What zEC12 Can Do That Intel-Based Systems Cannot

David Rhoderick

A Fundamental Shift In Costs, Driven By Intel Sprawl

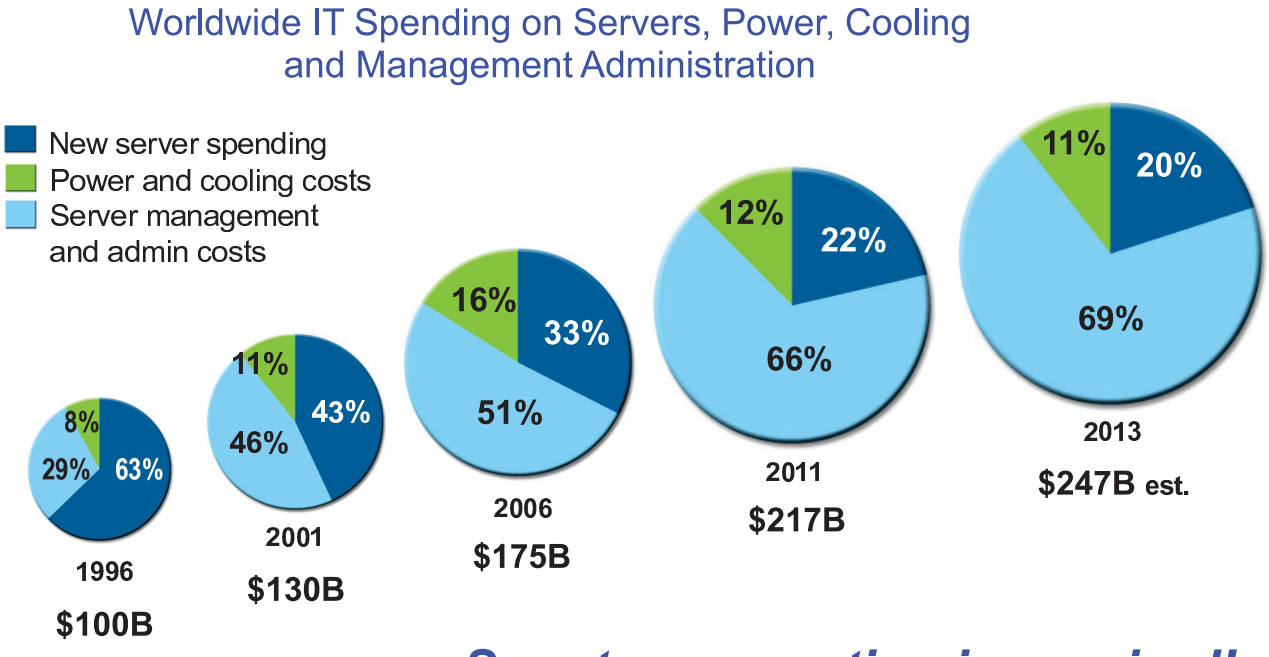

Smarter computing is required!

Source: IDC, 2012

Mainframe Offload Strategies Proliferate Intel Servers And Drive Up Cost

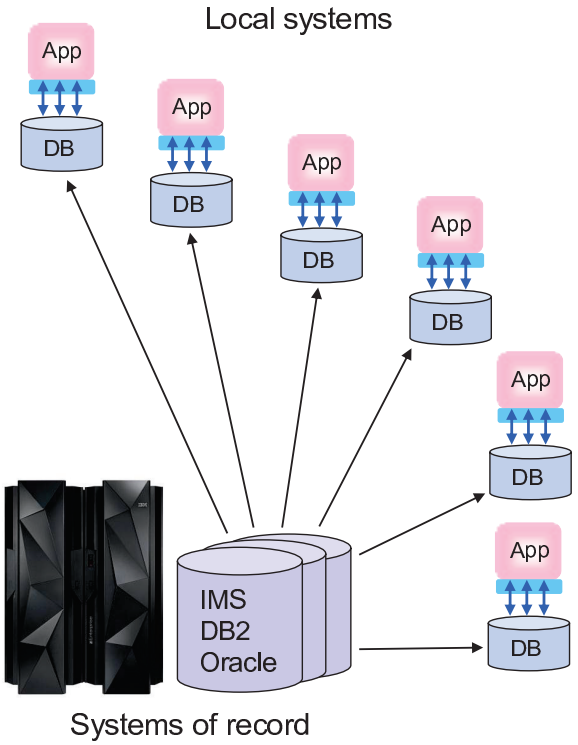

- Core proliferation
- Compromised performance
- Missing function
- Data duplication
- Lower quality of service
- More IT resources
- More cost

Intel Solution Required 128x More Cores

Typical Eagle TCO Study For A Financial Services Customer

Intel – 4 HP Proliant DL 980 G7 servers


Production


Development

256 cores total

Hardware	$1.6M
Software	$80.6M
Labor (additional)	$8.3M
Power and cooling	$0.04M
Space	$0.08M
Disaster Recovery	$4.2M
Migration Labor	$24M
Parallel Mainframe costs	$31.5M
Total (5yr TCO)	**$150M**

System z z/OS Sysplex

2,800 MIPS (2 zEC12* cores)

Hardware	$1.4M
Software	$49.7M
Labor	Baseline
Power and cooling	$0.03M
Space	$0.08M
Disaster recovery	$1.3M
Total (5yr TCO)	**$52M**

Intel solution cost 3x more!

* Estimated values as of July 2012

Intel Solution Compromised Performance

State Department of Motor Vehicles

- Offloaded registration system from CICS to Microsoft
  - Performed by Microsoft and Bearing Point
- Before: CICS had <1 second response times
- After: Microsoft has 30 second response times
- Cost of project: $28.3M, 3 years late

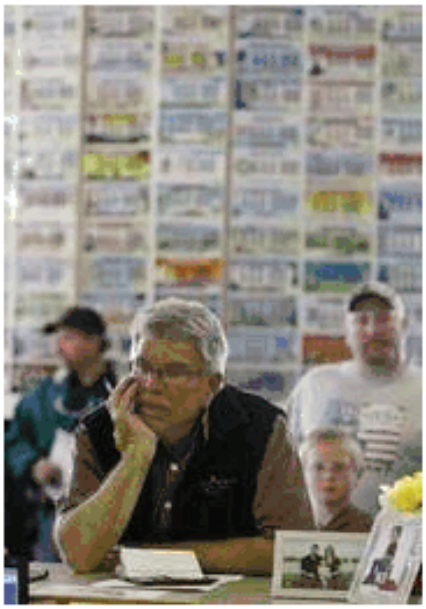

"Transferring titles is taking two to three hours instead of 15 minutes," Anderson said. One employee told him she had never heard so many "four-letter words" from customers.

zEC12 Advantages Over Intel

- zEC12 is designed to deliver more efficient processing capacity than Intel
  - Scales further for large enterprise workloads
  - Runs multiple concurrent workloads
  - Simplifies environments and reduces costs
- zEC12 has unique workload management capabilities
  - High utilization reduces cost
- zEC12 qualities of service are superior to typical Intel solutions
  - Administrator productivity, reliability, availability, capacity on demand, disaster recovery..

zEC12 Hardware Design Supports Shared Workloads Better Than Intel

- 101 configurable cores in a single server
- More on-chip cache per core
- Large L4 cache shared by ALL processors (across books)
- Uniform access times to memory

	zEC12	Westmere EP	Westmere EX	Sandy Bridge EP
On chip cache/core	8MB	2MB	3MB	2.5MB
Maximum L4 cache/server	1.5GB	None	None	None
Maximum memory/server	3TB	768GB	4TB	768GB

- Intel design works best for dedicated workloads
  - Performance degrades when workload not pinned to a core or socket

zEC12 Delivers More CPU Processing Capacity

z196

25% more performance, but constant pricing per MIPS and per specialty processor

zEC12

5.2 GHz	*Clock speed*	**5.5 GHz**
4	*Cores per chip*	**6**
96	*Total core*	**120**
80	*Configurable core*	**101**
1,202 MIPS	*General processor core performance*	**1,514 MIPS**
1,202 MIPS	*Specialty processor core performance*	**1,514 MIPS**
52,286 MIPS	*Total Capacity*	**78,426 MIPS**

zEC12 Parallel Sysplex Delivers Near Linear Processing Capacity – Intel Doesn't

- Specialized hardware - Coupling Facility
  - Dedicated processor with specialized microcode to coordinate shared resources
  - High speed inter-connect to clustered systems
  - Hardware invalidation of local cache copies
  - Special machine instructions
- Exploited by IMS, CICS, DB2, MQ, and other middleware on z/OS for transaction processing scale
- Unique hardware and software optimizations not found in Intel

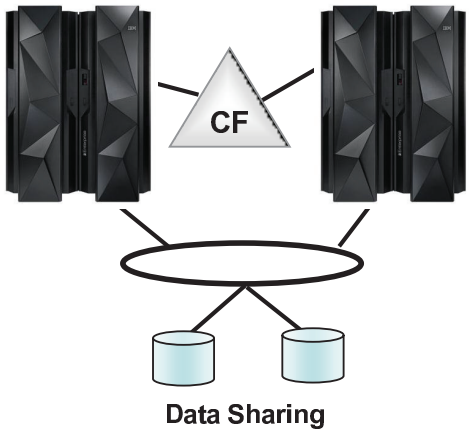

Cluster up to 32 nodes for a total of 3,232 configurable cores

78,426 x 32 = 2,509,632 MIPS

DB2 For z/OS In A Parallel Sysplex Scales Efficiently And Transparently

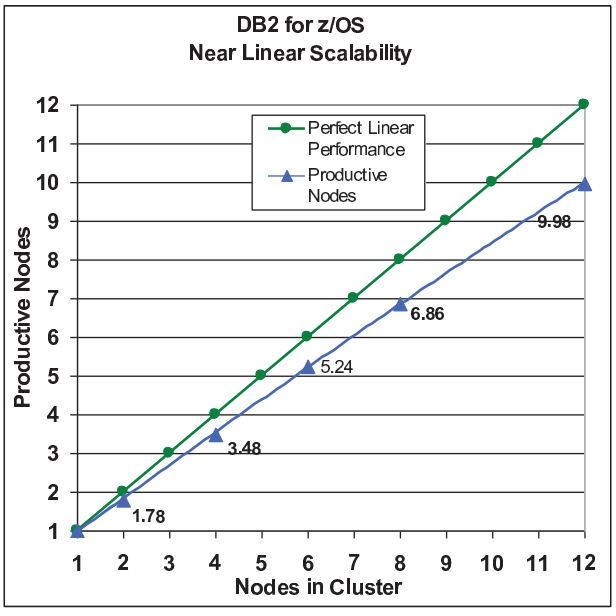

DB2 for z/OS OLTP result (ITG '03)

- DB2 leverages unique Parallel Sysplex clustering design to achieve near linear scaling
  - No data partitioning required
  - No transaction routing required
  - No cluster awareness required in applications
- Elastic processing capacity
  - Applications are not tied to database partitioning schemes
  - Automatically balances workload across cluster
- Virtually unlimited capacity
  - Up to 32 nodes can be clustered in a Parallel Sysplex
- Continuous availability during both planned and unplanned outages

The Only Option For Intel Is Oracle RAC

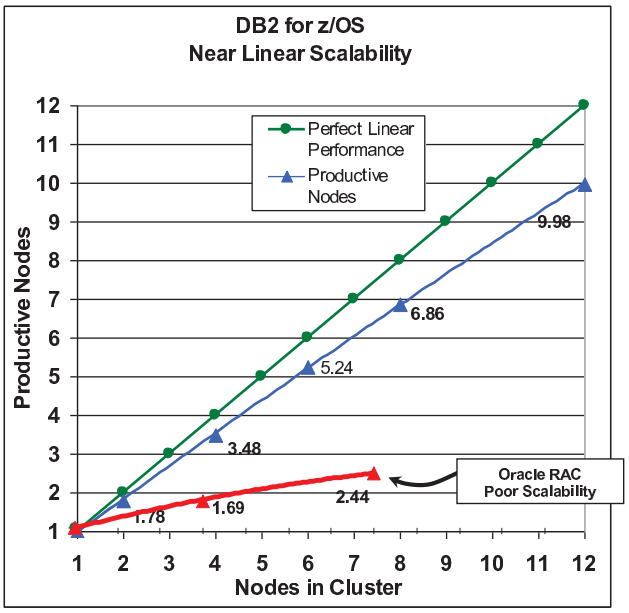

- Oracle RAC is inefficient by design
  - Network based lock and buffer management
  - Scaling RAC requires complex tuning and partitioning
  - Application partition awareness makes it difficult to add or remove nodes
- Published studies demonstrate difficult or poor scalability
  - Dell (shown in chart): Poor scalability despite using InfiniBand for RAC interconnect
  - CERN: Four month team effort to tune RAC, change database, change application
  - Insight Technology: Even a simple application on two node RAC requires complex tuning and partitioning to scale

Oracle RAC characteristics as shown in Dell RAC InfiniBand Study http://www.dell.com/downloads/global/power/ps2q07-20070279-Mahmood.pdf
CERN (European Organization for Nuclear Research) http://www.oracleracsig.org/pls/apex/RAC_SIG.download_my_file?p_file=1001900
Insight Technology http://www.insight-tec.com/en/mailmagazine/vol136.html

zEC12 Has A Dedicated I/O Subsystem For High I/O Bandwidth, Intel Doesn't

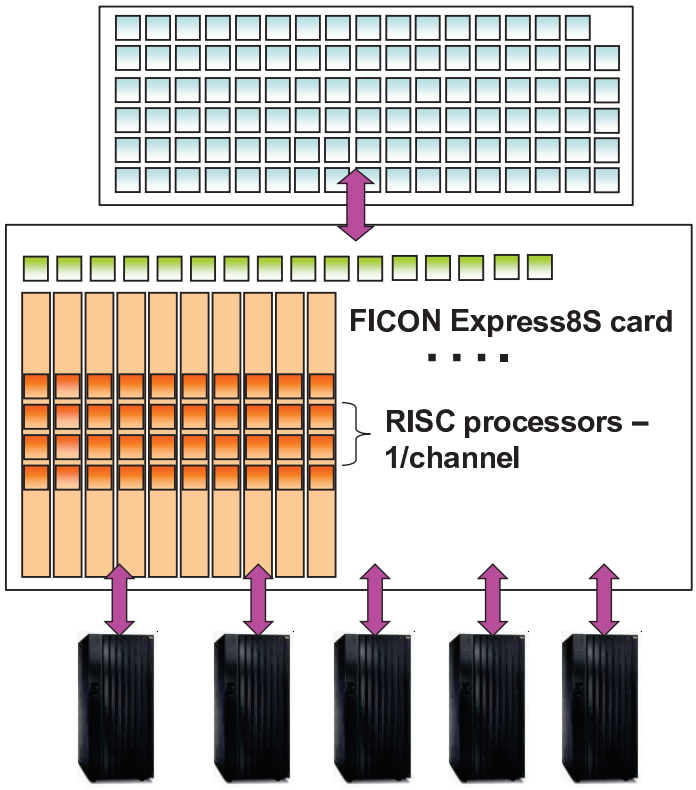

zEC12

- Up to 101 general purpose processors or Specialty Engines
  - Execute business logic
- Up to 16 System Assist Processors to manage I/O requests
  - Can sustain up to **2.4M IOPS***
- Up to 160 physical FICON cards for I/O transfers
  - Up to **320 RISC processors**
  - Up to **320 FICON channels**
- IBM DS8800 Storage System
  - Up to **440K IOPS capability**

* Recommend 70% max SAP Utilization – 1.7M IOPS
Numbers represent High Performance FICON traffic

Physical I/O Adapters And Channels In System z Are Virtualized And Shared By Workloads

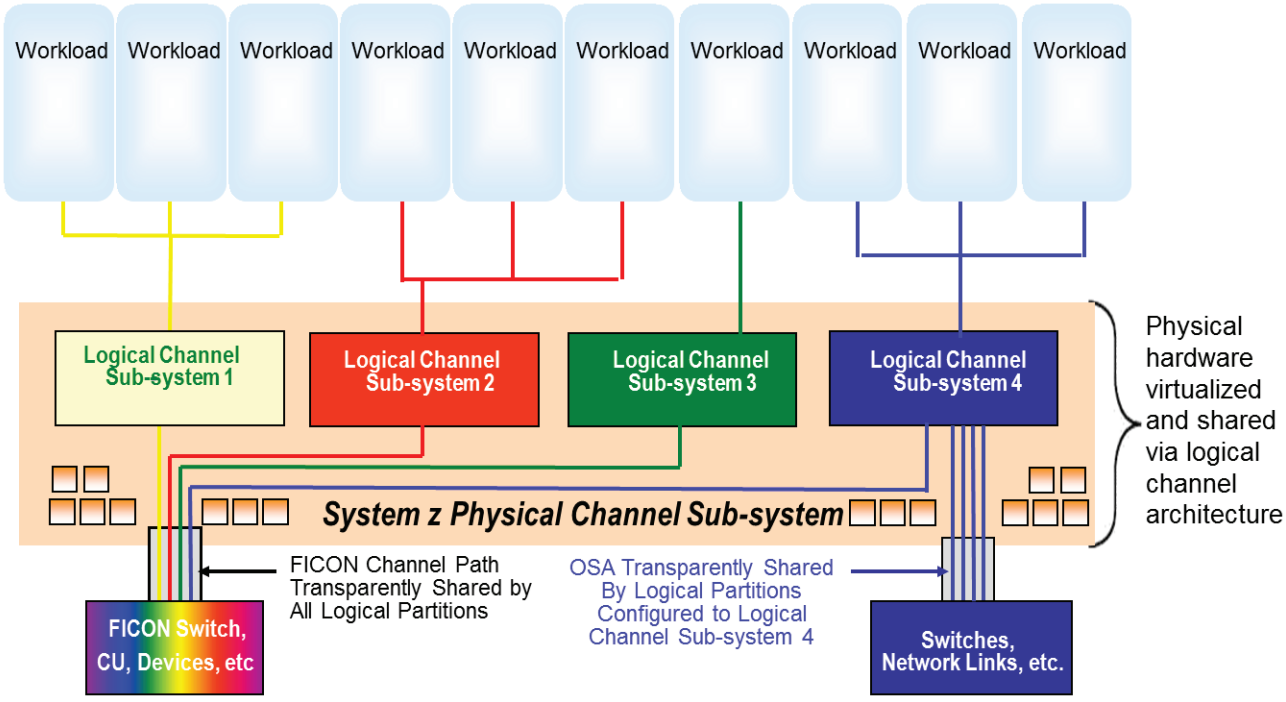

Advanced Capabilities In The zEC12 Channel System

- System z supports policy driven multi-pathing to dynamically satisfy I/O requirements
  - Intel Multipath I/O (MPIO) uses drivers configured at OS/hypervisor level
- Dynamic channel path ID management (DCM) manages a set of channel paths
  - Assigns channel paths dynamically to workloads
    - Based on the importance of the workload and the channel path delay measured for that workload
  - Chooses additional channel paths as workload scales
  - Automatically react to I/O hardware failures by switching to a different channel path

FICON Protocol Supports Advanced Storage Connectivity Features Not Found In Intel

- Extends workload management all the way down to storage
  - Supports DS8000 I/O Priority Manager integration with z/OS Workload Manager
- In-band I/O measurements to facilitate workload management
  - Fairness for I/O resource sharing
  - Dynamic assignment of I/O priority, dynamic throttling of I/O start rates

- End-to-end data integrity checking, transparent to applications and middleware, between the System z processor and storage controller
- First failure data capture for root-cause identification of failures
- Real-time configuration checking to detect cabling issues and prevent data corruption

Intel Performance Degrades As I/O Demand Increases

- Run multiple virtual machines on Intel server
- Each virtual machine has an average I/O rate
- Intel processor utilization is consumed as I/O rate increases

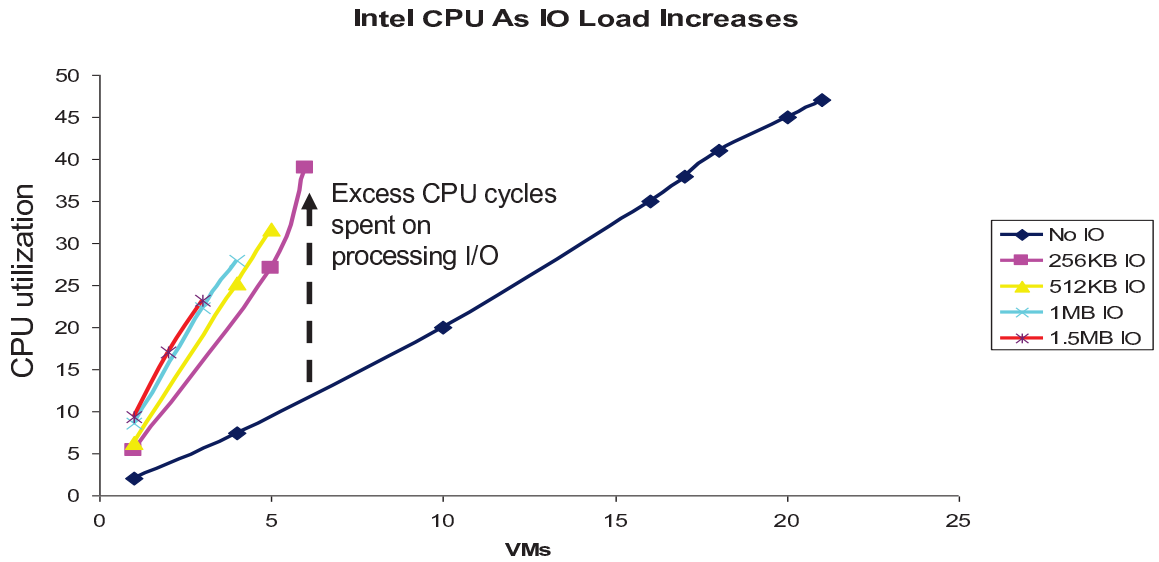

Source: CPO Internal Study. 12-core Westmere EP with KVM. FB at 22 tps with varying IO per transaction.

Database Workloads Benefit From High System z I/O Bandwidth

Competitor DB on Intel

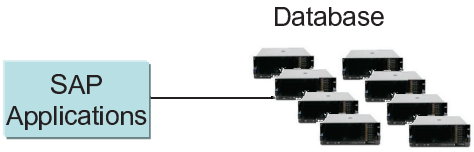

8x 3850 x5 with 32 cores
(dual active clusters)

256 DB cores

DB2 on z/OS

zEC12 2-way data sharing Sysplex *

2x22 DB cores
2x7CF cores

Database Unit Cost (5yr TCA)
$0.30/Postings per hour

Postings per Hour	42.0M
# of Accounts	90M
Hardware	$0.63M
Software	$11.98M
Total Cost	$12.61M

Database Unit Cost (5yr TCA)
$0.15/Postings per hour

Postings per Hour	59.1M
# of Accounts	150M
DB2 Solution Edition(HW+SW)	$7.49M
Capacity Backup (CBU)	$1.24M
Total Cost	$8.73M

A world record at half the cost!

Note: Cost of platform infrastructure for performance test transaction production. Cost of packaged application software not included. List prices used.

* Estimated values as of August 2012

Database Workloads Benefit From High System z I/O Bandwidth

Competitor DB on Intel

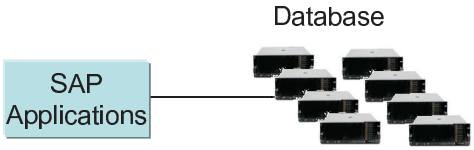

8x 3850 x5 with 32 cores
(dual active clusters)

256 DB cores

DB2 on z/OS

zEC12 2-way data sharing Sysplex *

2x22 DB cores
2x7CF cores

Database Unit Cost (5yr TCA)
$1.37/Accounts per hour

Accounts per Hour	18.0M
# of Accounts	90M
Hardware	$0.80M
Software	$23.94M
Total Cost	$24.74M

Database Unit Cost (5yr TCA)
$0.20/Accounts per hour

Accounts per Hour	37.2M
# of Accounts	160M
DB2 Solution Edition	$7.49M
Total Cost	$7.49M

84% less

Note: Cost of platform infrastructure for performance test transaction production. Cost of packaged application software not included. List prices used.

* Estimated values as of August 2012

zLinux Workloads Benefit From High System z I/O Bandwidth

Benchmark to compare consolidation platforms

workloads

Online banking workloads, each driving **22** transactions per second, with **1 MB I/O per transaction**

1 workload per 16-core Intel blade

Virtualized on Intel 16 core HX5 Blade

48 workloads per 32-way z/VM

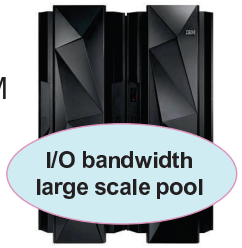

Virtualized on z/VM on zEC12* 32 IFLs

24x more workload density than Intel

* Estimated values as of August 2012

Oracle Workloads Benefit From High System z I/O Bandwidth

TCO study for a Media and Entertainment Industry customer

Oracle workload

Oracle Enterprise Edition
Oracle WebLogic Server
WebSphere MQ
SAS

107 total HP DL585 servers (1440 cores)

with 550 new network switches, routers, and cables

Which platform provides the lowest TCO over 5 years?

zEC12 with 48 IFLs* (Oracle EE)
Add zBX with:
1 PS701 (WMQ, WLS)
1 HX5 (SAS)

* Estimated values as of August 2012

Study Shows Oracle On zEC12 Is Half The Cost Of Oracle On Intel

TCO study for a Health Industry customer

107 HP servers

1,440 cores total

Hardware	$2.9M
Software	$24.2M
Labor	$7.9M
Space, Power and cooling	$1.2M
Disaster Recovery	$6.5M
Total (5yr TCO)	**$42.7M**

Oracle DB + App costs = $13.1M
(LIC + maint over 5 yrs.)

zEC12

48 IFLs*
1 PS701
1 HX5

Hardware	$4.9M
Software	$8.5M
Labor	$1.8M
Space, Power and cooling	$0.5M
Disaster recovery	$4.8M
Total (5yr TCO)	**$20.5M**

Oracle DB + App costs = $1.92M*
(LIC + maint over 5 yrs.)

* Estimated values as of August 2012

Email Sprawl At US Health Care Company

US Health Care Company Experiences Significant Email Outages With Windows Servers

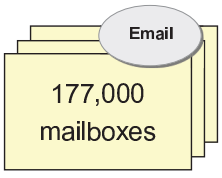

Lotus Domino Server
on 198 older x86
and p servers
(64% Windows, 36% AIX)

- Over 18 month period, customer's Windows servers experienced:
  - Almost 300 total email outages
  - 439 hours of total downtime
  - 1.5 hrs average duration of each outage
- Other hardware concerns included
  - Storage limitations
  - Increasing system complexity leading to increased operational costs
- Impending hardware refresh called for review of platform options
  - Existing mainframes had idle capacity

Domino Also Benefits From High System z I/O Bandwidth

Which platform provides the lowest TCO over 5 years?

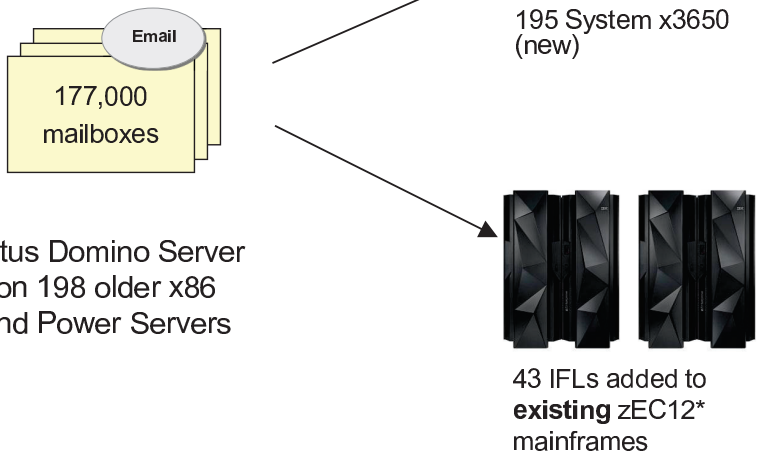

Windows on Intel x86

$27.5M (5 yr. TCO)

$155 per mailbox

Linux on System z

$9.33M (5 yr. TCO)

$53 per mailbox

1/3 the cost!

TCO includes hardware, software, labor, floor space, and power.

* Estimated values as of August 2012

zEC12 Advantages Over Intel

- zEC12 is designed to deliver more efficient processing capacity than Intel
  - Scales further for large enterprise workloads
  - Runs multiple concurrent workloads
  - Simplifies environments and reduces costs
- zEC12 has unique workload management capabilities
  - High utilization reduces cost
- zEC12 qualities of service are superior to typical Intel solutions
  - Administrator productivity, reliability, availability, capacity on demand, disaster recovery..

zEC12 Workload Management Optimizes The Use Of Hardware Resources

Example: Core banking workloads running on z/OS simultaneously with high utilization

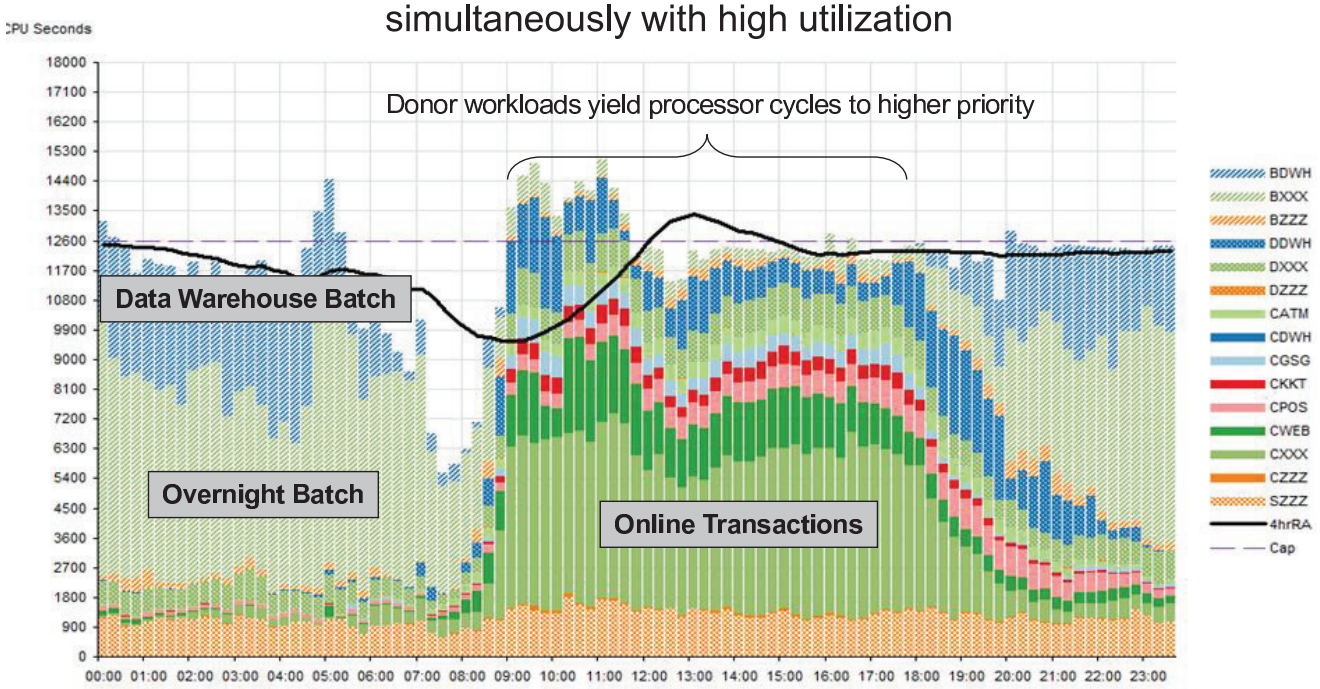

Efficient Workload Management Is Key To Mixing Different Workloads

- z/OS WLM efficiently prioritizes all computing resources **from top to bottom**
  - Processors to I/O to FICON to Storage
  - Isolates high priority workloads from lower priority workloads
  - Enables consolidation of database workloads while maintaining service level agreements (SLA)
    - Run at higher levels of utilization
- Oracle on Intel lacks top to bottom workload management
  - Processors are prioritized but I/O can interfere
  - Workload performance becomes highly variable
    - Analytics workload overwhelms OLTP
    - Even with extensive tuning not able to isolate workloads

Workload Management Is Critical In the Context Of Virtualization

- Multiple tenants could have differing levels of priorities
- Various mechanisms like OS priorities, VM share values, weights etc. are used to make sure the higher priority workloads are not impacted
- System z provides advanced workload management mechanisms to ensure the performance profiles for higher priority workloads are not affected
  - PR/SM provides advanced workload management at the LPAR level
- Limited workload management in Intel-based systems results in a degradation in the performance profile of higher priority workloads

Performance Profile Of High Priority Workload Is Maintained On System z

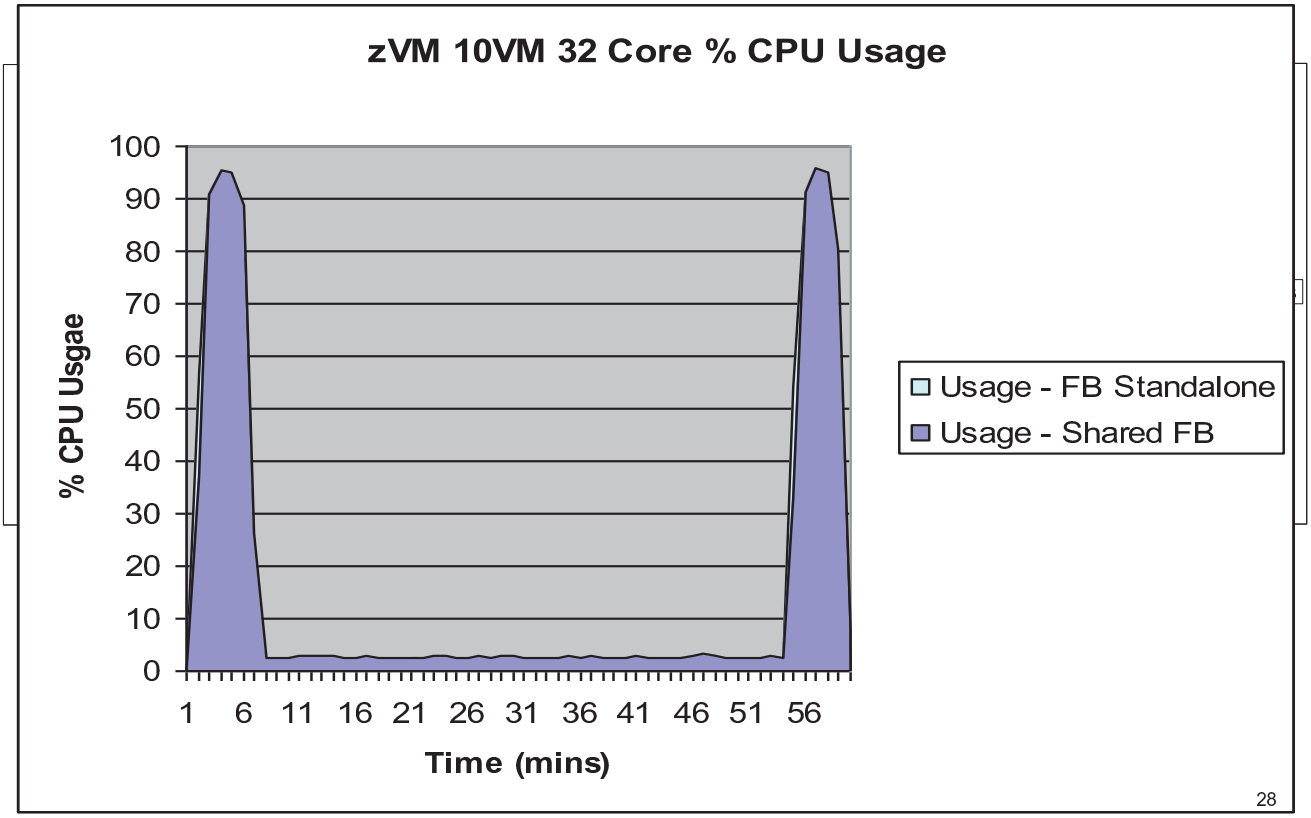

Performance Profile Of High Priority Workload Degrades On VMware

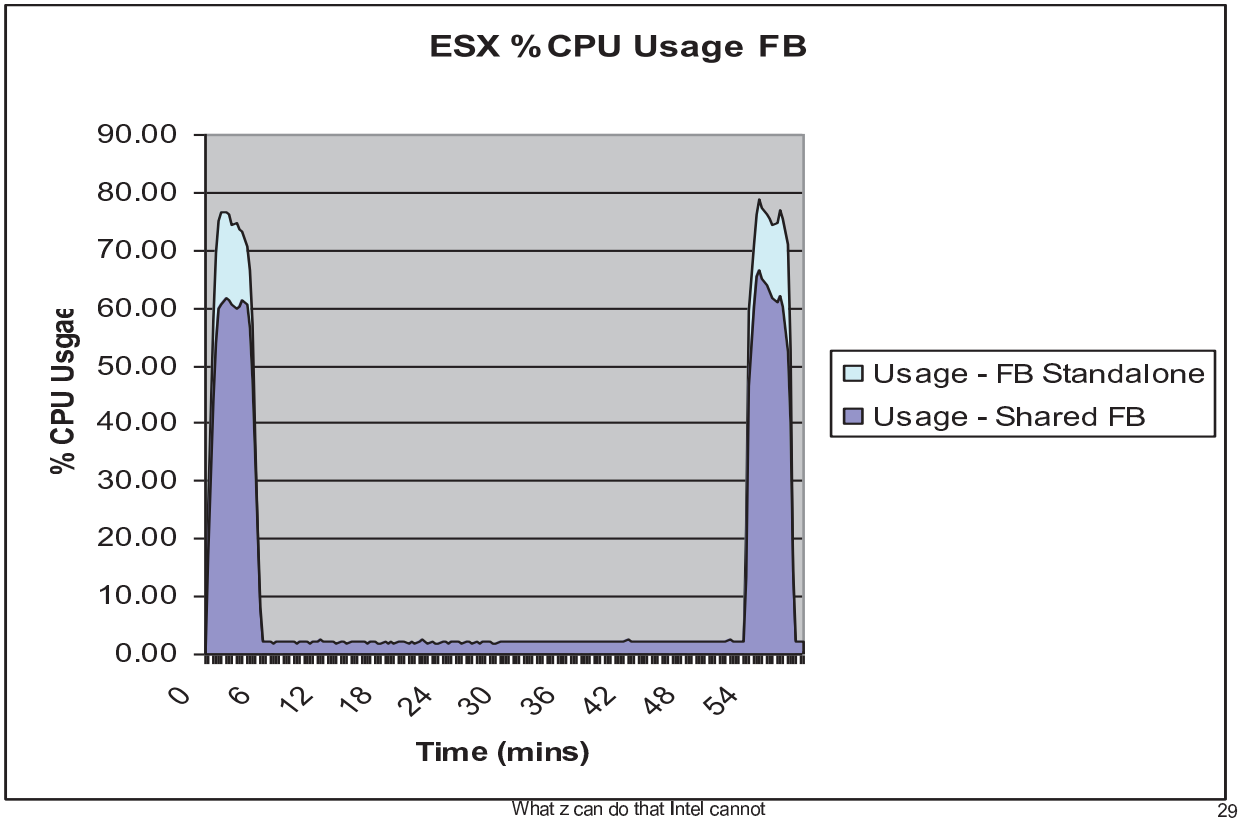

Inefficient Workload Management Has An Economic Impact

- Intel systems do not have an LPAR mechanism to efficiently separate workloads
  - VMware share and reservation mechanisms are limited
- Case study demonstrates 17% degradation in overall throughput and 18% degradation in maximum throughput with VMware on Intel
- It may not be practical to put high priority workloads and low priority workloads on the same Intel hypervisor
- Additional hardware needed to support different environments
- This translates into additional hardware and software acquisition costs as well as additional networking, management and environmental costs

zEC12 Advantages Over Intel

- zEC12 is designed to deliver more efficient processing capacity than Intel
  - Scales further for large enterprise workloads
  - Runs multiple concurrent workloads
  - Simplifies environments and reduces costs
- zEC12 has unique workload management capabilities
  - High utilization reduces cost
- zEC12 qualities of service are superior to typical Intel solutions
  - Administrator productivity, reliability, availability, capacity on demand, disaster recovery..

Large Systems With Centralized Management Deliver Better Labor Productivity

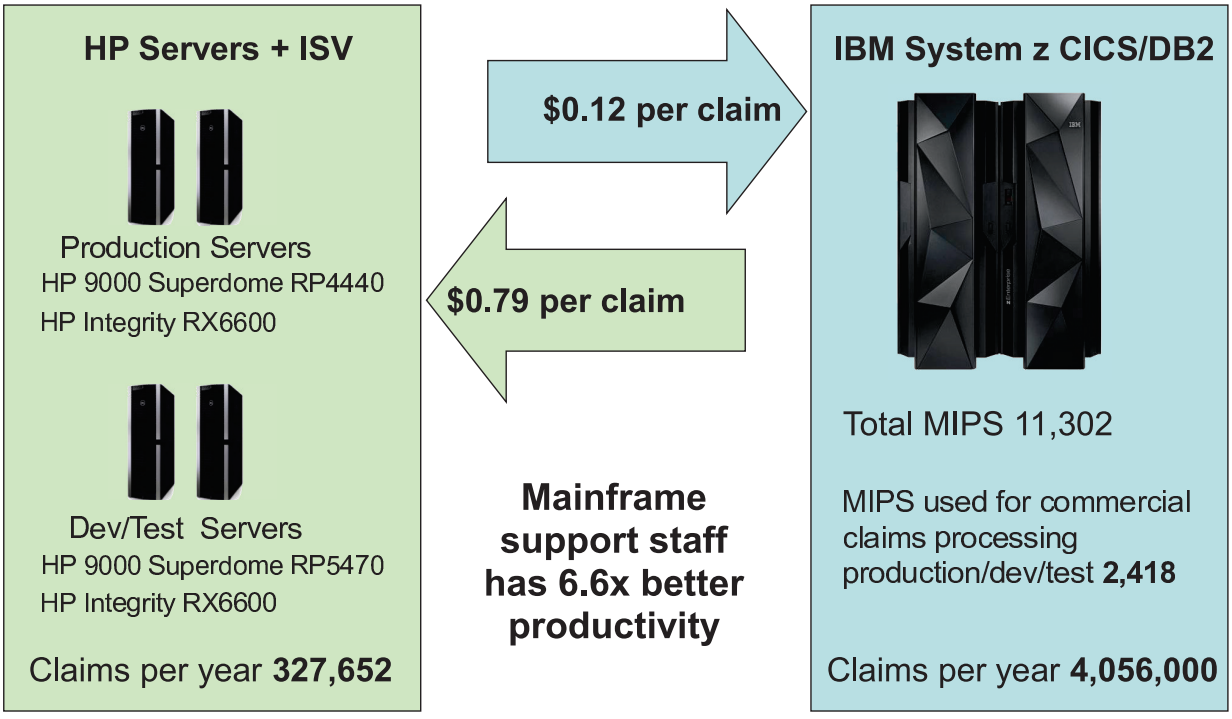

Source: CPO Eagle Team

zEC12 Availability Is Superior To Intel

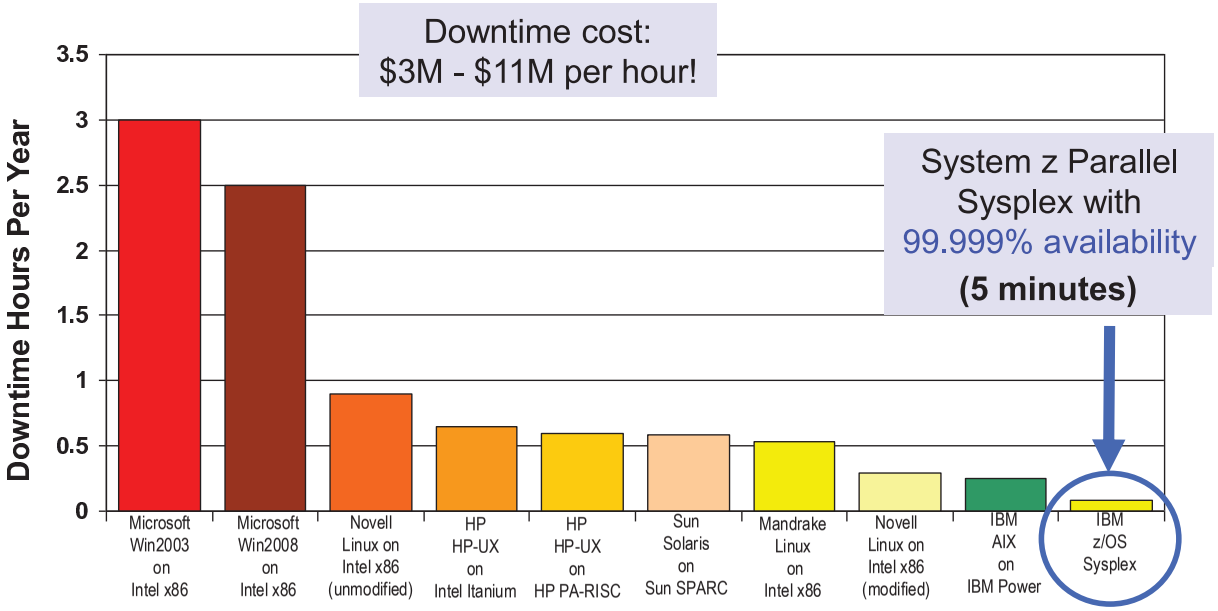

Source: ITIC: ITIC 2009 Global Server Hardware & Server OS Reliability Survey; July 2009; http://itic-corp.com/blog/2009/07/itic-2009-global-server-hardware-server-os-reliability-survey-results/; 400 participants in 20 countries; Results measured in minutes per year. All operating systems included in the survey are not included in this chart.

zEC12 Continues History Of Improvements To Reliability And Serviceability Not Found In Intel

Comprehensive, multi-layered strategy for reliability and serviceability

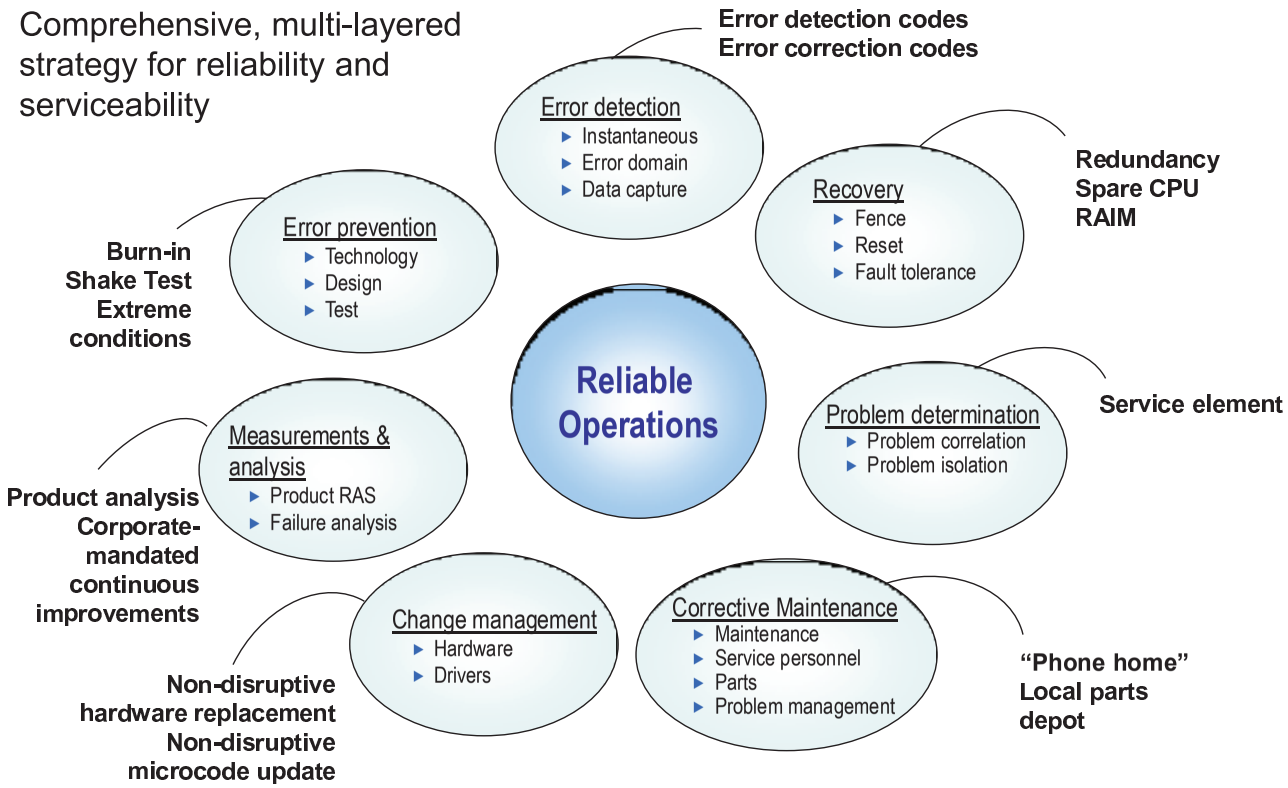

zEC12 Concurrent Operations With Hardware Repair Is Better Than Intel

Capability	zEC12	System x
ECC on Memory Control Circuitry	Transparent While Running	Can recognize/repair soft errors while running; limited ability with hard errors
Oscillator Failure	Transparent While Running	Must bring server down to replace
Core Sparing	Transparent While Running	Must bring server down to replace
Microcode Driver Updates	While Running	Some OS-level drivers can update while running, not firmware drivers; reboot often required
Book Additions, Replacement	While Running	Must bring server down to replace core, memory controllers, cache, etc.
Memory Replacement	While Running	Must bring server down to replace
Memory Bus Adaptor Replacement	While Running	Must bring server down to replace
I/O Upgrades	While Running	Must bring server down to replace (limited ability to replace I/O in some servers)
Concurrent Driver Maintenance	While Running	Limited – some drivers replaceable while running
Redundant Service Element	2 per System	"Support processors" can act as poor man's SE, but no redundancy

zEC12 Capacity on Demand - Not Available With Intel

- On/Off Capacity on Demand (On/Off CoD)
  - Flexible, easy, temporary additional capacity
  - Self-managed
  - Total flexibility within number of books installed
- New capacity is immediately available for work without service disruption
- Can be automated

One zEC12 Book with 36 Processors

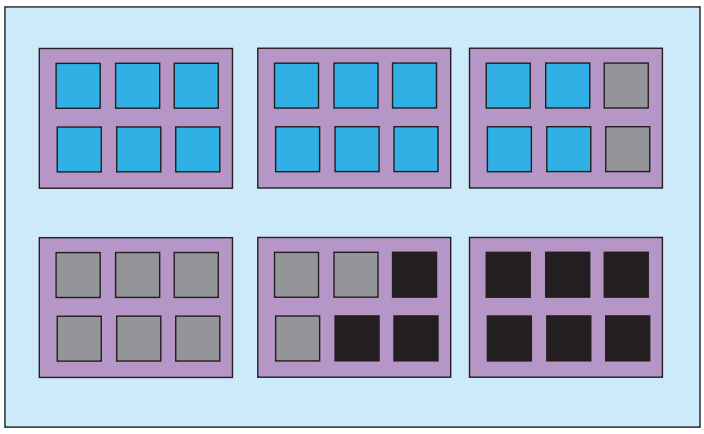

Active processors – pay full price

Inactive processors (On/Off CoD) – pay only 2% of full price

Dark processors (unused) – no charge

Parallel Sysplex Provides Capabilities Not Available With Intel

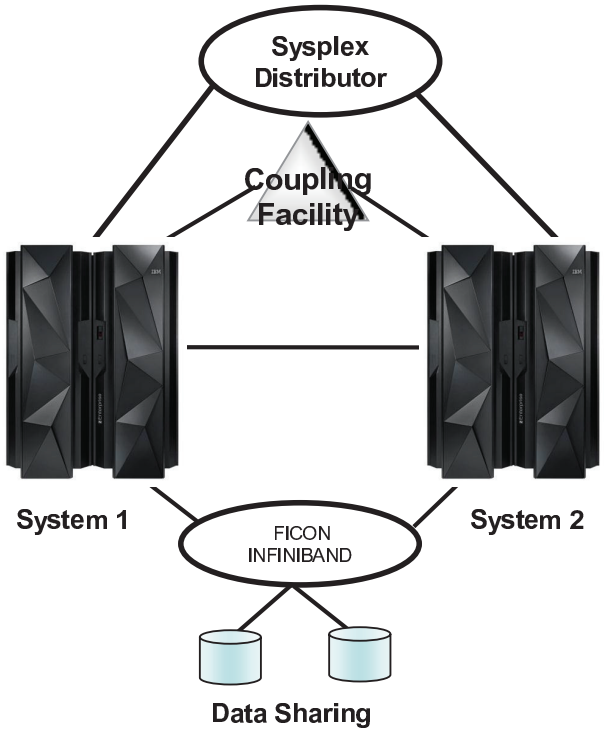

Rolling software update via automatic sysplex failover

- Shutdown LPAR on system 1 for maintenance
- Upgrade OS and middleware on LPAR
- Bring up LPAR on system 1
- Shutdown LPAR on system 2 for maintenance
- Upgrade middleware on LPAR
- Bring up LPAR on system 2

z/VM Consolidation Offers Options Unavailable on ESX

IBM

z/VM Hosting Linux on System z Guests

IBM zEnterprise running z/VM hosting Linux on System z guests offers options you can't find with ESX. z/VM offers:

- **Real CPU sharing -** Architecturally limitless; more than 60 VMs per CPU (workload dependent)
- **Architectural maximum number of VMs -** Thousands per copy of z/VM
- **Non-disruptive Real capacity on demand**
- **Run multiple hypervisor copies per server -** up to 60 copies of z/VM on one mainframe
- **Virtualization -** z/VM can run as a guest of itself to multiple levels (z/VM on z/VM on z/VM on …)

Competition

ESX hosting Linux guests

ESX can't match z/VM on:

- **Real CPU sharing –** Up to 25 VMs per CPU (workload dependent)
- **Architectural maximum number of VMs -** 512 per copy of ESX
- **No Real capacity on demand**
- **Only one hypervisor copy per server**
- **Virtualization –** No ability to run ESX as a guest in ESX

z/VM 6.2 information :http://publib.boulder.ibm.com/infocenter/zvm/v6r2/index.jsp

ESX 5.1 information: http://www.vmware.com/pdf/vsphere5/r51/vsphere-51-configuration-maximums.pdf

The information is based on a point in time (08/28/2012).

zEC12 GDPS – Systematic Failover Capability Beyond Typical Intel Solutions

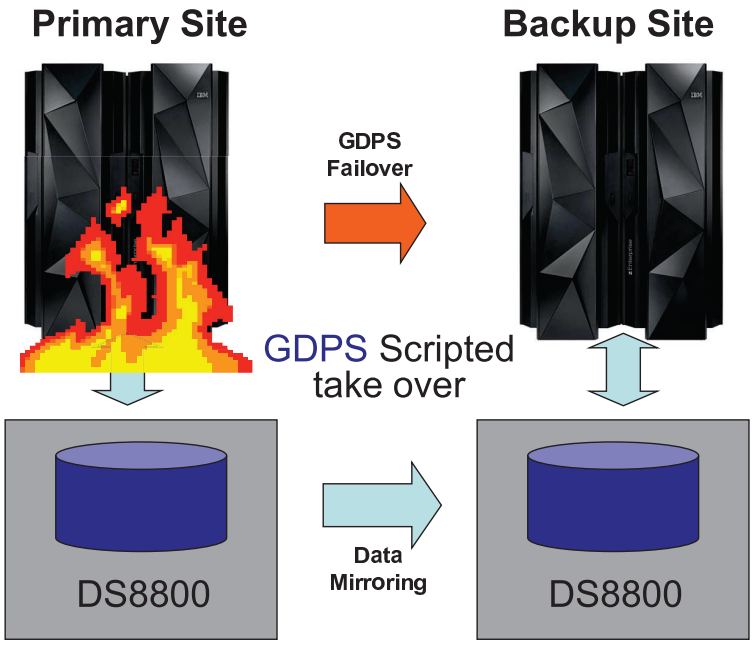

- Site Failover
  - Failover to secondary site in case of complete site failure
- Data Mirroring
  - Protect data in the event of a disk system failure

Supports systematic DR for virtualized Linux environments

zEC12 Advantages Over Intel

- **zEC12 is designed to deliver more efficient processing capacity than Intel**
  - More processors, cache, and I/O capability than Intel
  - Parallel Sysplex provides scale and reliability
  - Run larger and more workloads than Intel
  - Reduces cost per workload compared to Intel scale out
- **zEC12 has unique workload management capabilities**
  - Enables high utilization with mixed workloads
- **zEC12 qualities of service are superior to typical Intel solutions**
  - Administrator productivity, reliability, availability, capacity on demand, disaster recovery..

System Building Blocks Require Evolutionary Feedback

"Organic chemistry"	*"DNA"*	*"Body"*	*"Brain"*	*"Skills"*	*"Talent"*
• Raw components & technology of computer • Lithography, process	• The CPU chip • Cache • Memory • Buses • Plus associated support chips	• Holistic design of the hardware • Channels • Power systems • Connections • Storage • Cooling • RAS • Hybridization	• Software infrastructure • Enablement of new technology • Maximize utilization & efficiency • Hypervisors • Resource allocation	• Specific capabilities • Exploitation of hardware	• Differentiation • Provision of business componentry (DB, OLTP…) • APIs

Intel — **OEM, ISV1, ISV2** — **ISV3** — **ISV4**

IBM optimized

IBM zEnterprise EC12: An Optimized System

• eDRAM cache integrated on processor chip • Efficient packaging • Optimized high-K metal gate transistor design • 2x transistor density for logic and on-chip memory	• World's fastest chip at 5.5 GHz • More cache for superior data serving • On-chip crypto and data compression • Instruction co-optimized with compilers, Java, DB2 • Transactional Execution • Runtime Instrumentation • Decimal Format Conversions • 2 GB page frames	• Optimized I/O pathways & subsystems with massive scale • Best-of-breed reliability, availability & serviceability • RAIM memory for high availability • I/O cryptographic coprocessor • Highest commercial security certification in industry (EAL5+ & FIPS 140-2 level 4) • Multi-platform design of mainframe and distributed technologies • Designed for highest utilization with heterogeneous workloads	• Intelligent policy management of mainframe & distributed technologies • Granular, scalable virtualization of servers, memory and I/O with low overhead • PR/SM and z/VM offer two-tier approach for superior virtualization supporting native Linux • Dispatching on z/OS keeps software close to cache for optimized performance • Highly available and resistant partition architecture suitable for private cloud	• Java, ASM, C/C++, PL/I Compilers optimized to exploit hardware architecture • Java Runtime Environment exploits dynamic optimization (transactional execution)	• WebSphere, MQSeries, DB2 exploit multi-system workload management, scale and availability • Middleware optimized and tuned to scale up • Tivoli optimized for operations management and security • High volume, high performance, low cost transaction processing with CICS and IMS

There's More… System z Has The Potential To Be A Complete Environment

Operational Efficiency

- Mature processes
- Change management
- SLAs
- Governance
- ISV eco-system
- Printing, tapes

Pro-growth Economics

- Decreasing incremental charges
- Infinite capacity
- Elastic scalability
- Comprehensive instrumentation