3.8 TB storage

Prod



40 Chip
80 Core

3 year
cost of
acquisition
\$12.05M

HP DR solution is used in
software and hardware

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

Content Management Requirements


- Capture unstructured data

FileNet Content Manager

- Provide on-line access

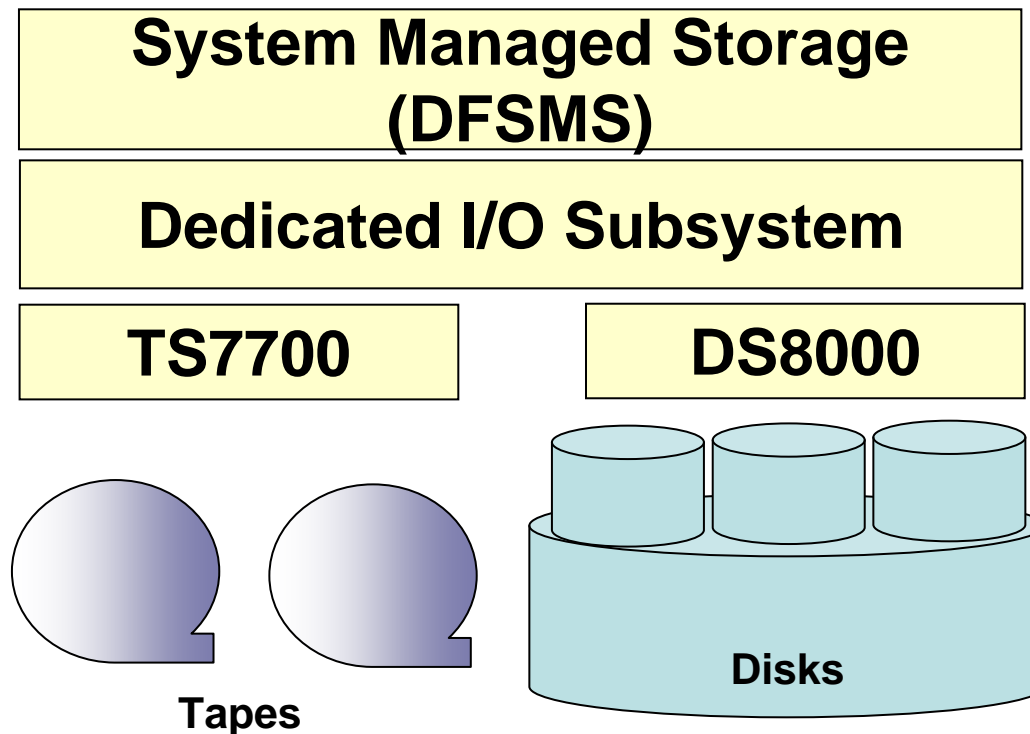
- Store and archive massive amounts of data

System z



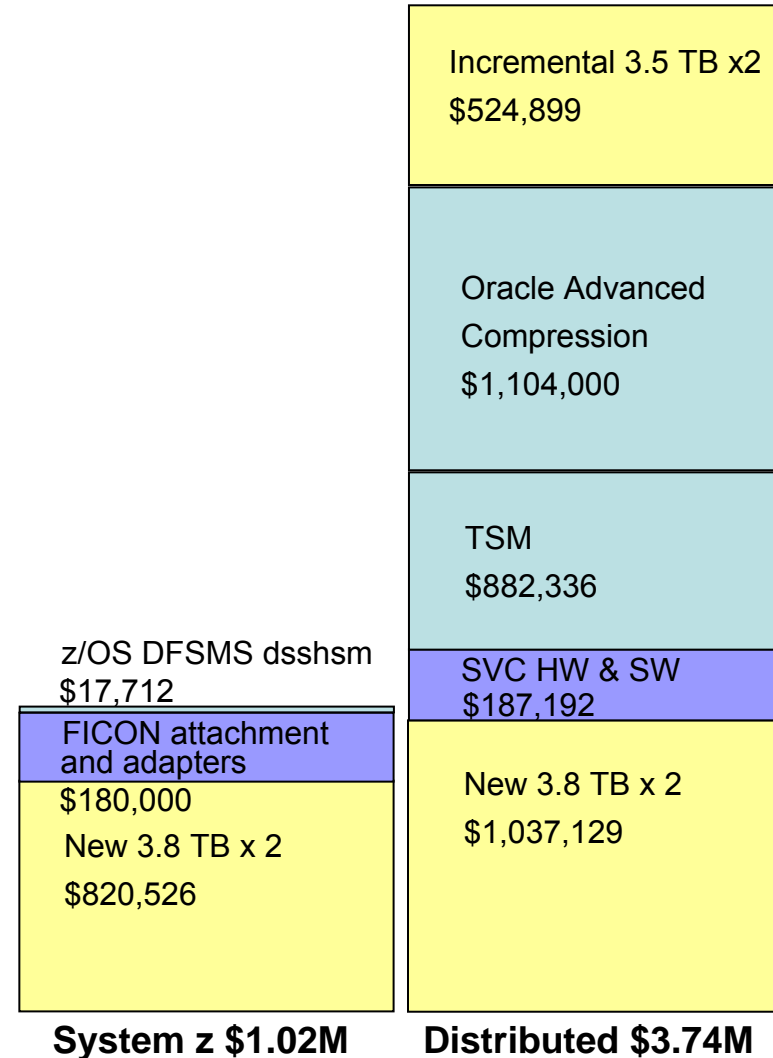
System z Storage Management Is Cost Efficient And Best Of Breed

- Virtualized and System managed storage functionality is built-in to z/OS
- Unique I/O subsystem providing data sharing, I/O virtualization and offloading
- Scalable, energy-efficient and tiered architecture for massive data storage and archival




System z Storage Is 73% Cheaper Than Distributed For A New 10TB Database

- For new storage capacity, 3.8TB x 2 (Primary+Secondary), DS8100 costs less than HPXP2400
- Storage Management (HSM) and Virtualization (Data Sharing)
 - ▶ System z – DFSMS and FICON attachment and adapters
 - ▶ Distributed – San Volume Controller (SVC) and Tivoli Storage Management (TSM)
- 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%




Virtualization Simplifies Management And Storage Cost Through Higher Utilization

- Individual disks over-provisioned
- Difficult to use stranded storage
- New storage added only when required
- Minimize stranded storage



Direct-Attached Storage

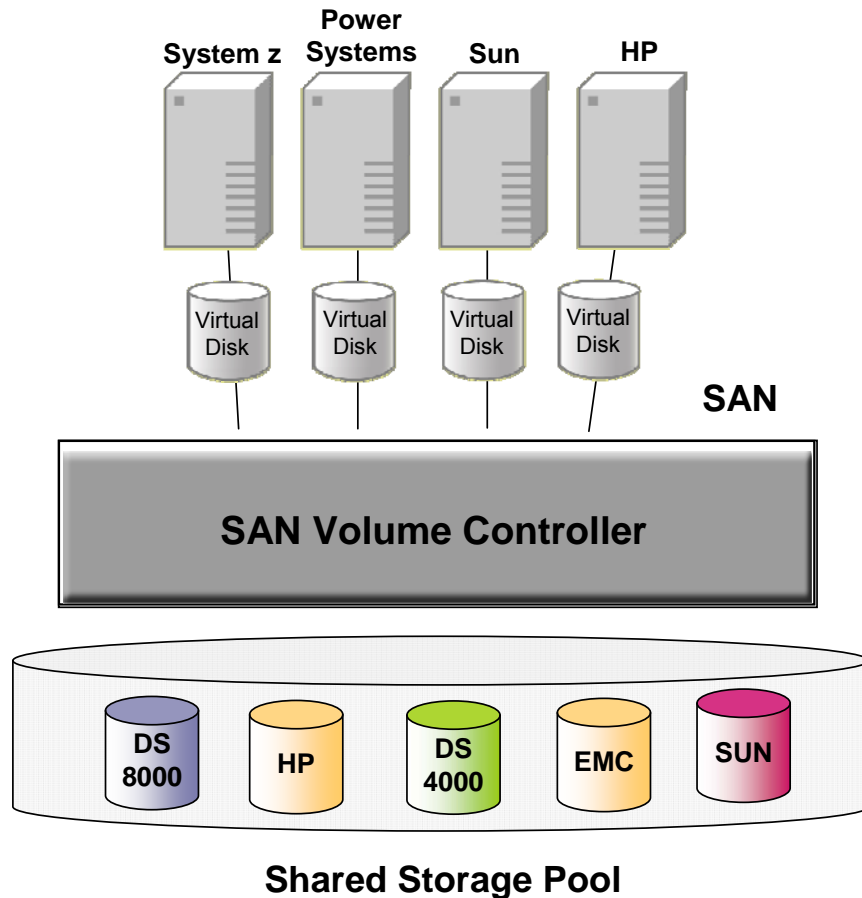


Shared Storage

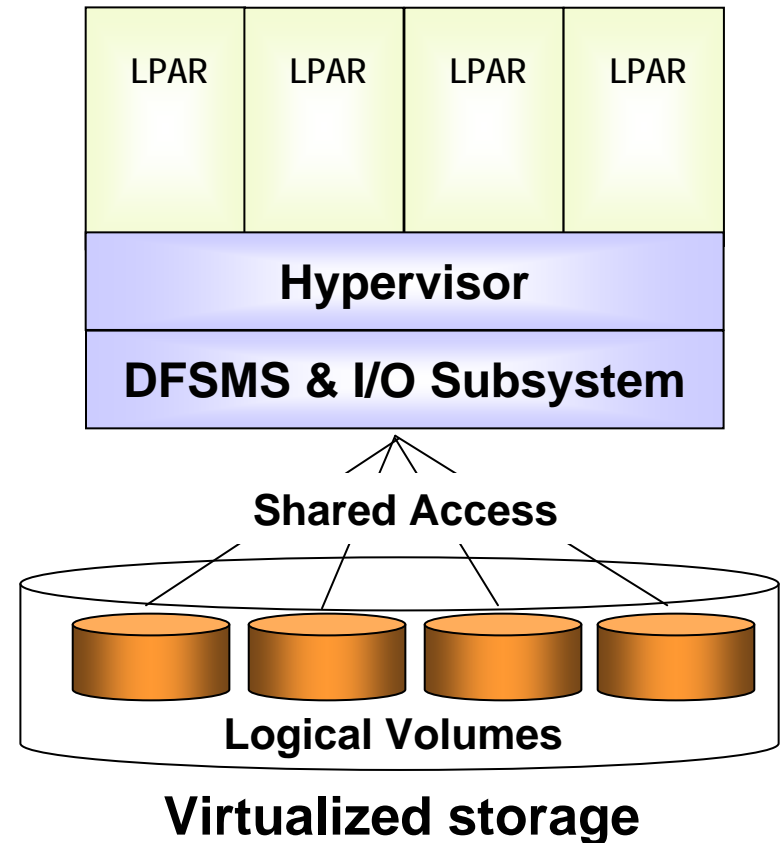
In this example, storage virtualization uses half the hardware!

You Need Additional Software For Shared Storage In Distributed, But It Is Included With System z

Distributed



System z



Storage Management Is Built Into z/OS Via Data Facility Storage Management Subsystem (DFSMS)

**dfp – automates management
of datasets, catalogs, objects
and z/OS UNIX files and
logical volumes**



**dss – copy, backup
and automatic
space management**

**hsm – manages
movement of data in
storage hierarchy**


**rmm – manage
removable data
(tape libraries)**

**tvS – concurrent
update of shared
VSAM data**

System Managed Storage Simplifies Management And Reduces Labor Costs

- Separate application, data and storage media
- Group physical volumes in storage pools
- Assign data sets/applications to pools based on policies
- z/OS automates data set allocation, retrieval and mounting
- Advanced Hierarchical Storage Management monitors data usage and automatically migrates data from high to low-cost media

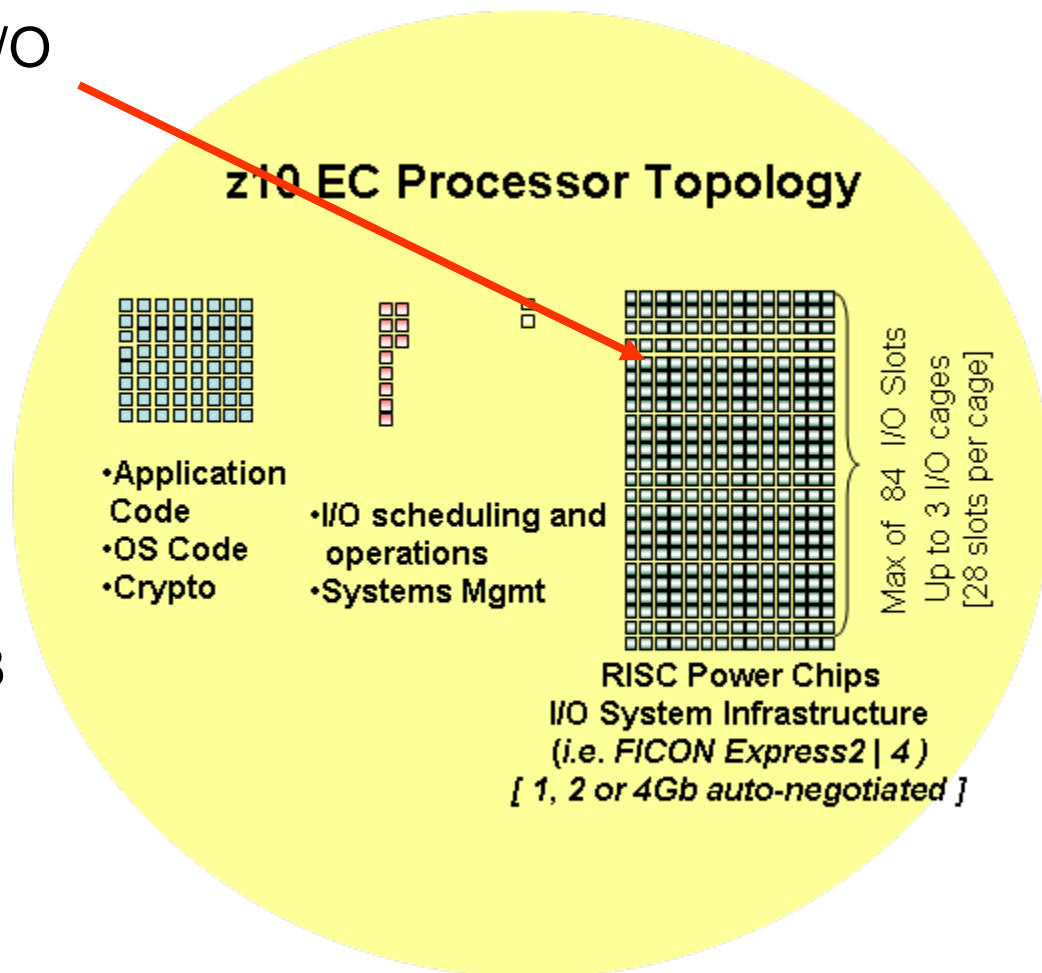
Hierarchical Storage Management (HSM) Autonomically Migrates Data For Archival



- 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 and Tivoli Storage Manager for Space Management is required for HSM in distributed environment

System z Has Dedicated I/O Subsystem To Enhance Access Performance

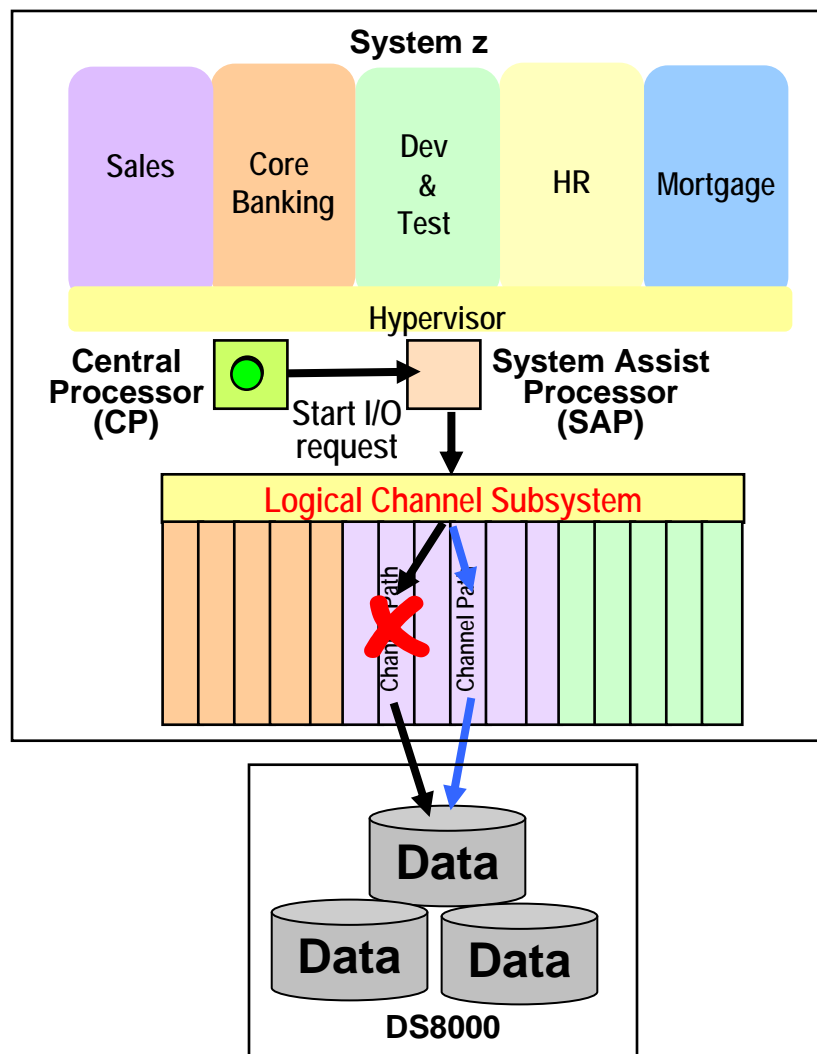
- Up to 336 RISC CPUs handle I/O
- Offload I/O operations to dedicated System Assist Processors, saving CPU
- Specialty zIIP processors for data mirroring
- Policy goals drive I/O priority (defined through DFSMS)
- Maximum I/O Bandwidth of 288 GB/sec
 - ▶ HP Superdome maximum bandwidth is 173 GB/sec and 122 GB/sec sustained



You don't get this with a distributed server

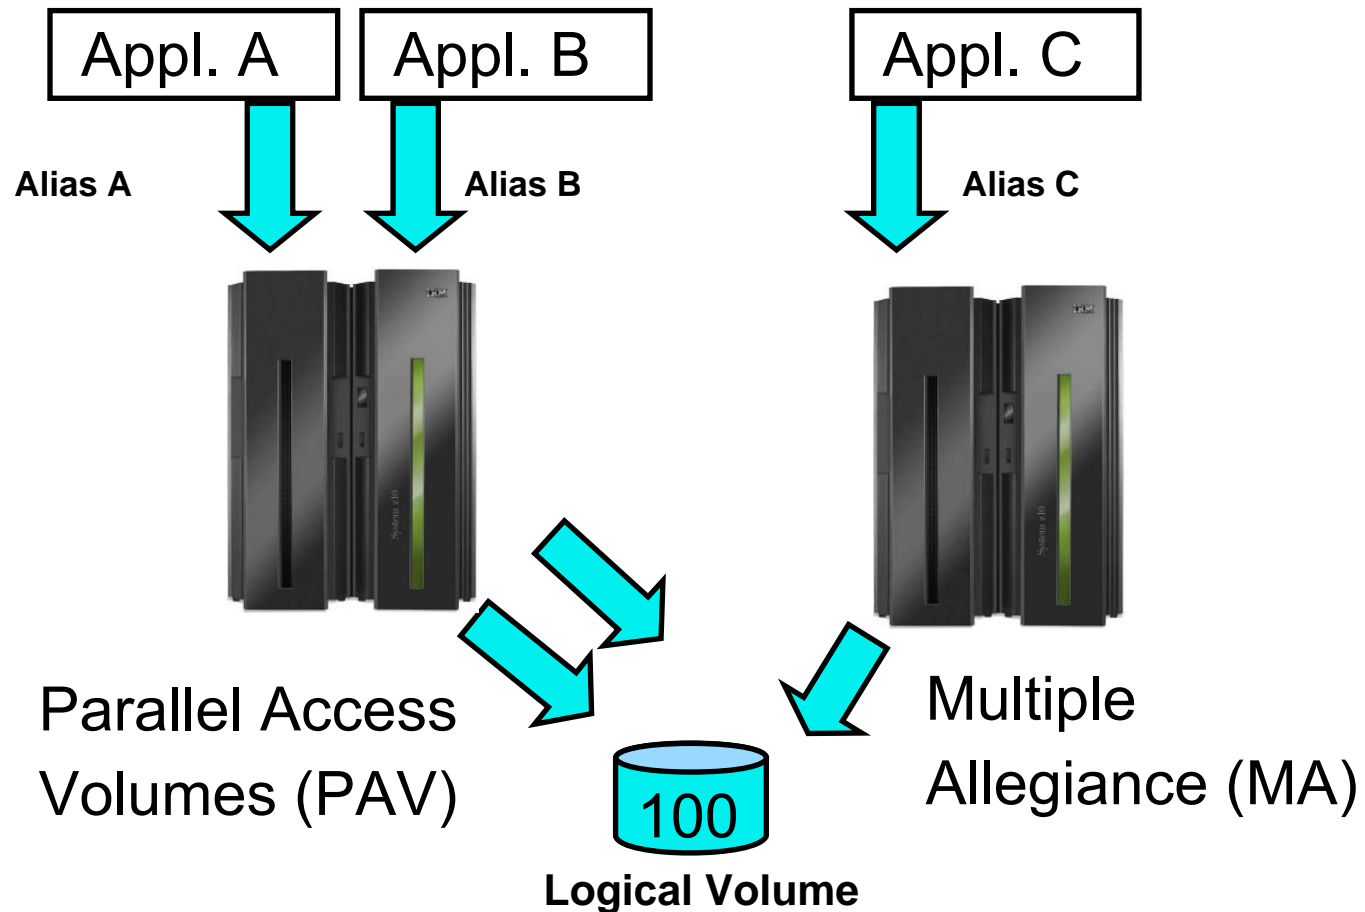
Virtualization Of I/O Enables Redundant I/O Paths

- I/O Virtualization provided by Logical Channel Subsystem
 - ▶ Up to 4 Logical Channel Subsystems, each with 256 channel paths and supporting up to 15 logical partitions (LPARs)
- Virtualization enables optimal Physical I/O path to be used
 - ▶ Dynamic path selection
 - ▶ Load balances I/O traffic
- Transparent Failover
 - ▶ Recover I/O operations in progress and switch to alternate path



System z provides uniform I/O access across the channel paths – better than HP Superdome's non-uniform I/O access across drawers

System z Disk I/O Operations Are Executed In Parallel Reducing Bottlenecks




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

Solid State Disk (SSD) Drives Are Here To Revolutionize Storage

- Semiconductor
 - ▶ Electronically erasable medium
 - ▶ Random access
 - ▶ No mechanical read/write interface
- Response times is around 0.8 milliseconds in contrast to 6 milliseconds for a typical 15,000 RPM (revolutions per minute) enterprise fiber channel hard disk drive (HDD)
 - ▶ 5-10x improvement in throughput & queries
 - ▶ SSD drives can sustain I/O rates of many ten's of thousands of I/O's per second while traditional spinning disk can handle hundreds of I/O operations per second
 - ▶ Reduce the "batch window"
- 75% reduction in number of disks switching from HDD to SSD
- Cost reductions due to no moving parts like spinning disks
 - ▶ 75% reduction in space
 - ▶ 80+% reduction in power and cooling
 - ▶ Reduce RAM requirements
 - ▶ Reducing peak MIPS saves \$3K/MIPS/year
- Avoidance of application changes




DFSMS Policy Based Storage Management Automatically Controls Which New Datasets Gets Allocated On HDD Versus SSD



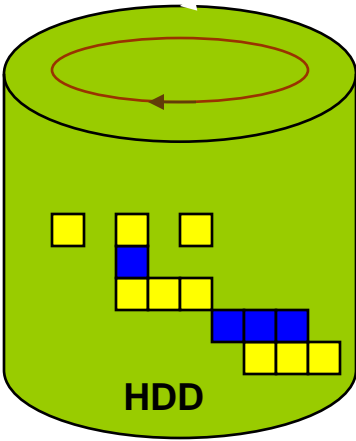
OLTP Application




z/OS SMS Policy for new datasets



Batch Application




HDD



SSD

HDD – Hard Disk Drive
SSD – Solid State Drive


DFSMS Policy Based Storage Management Automatically Controls Which New Datasets Gets Allocated On HDD Versus SSD




OLTP Application




z/OS SMS Policy for new datasets



Batch Application



HDD



SSD

HDD – Hard Disk Drive
SSD – Solid State Drive

IBM DS8000 Provides High Capacity, Secure, And Resilient Storage To Store And Process Growing Data Volumes

- 1.1 Terabyte (TB) to over **1000 Terabyte** on a single system with **max 1024 disk drives**
- DS8000 supports tiered storage with intermix of disk drives
 - ▶ 146GB (SSD), 450 GB (FC HDD, 15K rpm), 1TB (SATA, 7200 rpm)
- 4 Gbps fiber channel adapters each transmitting 154,000 Input/Output operations (IOPs) per second for a total of up to 4.9 million IOPs
- Equipped with up to 256 GB of cache with Intelligent Write Caching and Adaptive Multistream Pre-fetching
- Powered by dual 4-way Power5+ processors
 - ▶ Supports two virtual storage subsystems
- Stripe data across multiple RAID arrays
 - ▶ Minimize disk “hot spots”
- Mirroring for business resilience
 - ▶ Synchronous copies up to 300 km apart
 - ▶ Asynchronous copies over virtually unlimited distances




IBM DS8000 Supports Unique Features For Improved Scale And Capacity On System z


- **HyperPav** supports an on-demand I/O model, where base addresses for physical I/O resources are dynamically assigned per I/O operation to logical aliases from a pool based on workload requirements
- **Modified Indirect Data Address Word (MIDAW)** reduces channel traffic by allowing same command to read/write multiple discontinuous storage locations. **z High Performance FICON (zHPF)** doubles I/O bandwidth by shrinking channel program sizes, replacing multiple channel command words (CCW) with a single task control word (TCW)
- Simplifies disk management with reduced number of volumes by using “**Extended Address Volume**” feature of z/OS, which allows up to 223 GB per volume
- Disk availability is increased as migration to larger volumes is done online without disrupting operations with “**Dynamic Volume Expansion**”
- Administrator work for system load balancing is reduced as “**Storage Pool Striping**” feature puts data fragments on different disks automatically

TS7700 Provides A Virtual Tape Solution Supporting A Tiered Storage Hierarchy Of Disk And Tape

Virtual
Volumes
and Drives



Tape
Volume
Cache



TS3500
Tape
Library



- Provides transparent access to virtual tape volumes or drives
- Better performance with up to 70TB Disk Cache and integrated hierarchical storage management
- Scalable and Automated Tape solution supporting 16 tape drives and can archive up to 11 PB with Ultrium 4 cartridges

Philippine Airlines Selects 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
- In 2008, PAL celebrated the 25th anniversary of the Philippine Airlines Commitment to Excellence to Reservations Systems on IBM System z

New Intelligence Is The Next Step For A Smarter Information Strategy

