What System z Can Do That Intel-Based Systems Cannot

Chicago Area VM (and Linux) Enthusiasts (CAVMEN)
17 October 2013

What System z Can Do That Intel-Based Systems Can't

1. Transaction processing at scale

Servers Based on latest Intel technology (Sandy Bridge)

System z

Transaction Processing At Scale Requires A Balance Of Capabilities

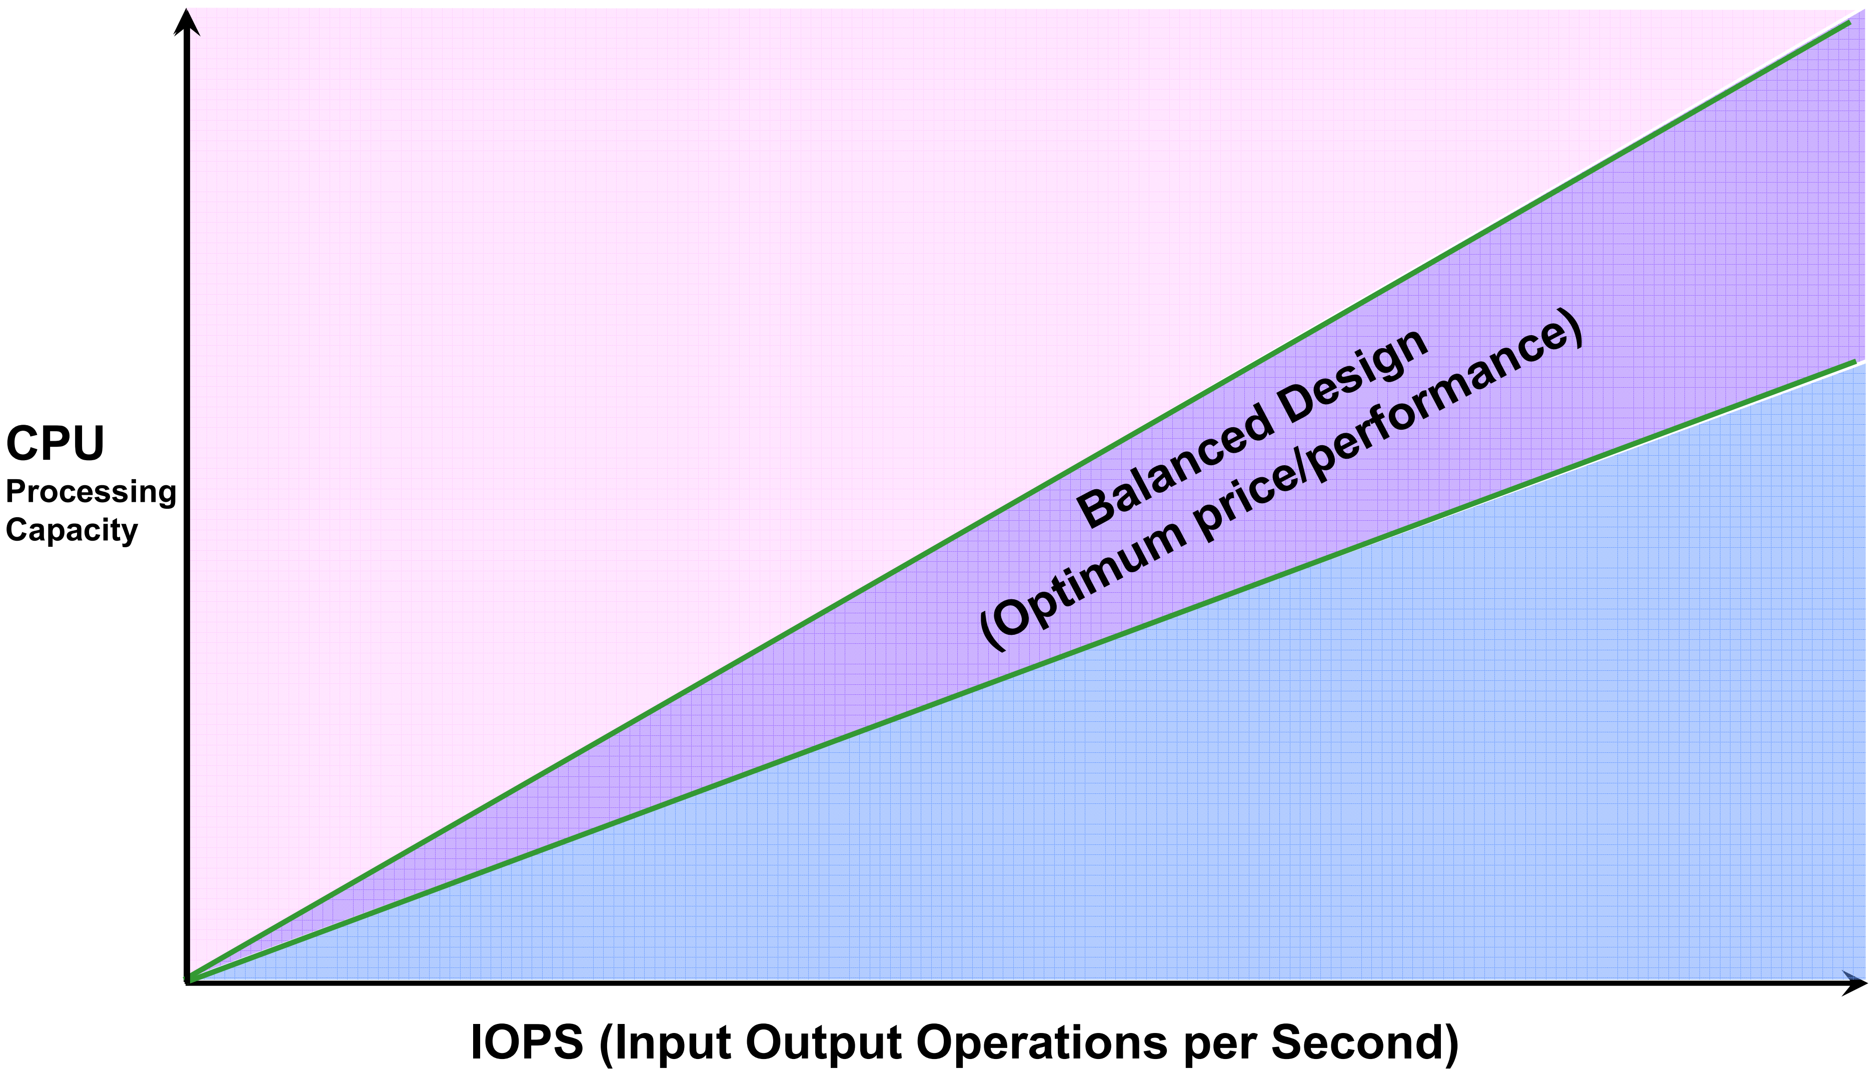

zEnterprise EC12 Design Is Unique And Optimized For Transaction Efficiency At Scale

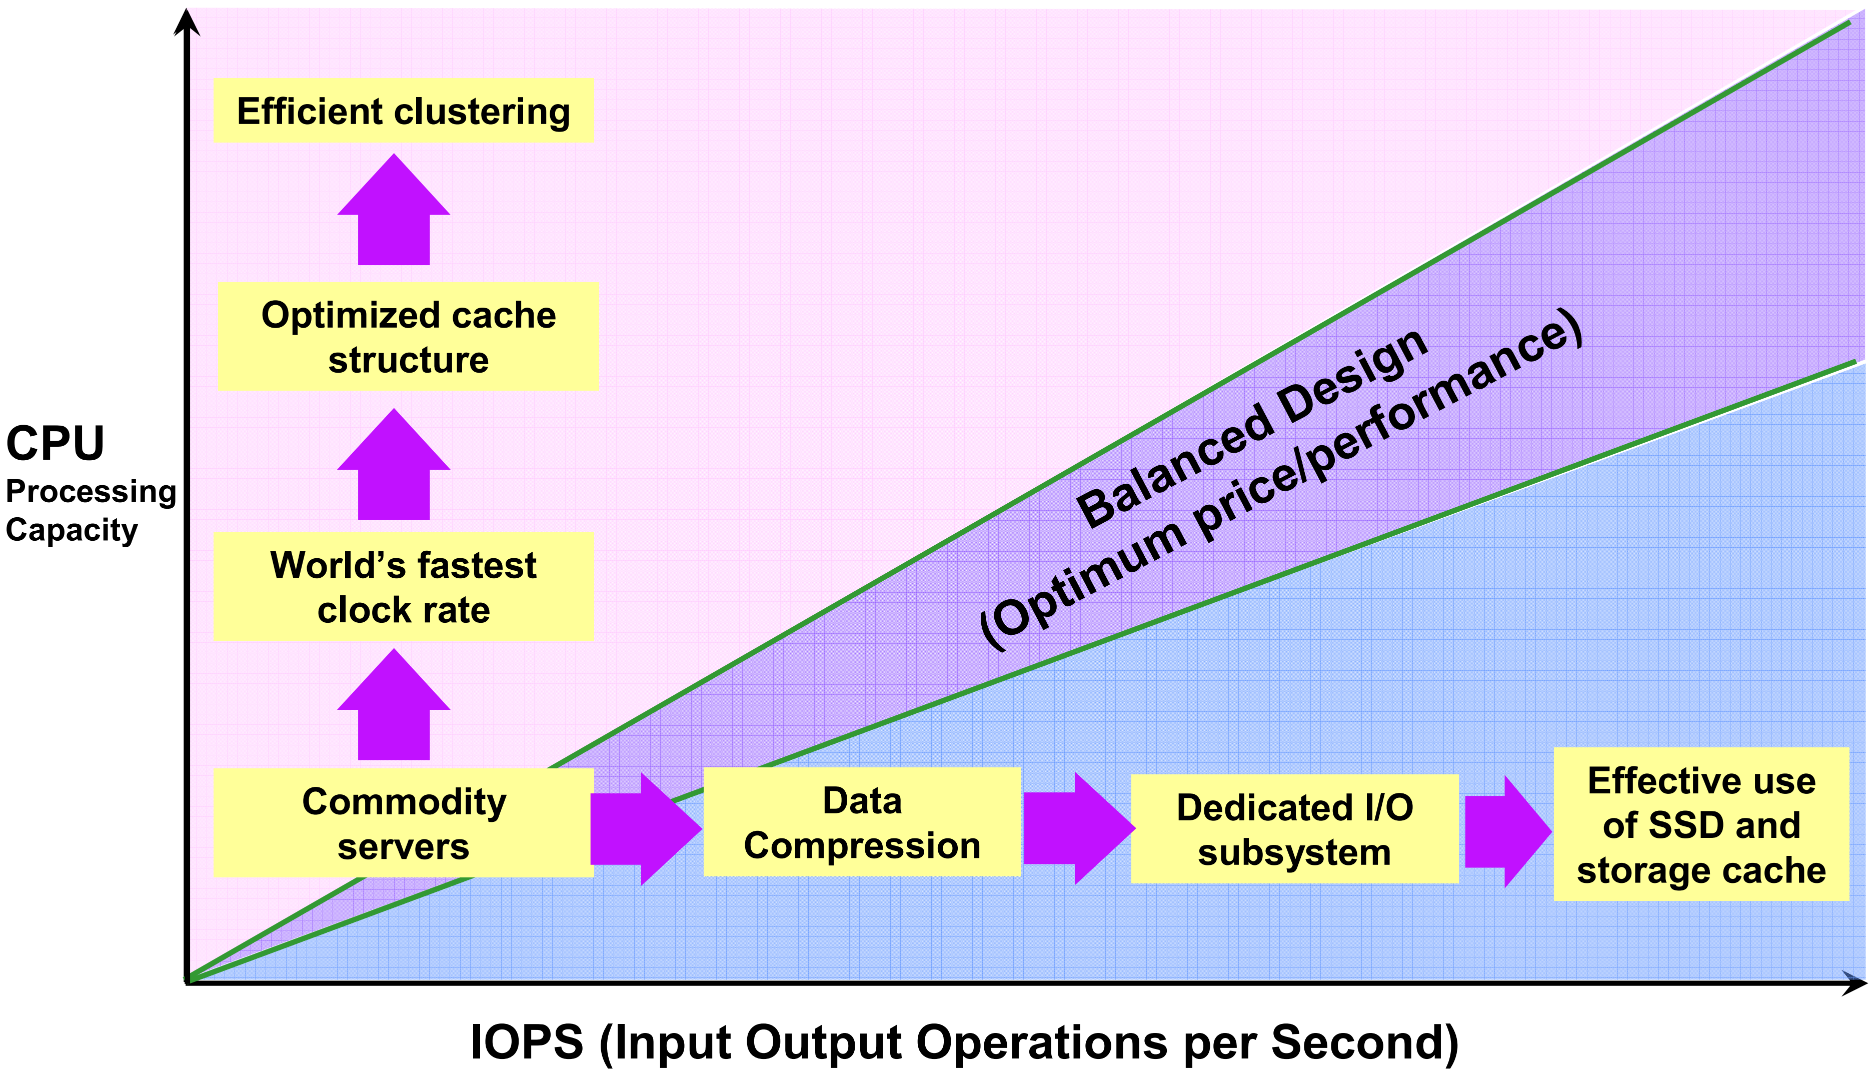

System z Delivers More Raw Processing Capacity Than Intel

	zBC12	zEC12
*World's **fastest** clock speed*	**4.2 GHz**	**5.5 GHz**
Total cores	**18**	**120**
Configurable cores	**6 GPs / 13 IFLs**	**101**
General processor core performance	**1064 MIPS**	**1,514 MIPS**
Specialty processor core performance	**1064 MIPS**	**1,514 MIPS**
Total Capacity	**4958 / 8733 MIPS**	**78,426 MIPS**

Maximum x86 clock speed = **3.4 GHz**

Maximum x86 cores = **32**

Intel Sandy Bridge

System z Has More Cache Than Intel To Support Multiple Concurrent Workloads

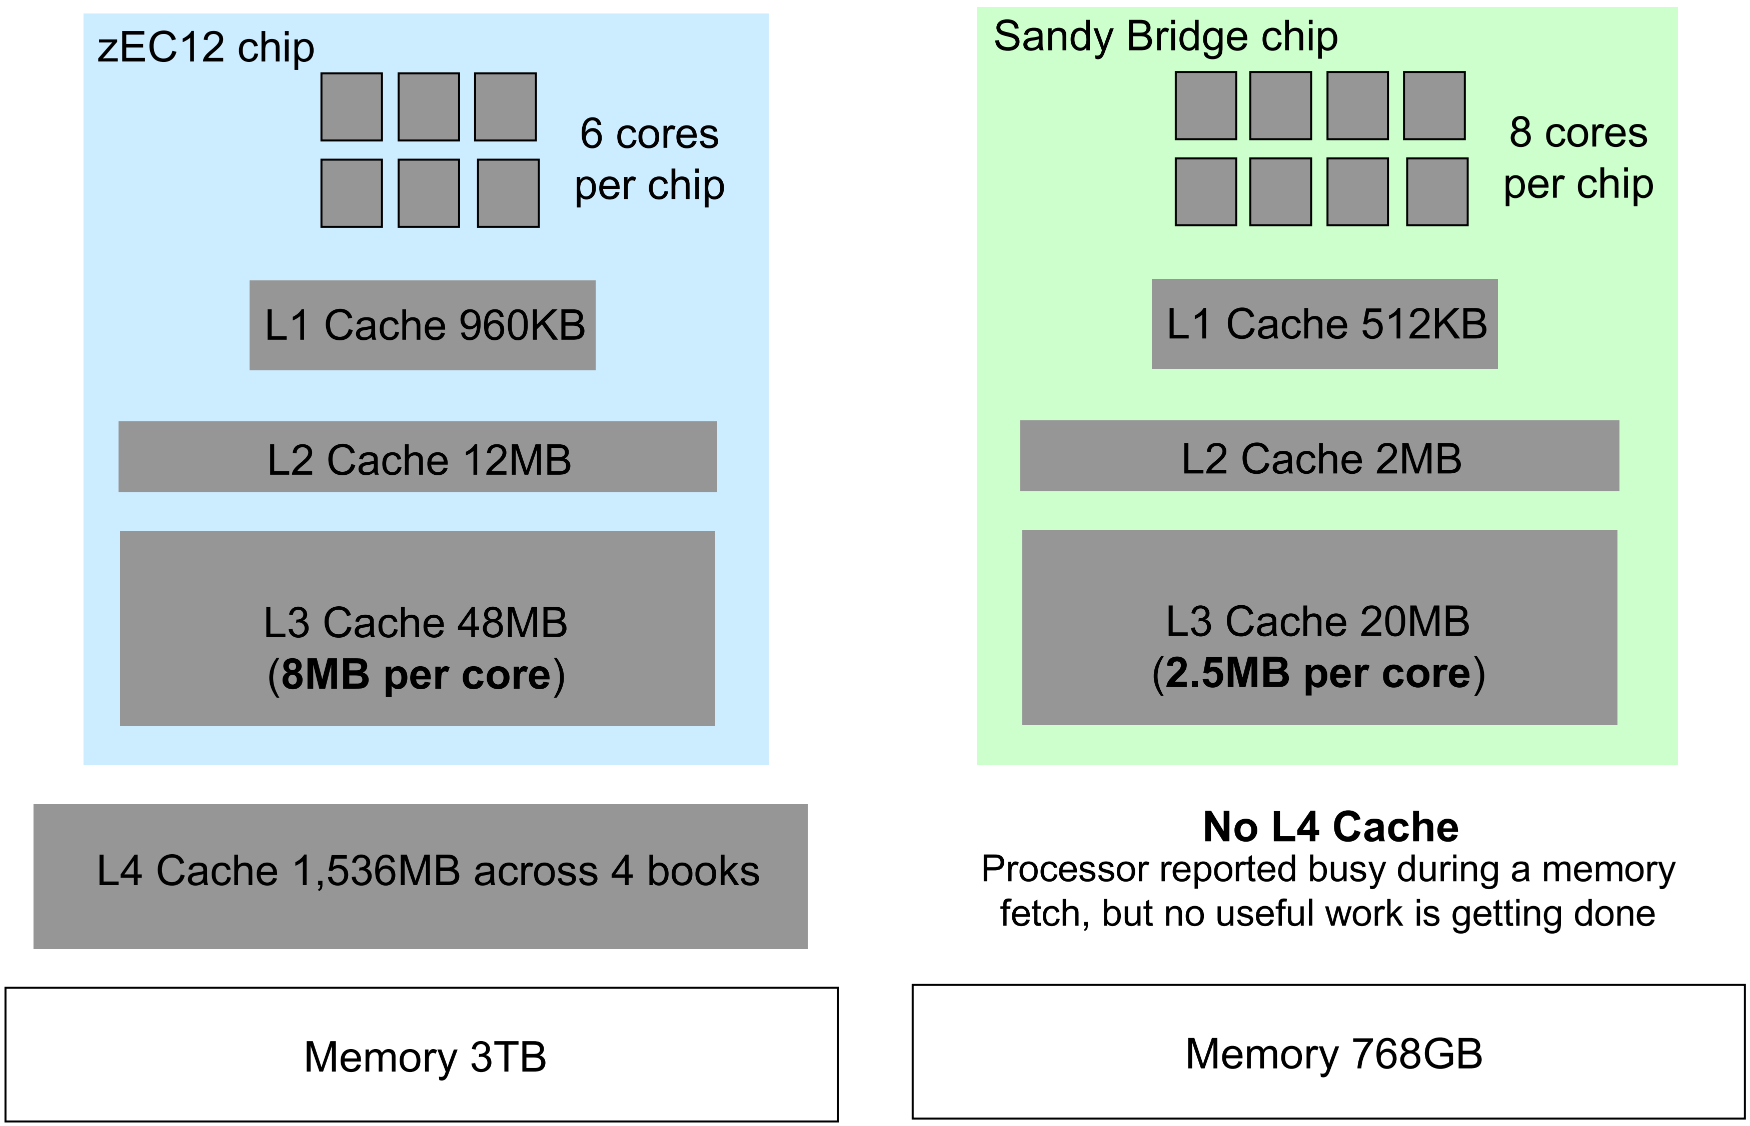

Intel Servers Slow Down Under Cache Intensive Workloads

- Multiple concurrent processes introduces cache contention
  - Example: 5 processes each with 70MB working set size
- Intel workloads significantly slowed due to cache contention
- System z with z/OS showed results 8X faster than Intel system

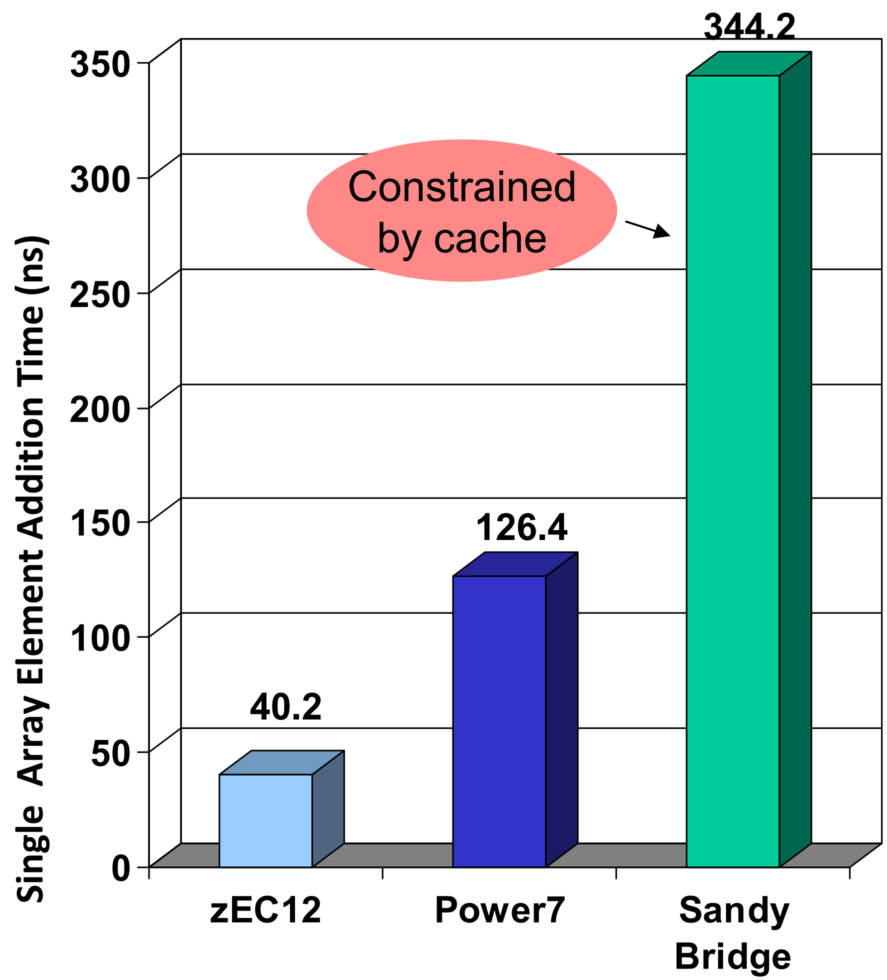

Parallel Sysplex Enables System z To Scale To Capacities Far Beyond What Intel Can

Single System Sysplex

Parallel sysplex clustering delivers **highest availability**

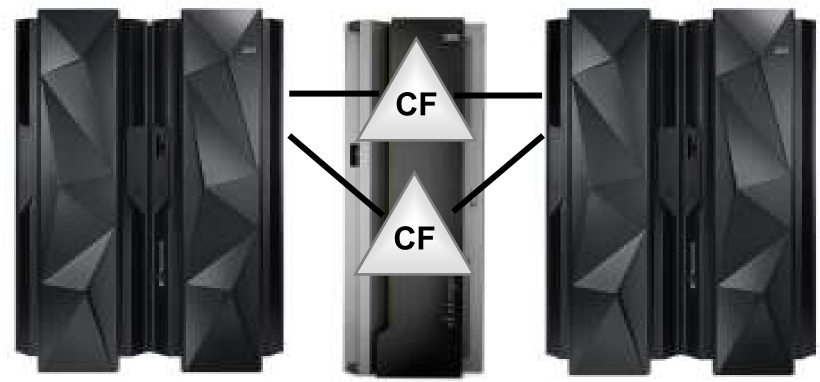

External Coupling Facility
(Can be different class server)

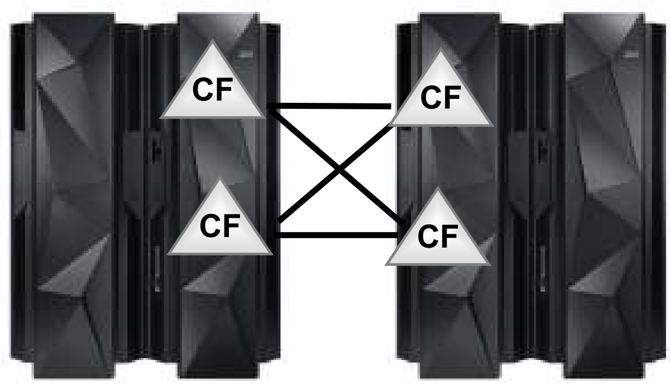

Cross Connected Servers with internal Coupling Facilities

Potentially 2.5 million MIPS per 32-way cluster

Supports rolling software updates via automatic sysplex failover

Real-World Benchmarks Show System z Runs Bigger Workloads Than Intel

- **Kookmin Bank**
  - **IBM System z and DB2**
  - TCS BaNCS
  - **15,353** Transactions/second
  - **50** Million Accounts
  - IBM benchmark for customer
  - DB2 V9, CICS 3.1, z/OS V1.8

- **State Bank of India**[1]
  - **HP Superdome**
  - **TCS BaNCS**
  - **10,716** Transactions/second
  - **500** Million Accounts
  - Largest banking benchmark performance claimed by HP

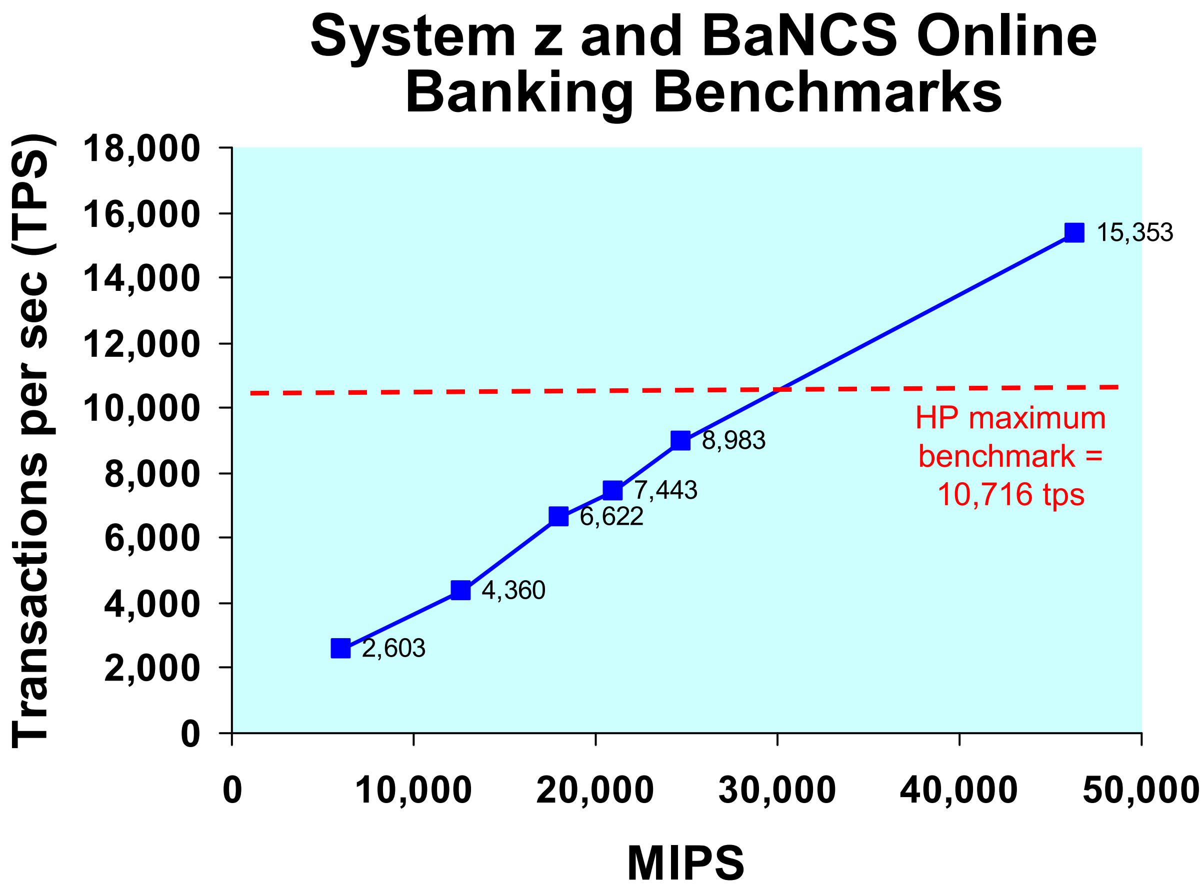

[1] Source: http://www.enterprisenetworksandservers.com/monthly/art.php?2976 and *InfoSizing FNS BANCS Scalability on IBM System z – Report Date: September 20, 2006;* Clement Report; http://h20195.www2.hp.com/v2/GetPDF.aspx/4AA1-4027ENW.pdf Feb 2010

Sustains High CPU Utilizations With No Service Degradation

Recent Daily Samples Indicating Stress Extended Periods above 97%

Monday June 10, 2013

8:15 to Noon

06/10	8.15	91.90
06/10	9.15	91.00
06/10	9.30	95.30
06/10	9.45	99.10
06/10	10.00	99.90
06/10	10.15	98.80
06/10	10.30	99.80
06/10	10.45	99.80
06/10	11.00	100.10
06/10	11.15	100.10
06/10	11.30	99.40
06/10	11.45	97.60
06/10	12.00	99.70

Noon to 5PM

06/10	12.15	97.00
06/10	12.30	97.00
06/10	12.45	90.40
06/10	13.00	91.20
06/10	13.15	92.90
06/10	13.30	95.70
06/10	13.45	96.00
06/10	14.00	99.40
06/10	14.15	100.00
06/10	14.30	99.10
06/10	14.45	98.10
06/10	15.00	97.60
06/10	15.15	99.60
06/10	15.30	99.80
06/10	15.45	99.00
06/10	16.00	99.60
06/10	16.15	97.40
06/10	16.30	92.40
06/10	16.45	90.50
06/10	17.00	94.20

Sustained Peaks	Number of Days	Dates
100%	2	4,10
99%+	6	3,4,6,10,19, 25
97%+	9	3,4,6,10,11, 19,24,25,26

zEC12 Has A Dedicated I/O Subsystem Which Can Deliver 1.7M I/O Operations Per Second

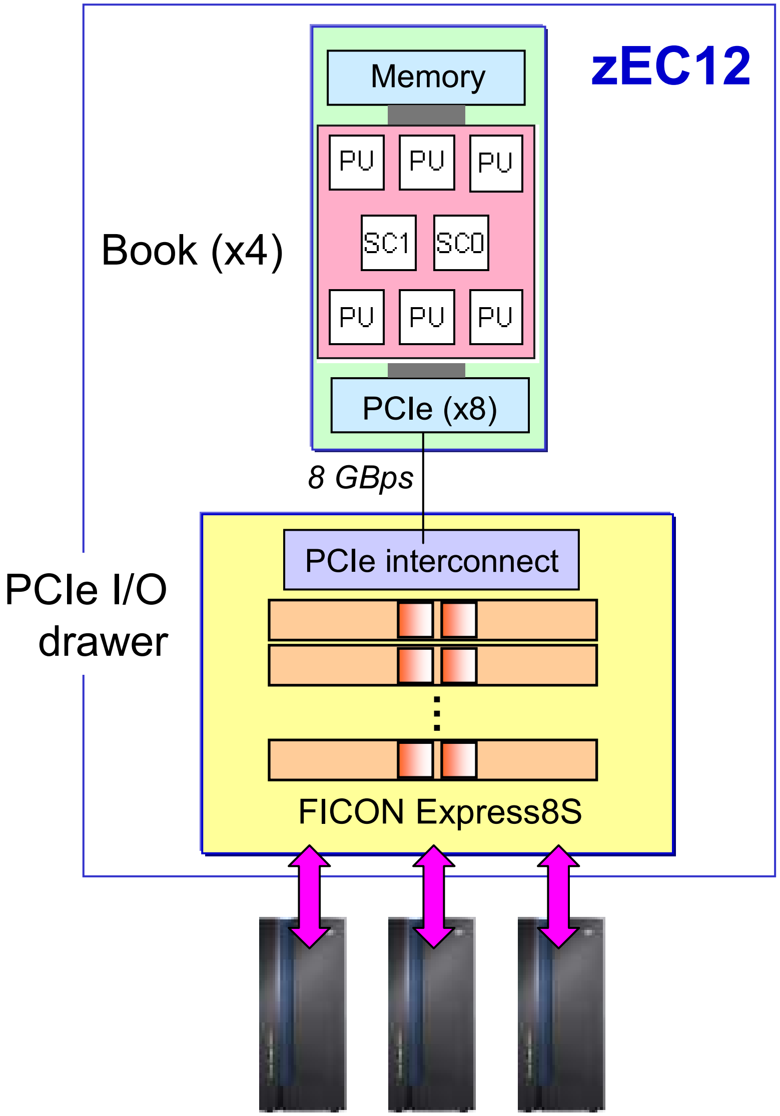

- **Up to 16 dedicated System Assist Processors (SAPs)**
  - I/O requests are offloaded to SAPs
  - SAPs can sustain up to 2.4M IOPS*
  - I/O subsystem bus speed of 8 GBps
  - Number of SAPs increases from 2 to 16 according to system size
- **Up to 160 physical FICON cards for I/O transfers**
  - Up to 320 RISC processors (2 per card)
  - Up to 320 FICON channels (2 per card)
  - 8 Gbps per link, 288 GB/Sec I/O aggregate per zEC12
- **IBM DS8800 Storage System**
  - Up to 440K IOPS capability
- **Delivers I/O efficiency at scale**

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

I/O Bandwidth Is Important For Critical Data Workloads

Bank of China System z Benchmark required huge I/O bandwidth capacity

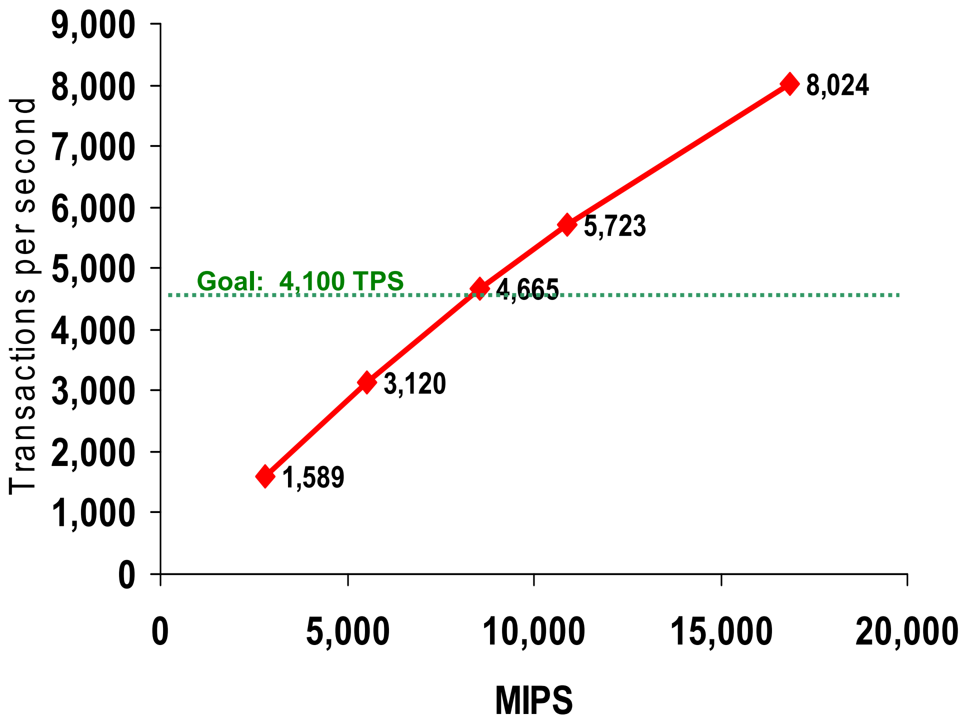

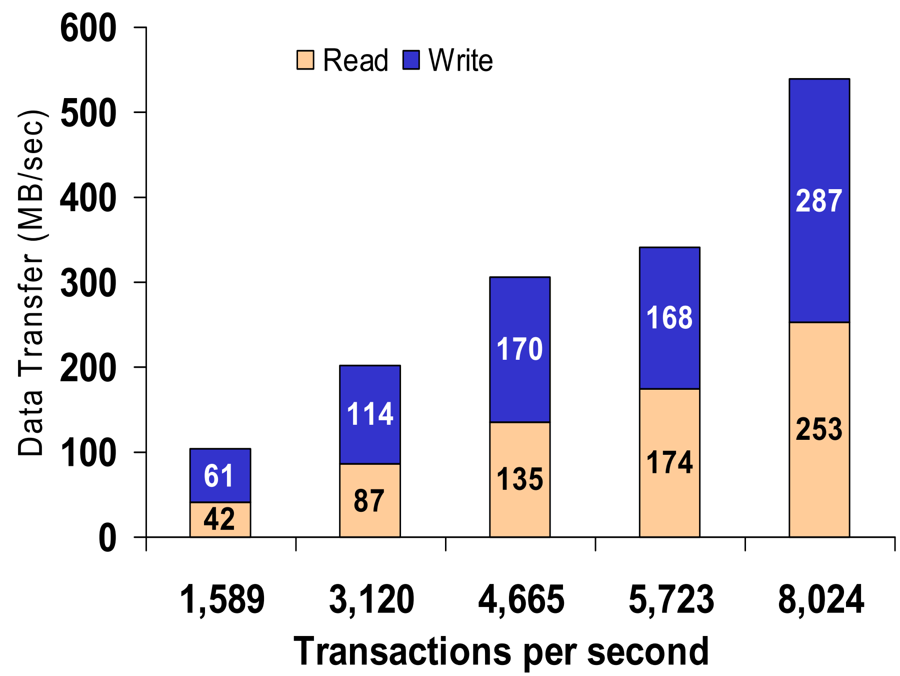

Bigger workloads require more I/O bandwidth capacity

System z's Dedicated I/O Subsystem Delivers More I/O Processing Capacity Than Intel

- Intel's performance degrades as I/O demand increases
  - No dedicated I/O subsystem
- Test case scenario: Run multiple virtual machines on x86 server
  - Each virtual machine has an average I/O rate
  - x86 processor utilization is consumed as I/O rate increases

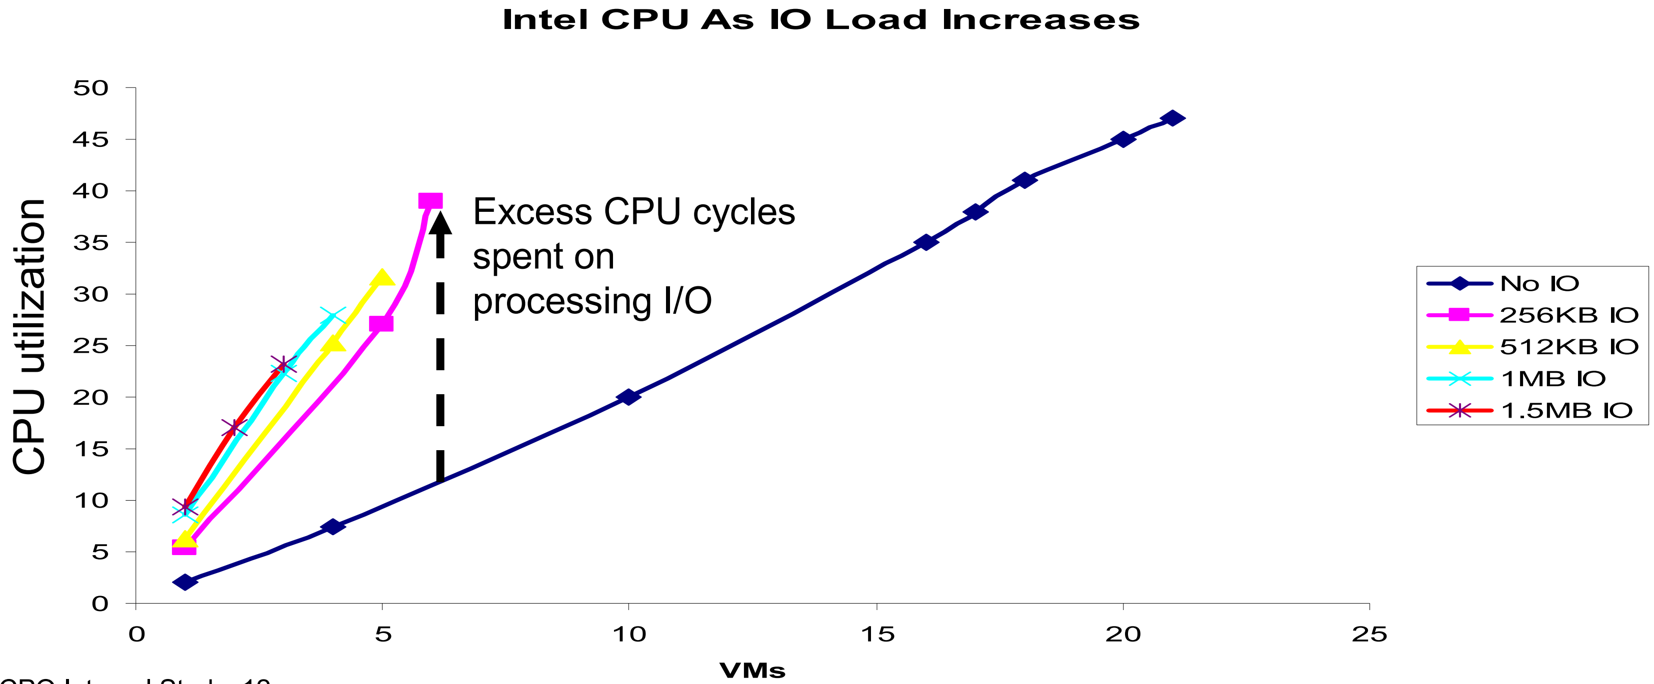

Source: CPO Internal Study. 12-core Westmere EP with KVM. FB at 22 tps with varying IO per transaction.

Test Configuration

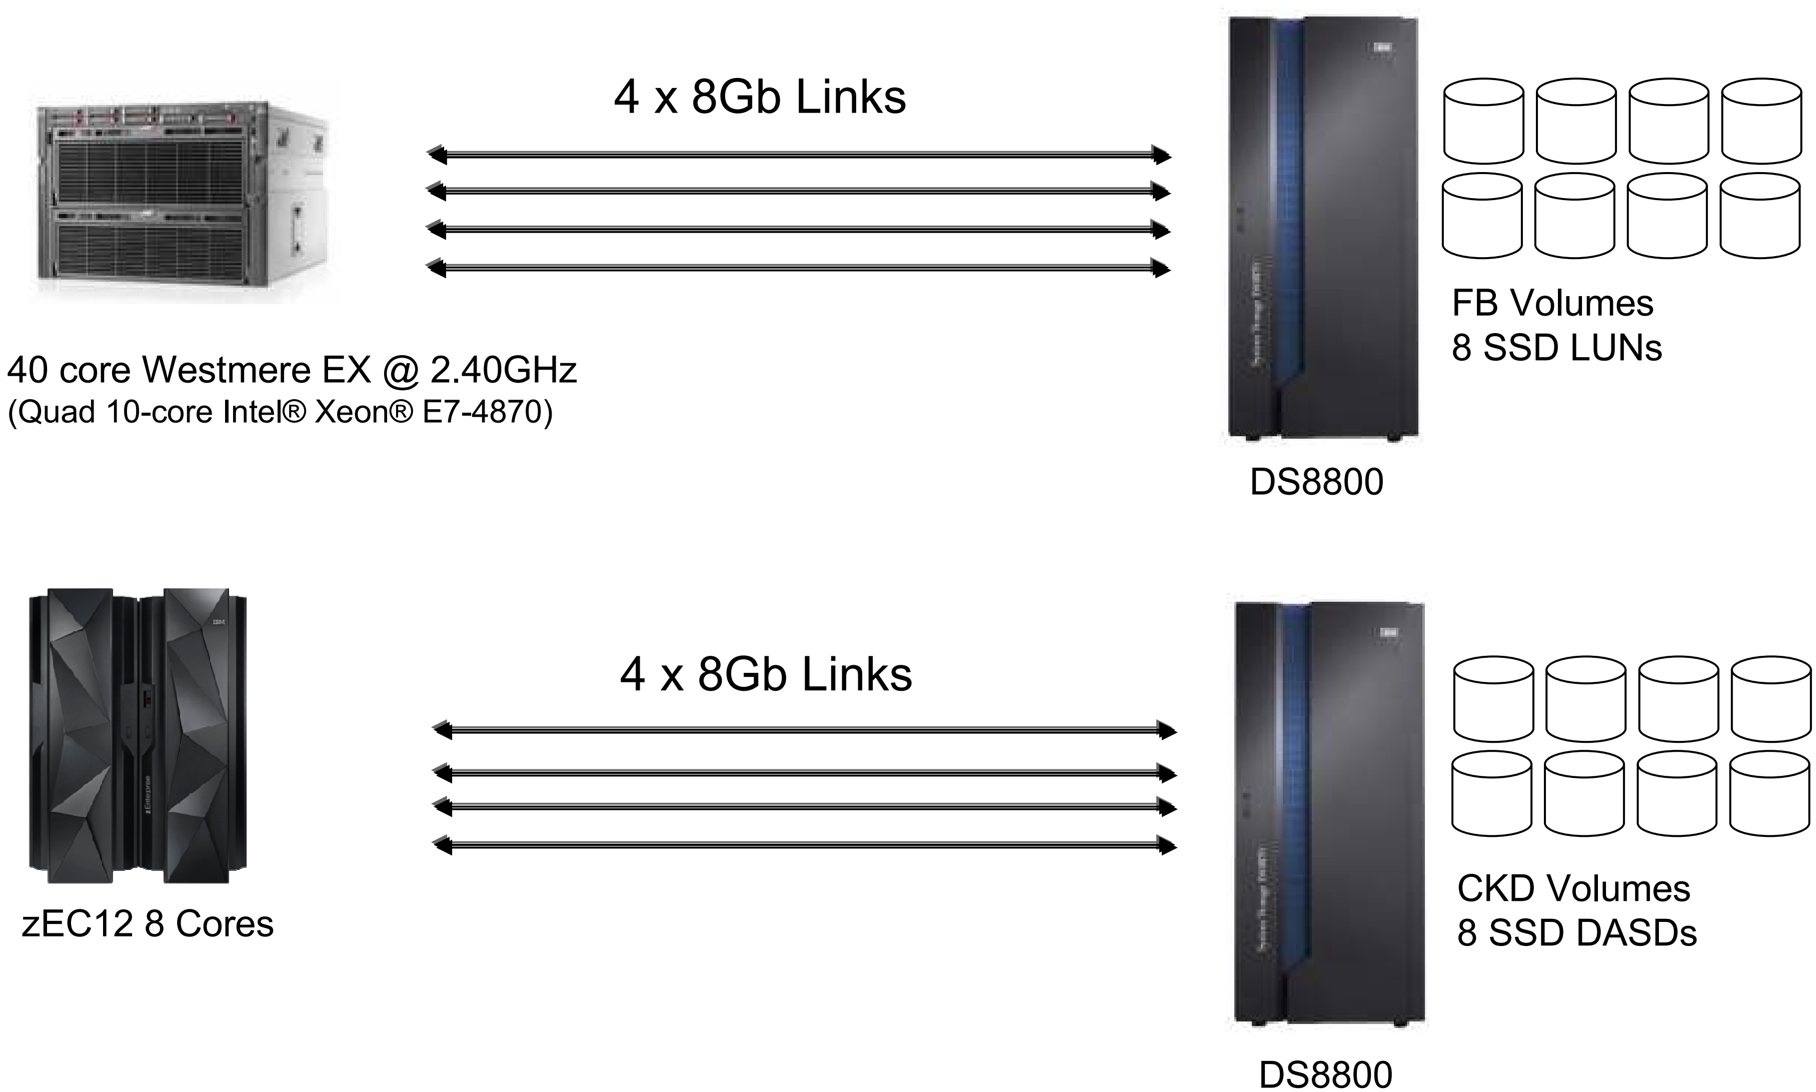

Elapsed Time Comparison

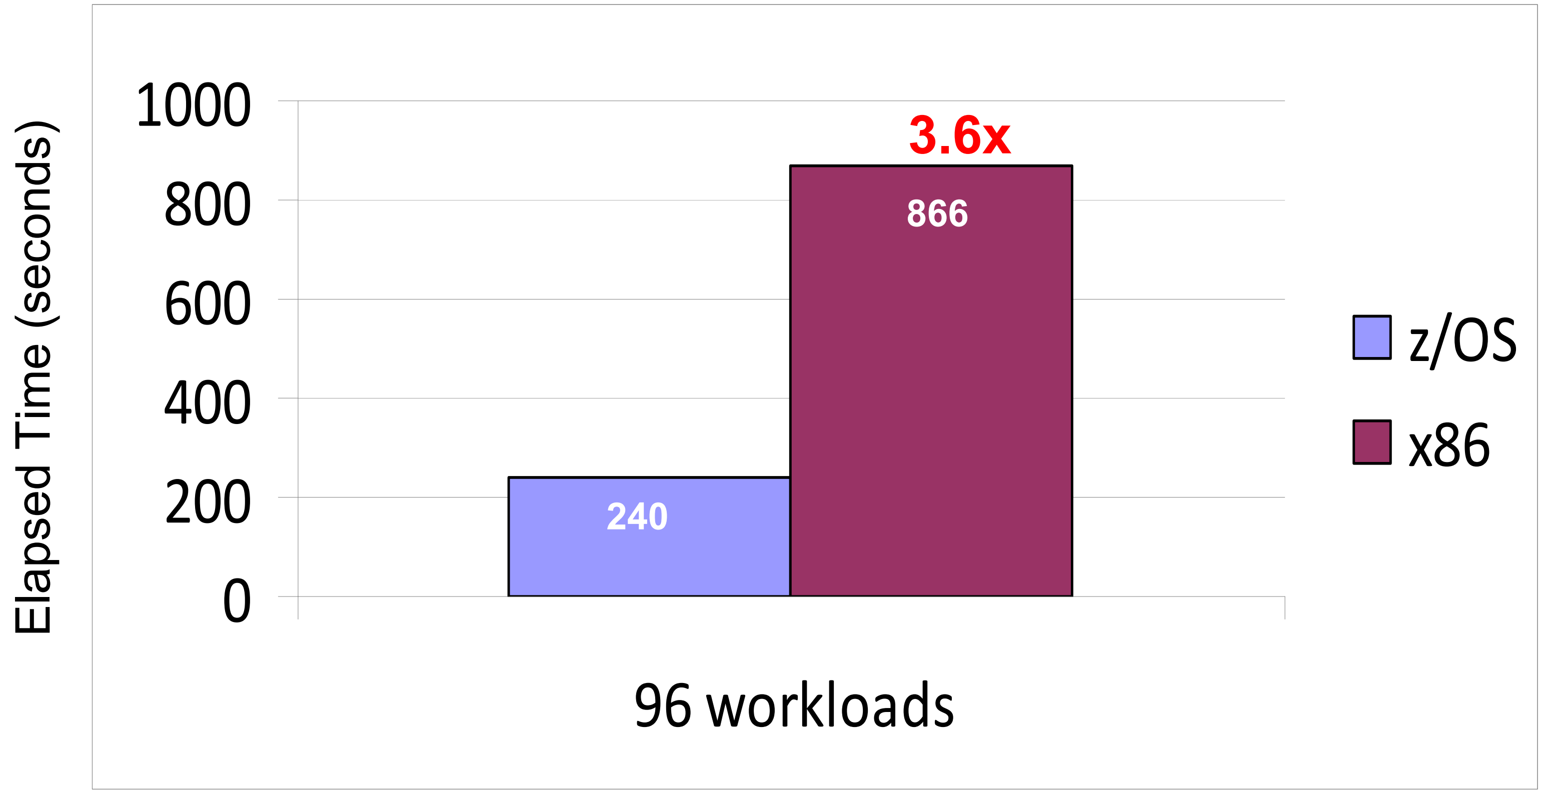

Important note: 96-wide was the max point to which the test cases could be run before the backend storage available for this study became a bottleneck. System z channel utilization is at 27% and the IOPs are only at 2.5%, leaving z considerable room to grow

Elapsed Time Normalized to CPU

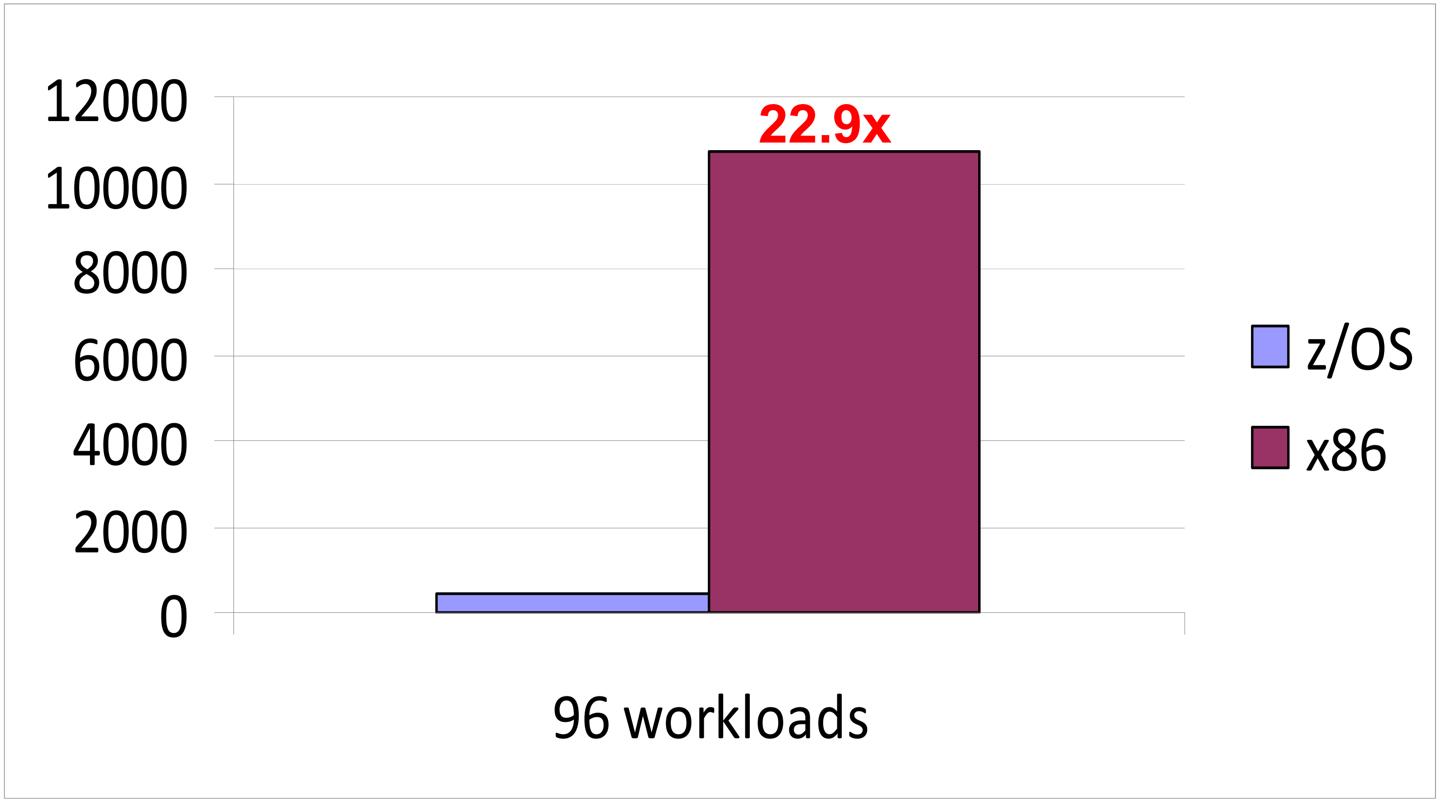

Important note: 96-wide was the max point to which the test cases could be run before the backend storage available for this study became a bottleneck. System z channel utilization is at 27% and the IOPs are only at 2.5%, leaving z considerable room to grow

Compare Impact of Virtualization

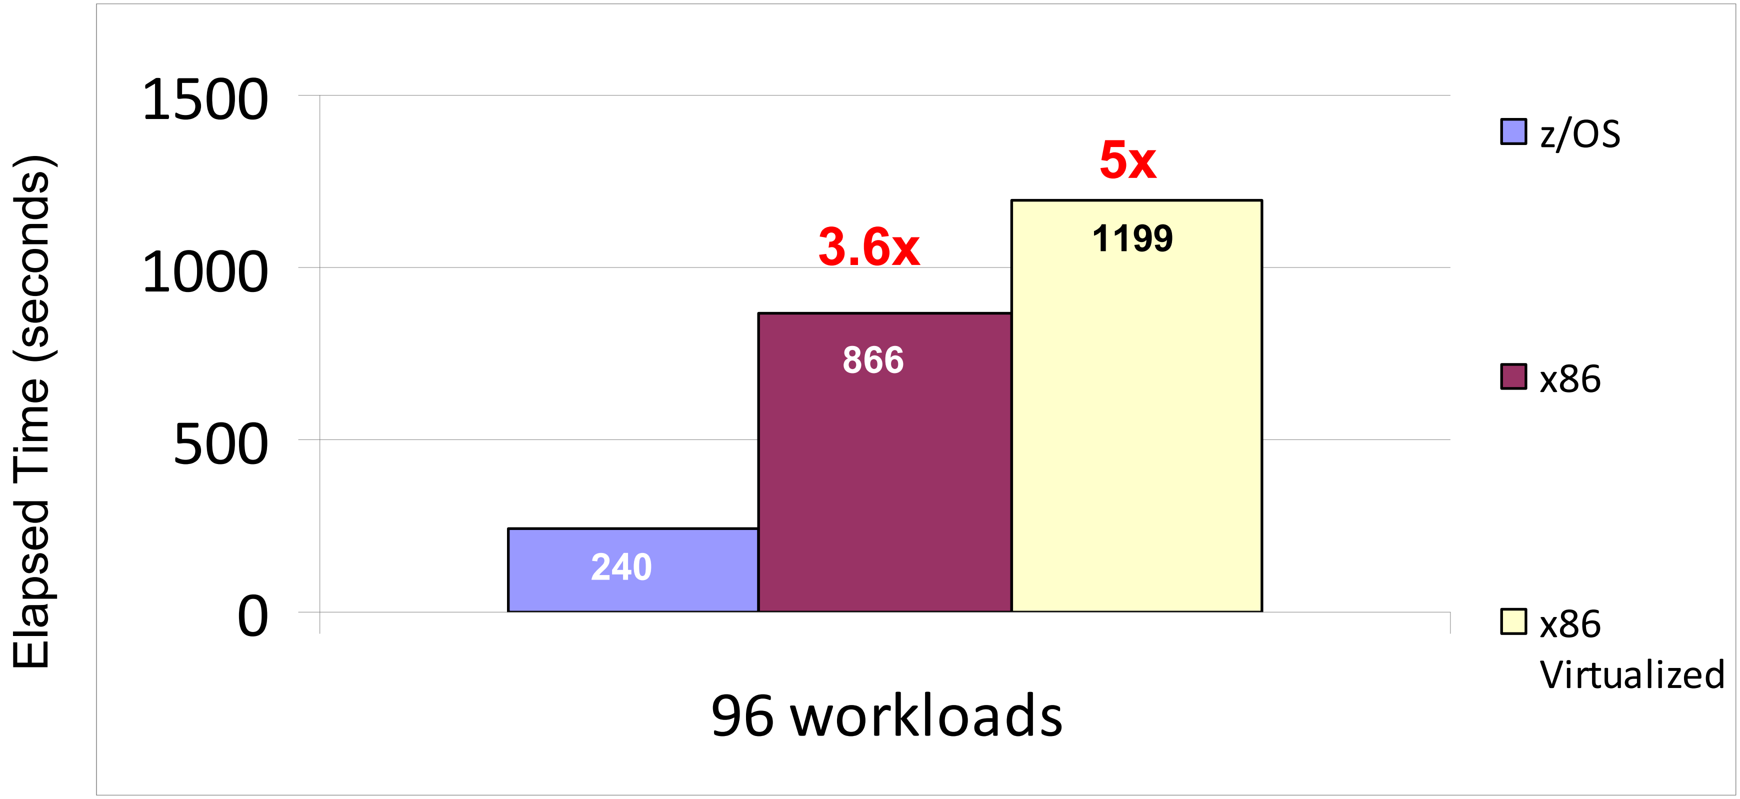

Important note: 96-wide was the max point to which the test cases could be run before the backend storage available for this study became a bottleneck. System z channel utilization is at 27% and the IOPs are only at 2.5%, leaving z considerable room to grow

z/OS Database Workloads Benefit From Higher I/O Bandwidth

Competitor DB on Intel

SAP Applications

Database

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

128 DB cores

DB2 on z/OS

SAP Applications

Database

zEC12 2-way data
sharing Sysplex

44 DB cores
14 ICF cores

*A world record at **half** the cost!*

Database Unit Cost
$0.30/Postings per hour

Postings per Hour	42.0M
# of Accounts	90M
Hardware	$0.63M
Software	$11.98M
Total (5 yr. TCA)	$12.61M

Database Unit Cost
$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 (5 yr. TCA)	$8.73M

Cost of platform infrastructure for comparative transaction production. Cost of packaged application software not included. List prices used.

Consolidated Oracle DB Workloads Benefit From Linux On System z's I/O Bandwidth

Which platform provides the lowest TCA over 3 years?

Oracle DB workload

Customer Database Workloads each supporting 18K tps

Oracle Enterprise Edition
Oracle Real Application Cluster

3 Oracle RAC clusters
4 server nodes per cluster

12 total HP DL580 servers (192 cores)

$13.2M (3 yr. TCA)

3 Oracle RAC clusters
4 nodes per cluster
Each node is a Linux guest
zEC12 with 27 IFLs

$5.7M (3 yr. TCA)

Half *the cost*

TCA includes hardware, software, maintenance, support and subscription.
Workload Equivalence derived from a proof-of-concept study conducted at a large Cooperative Bank.

Intel Does Not Have The Physical Capacity For State-of-the-Art Systems Of This Magnitude

- 1B CICS trans/day
- 4,000 IMS trans/sec
- 14M ACH transactions in 2.5 hours
  - 6-way sysplex
  - 30ms response
  - 216 CPU's at primary site
  - 200K MIPS

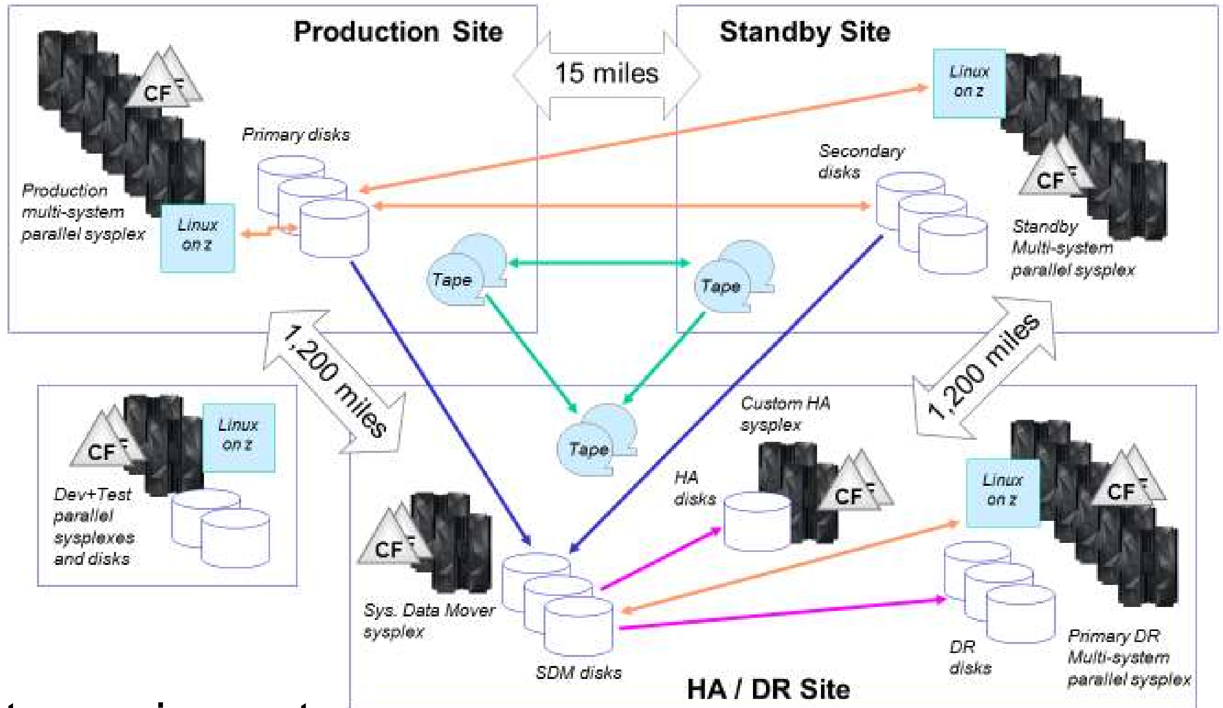

- Zero outages, zero customer impact
- Linux is Active-Active in the two data centers, with zero downtime
  - 15% Linux, growing at 30%
- *"Crazy about security overall, and the z system has a fortress around it"*

What System z Can Do That Intel-Based Systems Can't

1. Transaction processing at scale
2. Perfect Workload Management

Servers Based on latest Intel technology (Sandy Bridge)

System z

Workload Management

- Hosting platforms must be able to support high priority and low priority workloads together when sharing resources
  - Enables maximum utilization of the hosting platform

- Particularly relevant in a Private cloud environment
  - Multiple tenants with different priorities

- Desired behavior when mixing workloads
  - Low priority workloads “give up” resources to high priority workloads when required, soak up unused resources when available
  - High priority workload performance must not degrade

Priority Transactional Workload With Constant Demand Running Standalone On z/OS

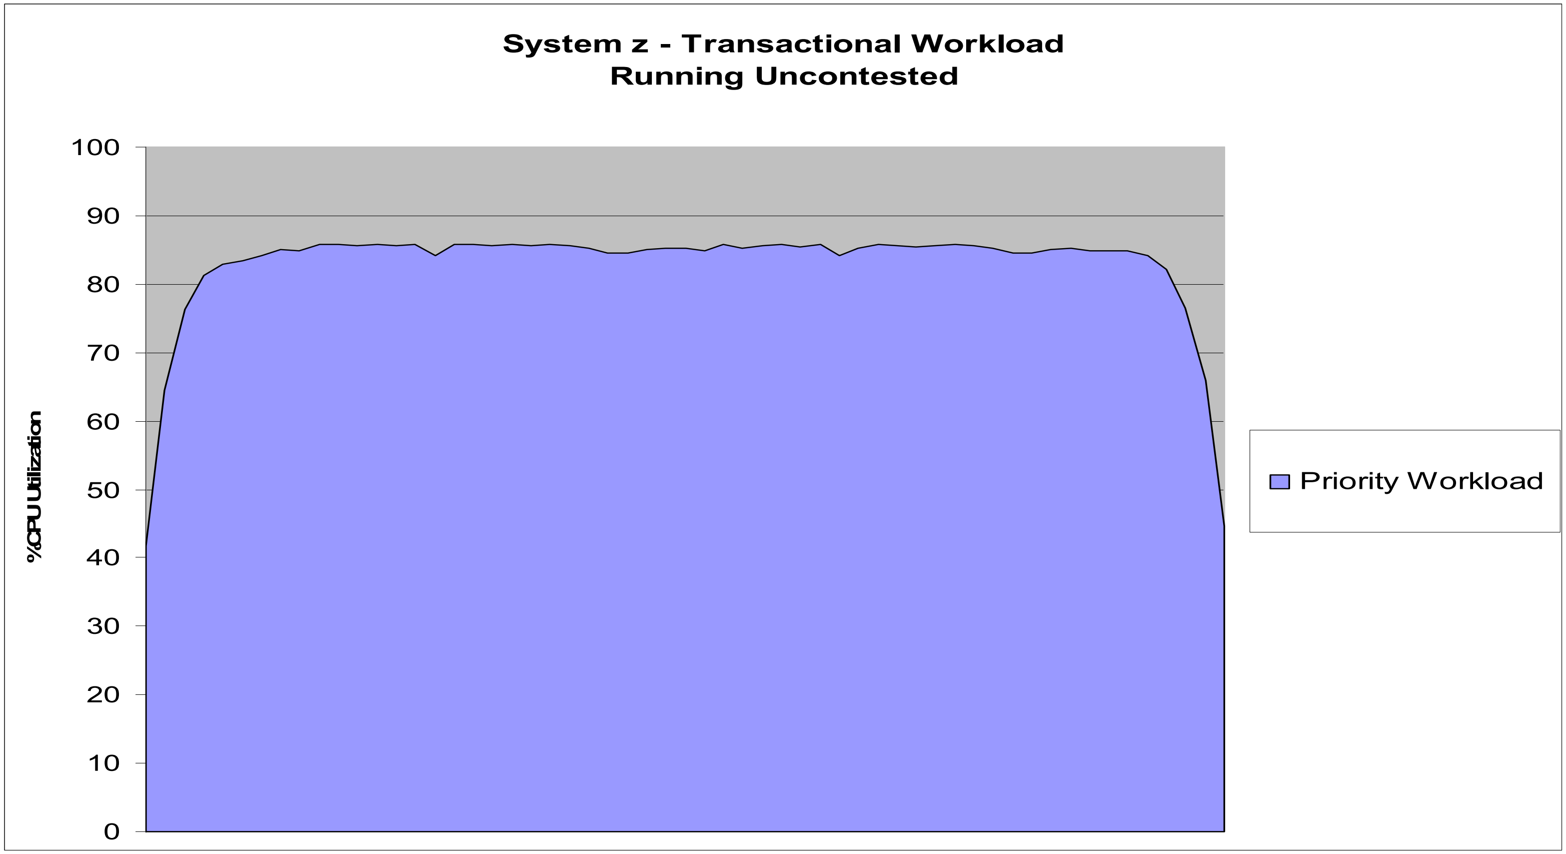

Capacity Used
High Priority Steady State - 85.2% CPU Minutes
Unused (wasted) - 14.8% CPU Minutes

Priority Workload Metrics
Total Throughput: 417.8K
Maximum TPS 129.7

Priority Transactional Workload On z/OS Does Not Degrade When Low Priority Donor Workload Is Added

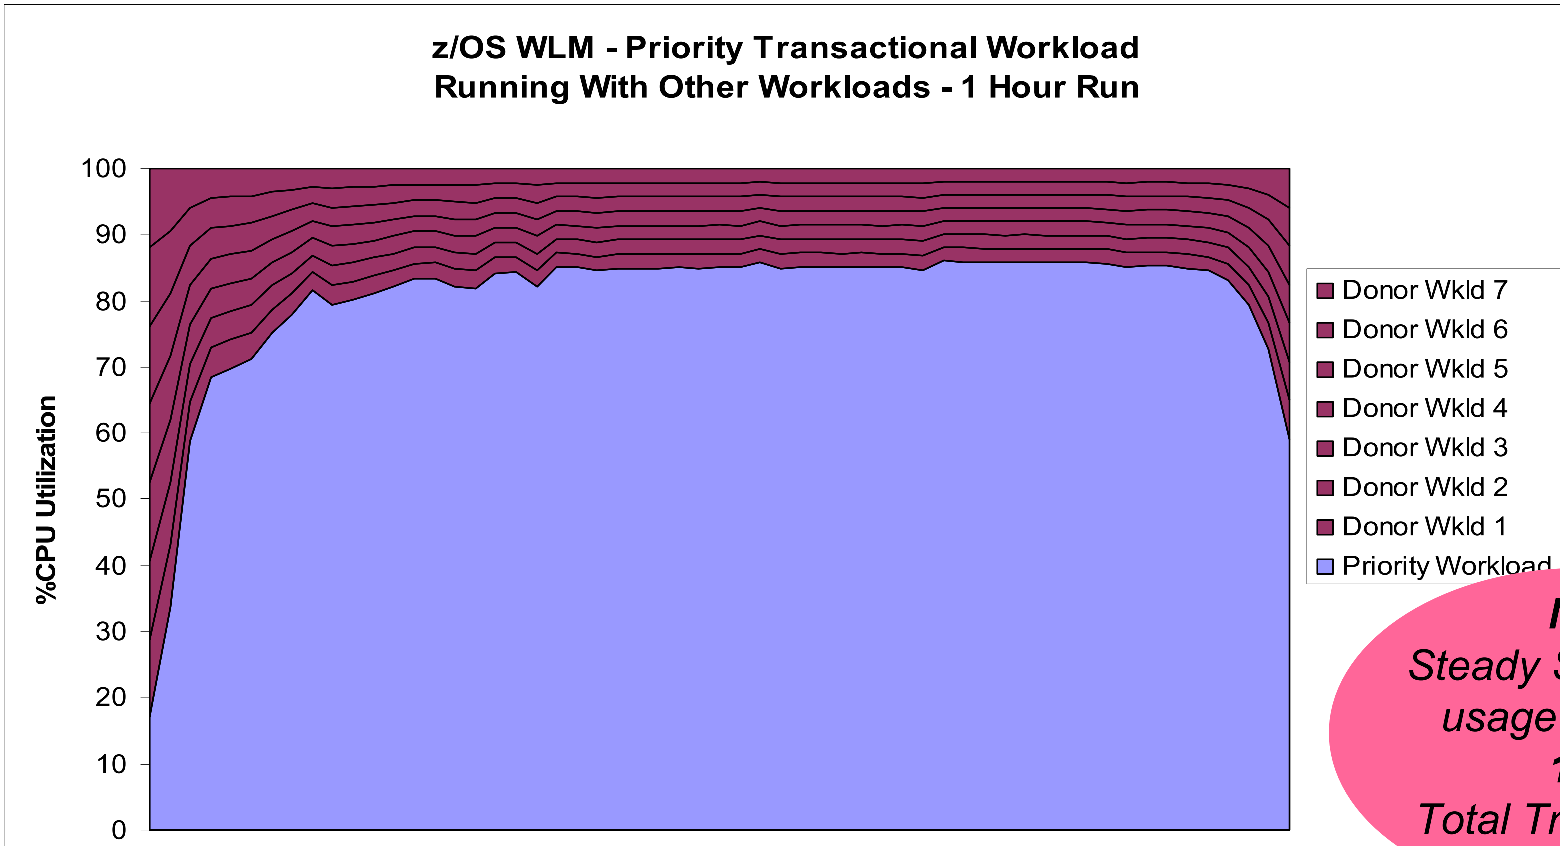

NO
Steady State CPU usage leakage
1%
Total Transaction leakage

Capacity Used
High Priority Steady State - 85.3% CPU Minutes
Unused (wasted) - 0% CPU Minutes

Priority Workload Metrics
Total Throughput: 414.7K
Maximum TPS 128.1

z/OS Workload Management Extends Priority All The Way Down To Storage

- FICON protocol supports advanced storage connectivity features not found in x86
- Priority Queuing:
  - Priority of the low-priority programs will be increased to prevent high-priority channel programs from dominating lower priority ones

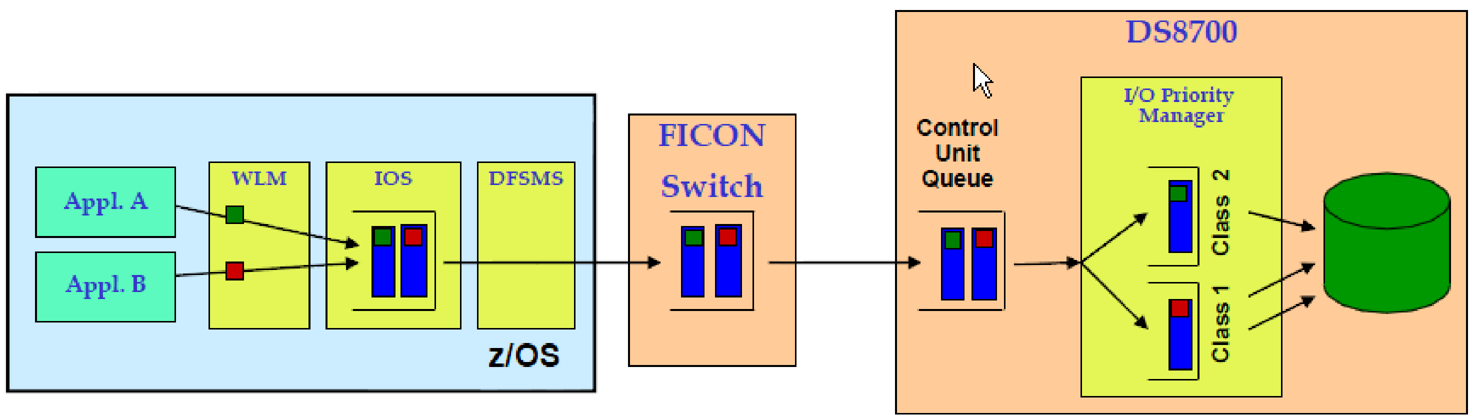

Intel can't do this

Comparison of System z To Intel Common Virtualization Environments

- High Priority web workload has defined demand over time
- SLA requires that response time does not degrade

- Low Priority web workload has unlimited demand
- It “soaks up” unused CPU minutes

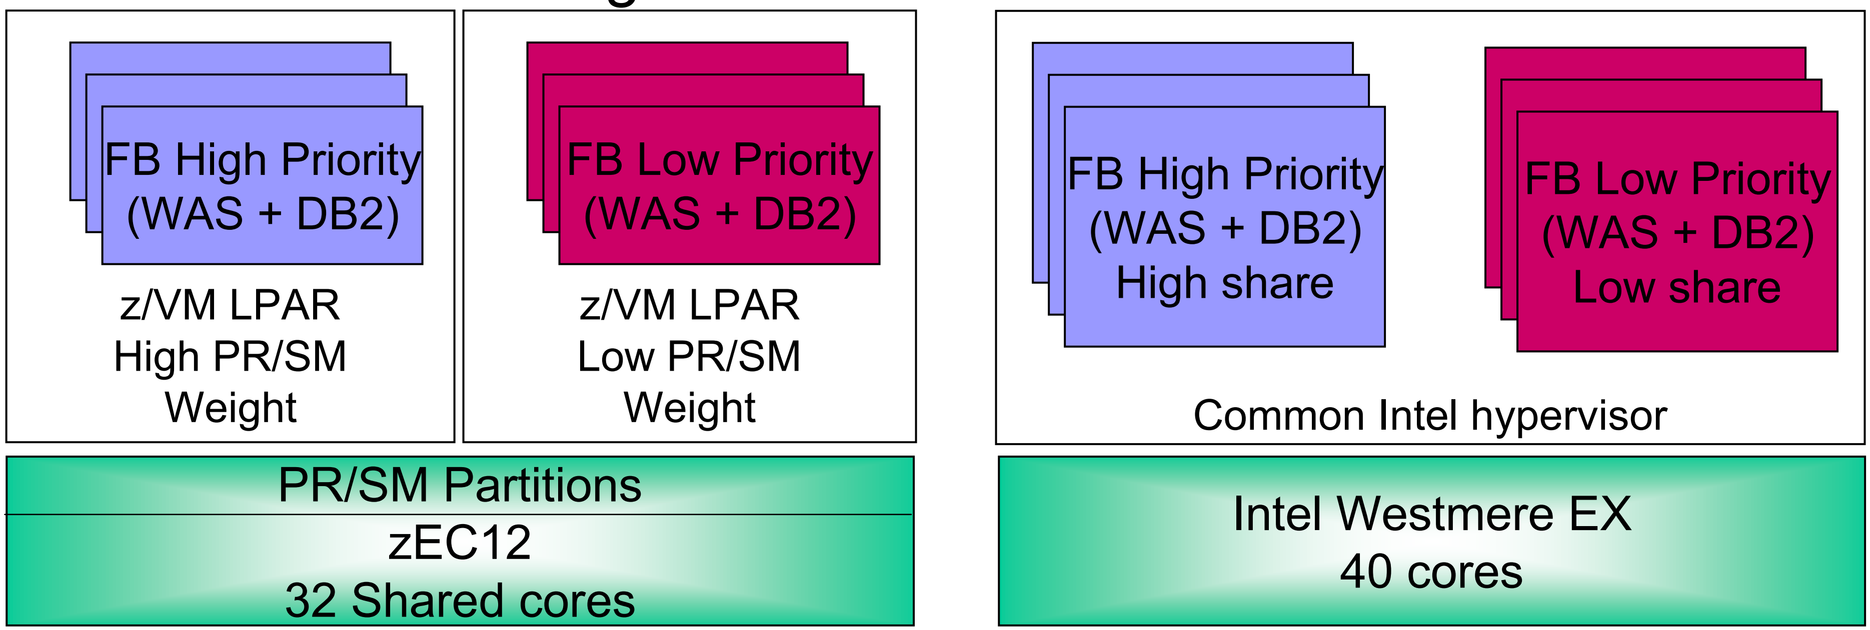

Priority Workload With Varying Demand Running Standalone On System z PR/SM

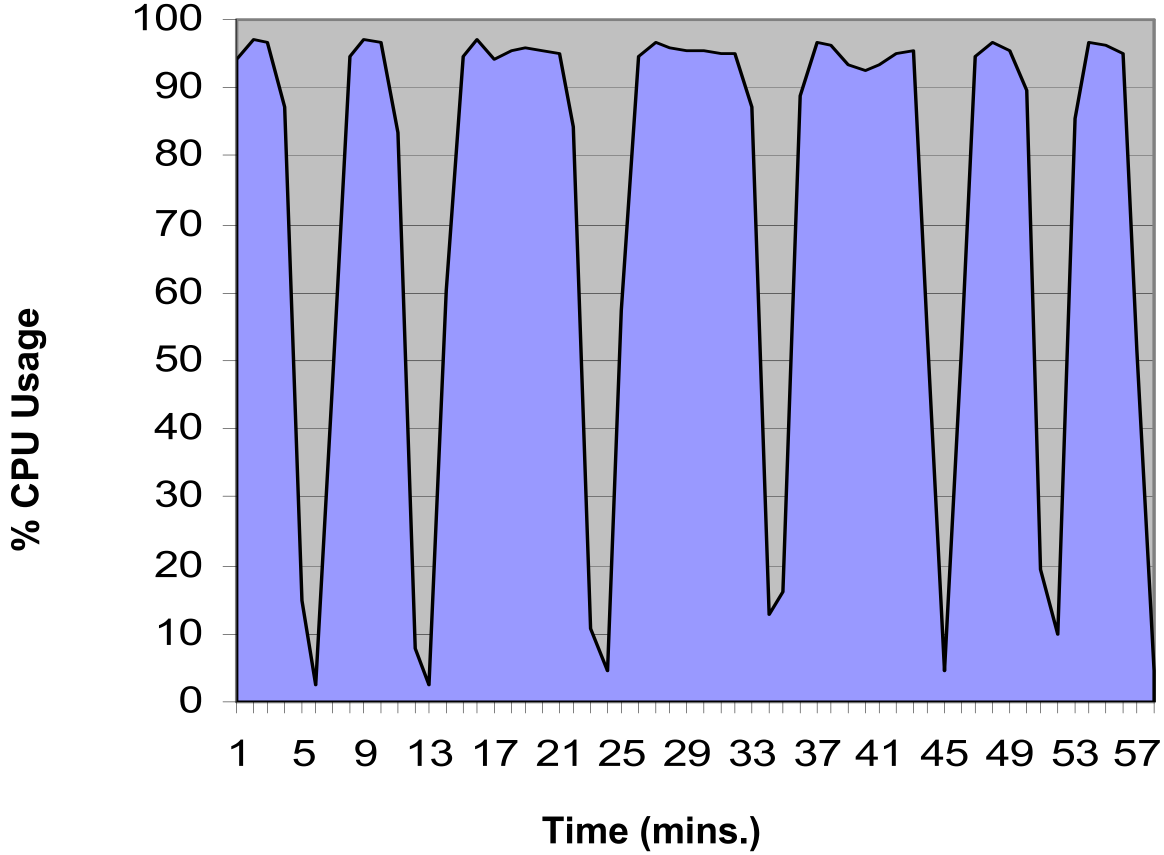

High Priority Workload Demand Curve

Priority Workload

Capacity Used
High Priority - 72.2% CPU Minutes
Unused (wasted) - 27.8% CPU Minutes

Priority Workload Metrics
Total Throughput: 9.125M
Avg Response Time: 140ms

Priority Workload On System z Does Not Degrade When Low Priority Donor Workload Is Added

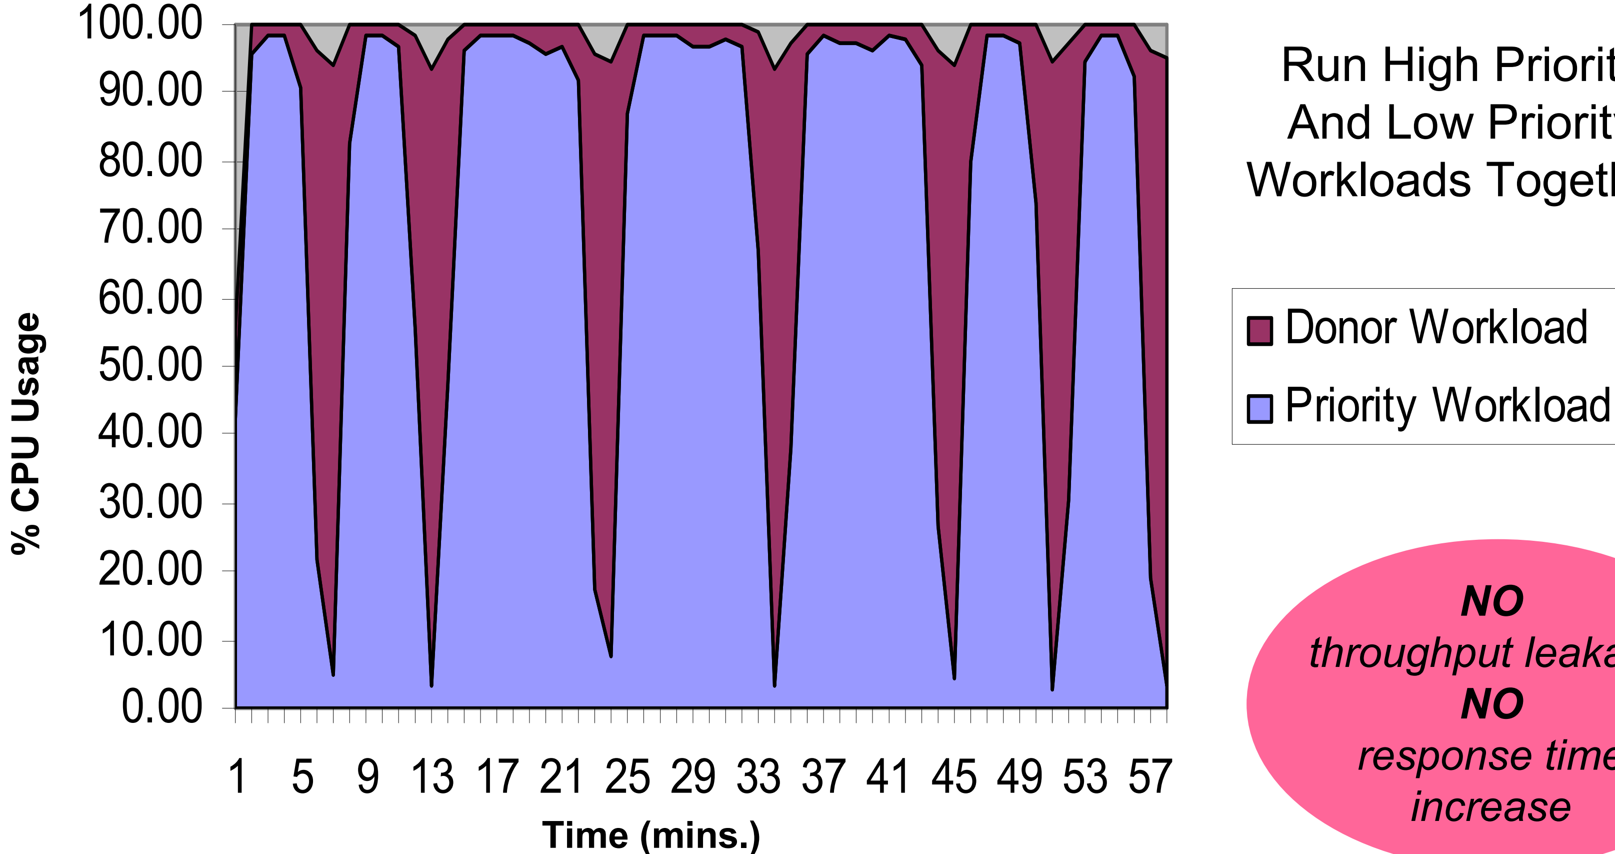

Run High Priority And Low Priority Workloads Together

***NO* throughput leakage**
***NO* response time increase**

Capacity Used
High Priority - 74.2% CPU Minutes
Low Priority - 23.9% CPU Minutes
Wasted – 1.9% CPU Minutes

Priority Workload Metrics
Total Throughput: 9.125M
Avg Response Time: 140ms

Priority Workload With Varying Demand Running Standalone On x86 Hypervisor

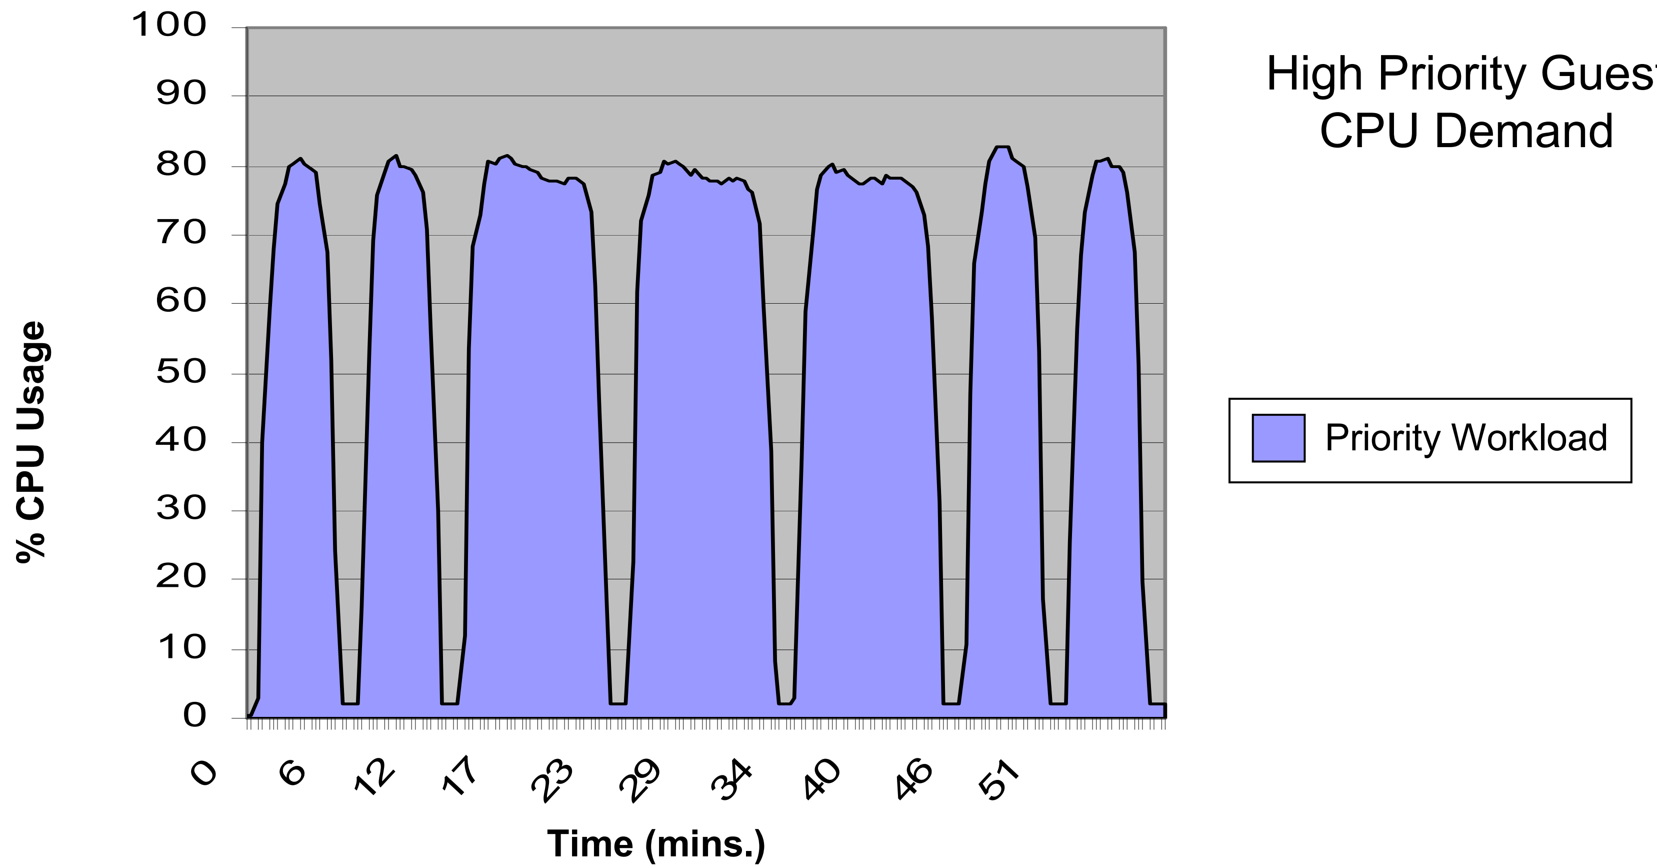

Capacity Used
High Priority - 57.5% CPU Minutes
Unused (wasted) – 42.5% CPU Minutes

Priority Workload Metrics
Total Throughput: 6.47M
Avg Response Time: 153ms

Priority Workload On x86 Hypervisor Degrades Severely When Low Priority Workload Is Added

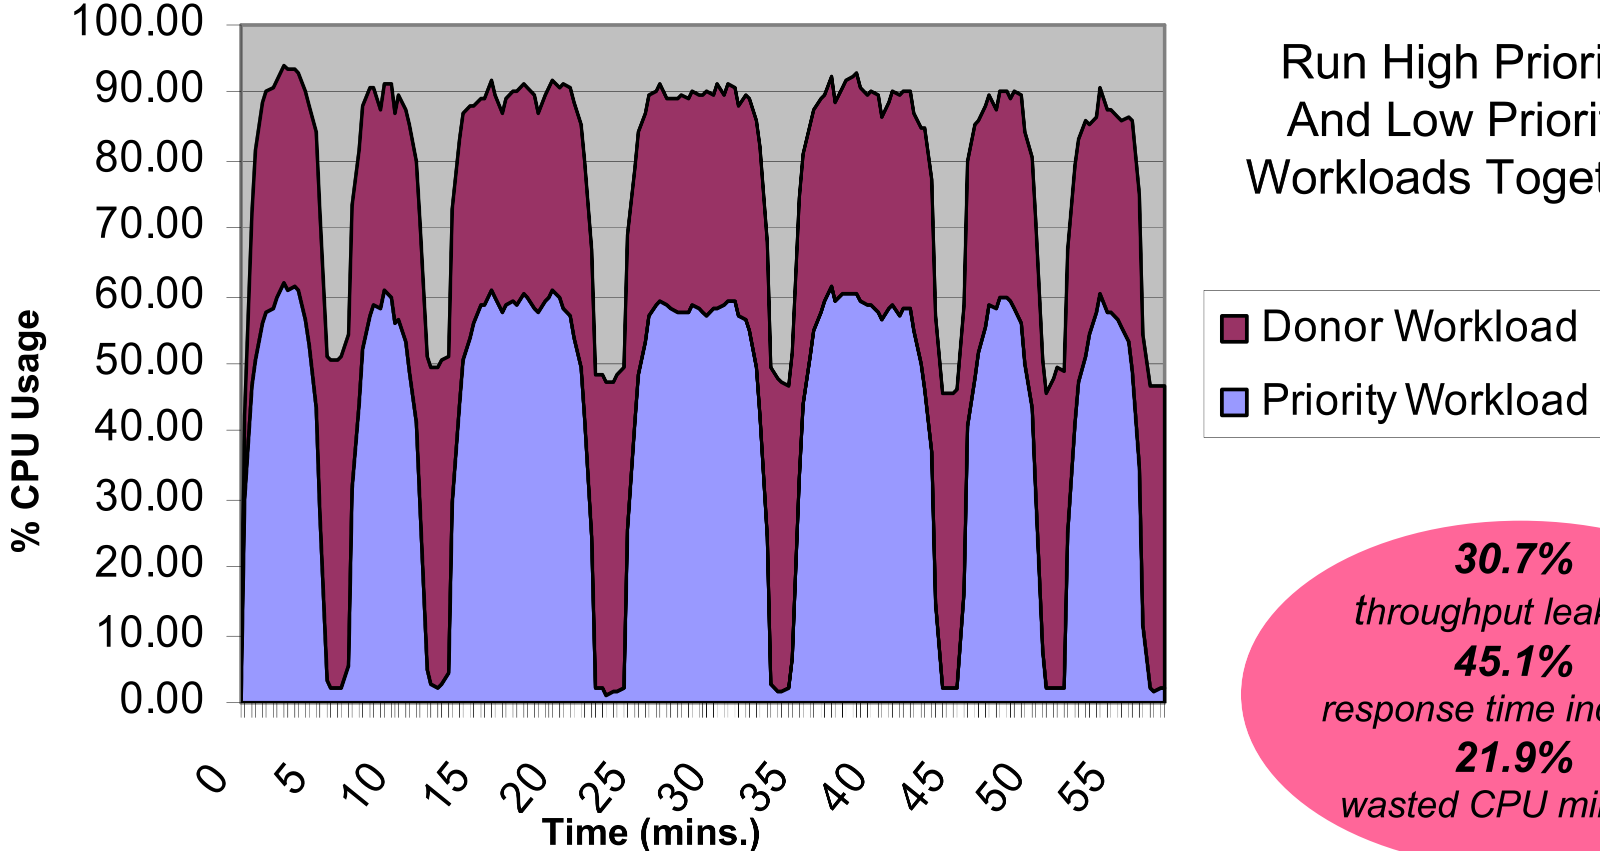

Run High Priority
And Low Priority
Workloads Together

30.7%
throughput leakage
45.1%
response time increase
21.9%
wasted CPU minutes

Capacity Used
High Priority - 42.3% CPU Minutes
Low Priority – 35.8% CPU Minutes
Wasted – 21.9% CPU Minutes

Priority Workload Metrics
Total Throughput: 4.48M
Avg Response Time: 220ms

System z Virtualization Enables Mixing Of High And Low Priority Workloads Without Penalty

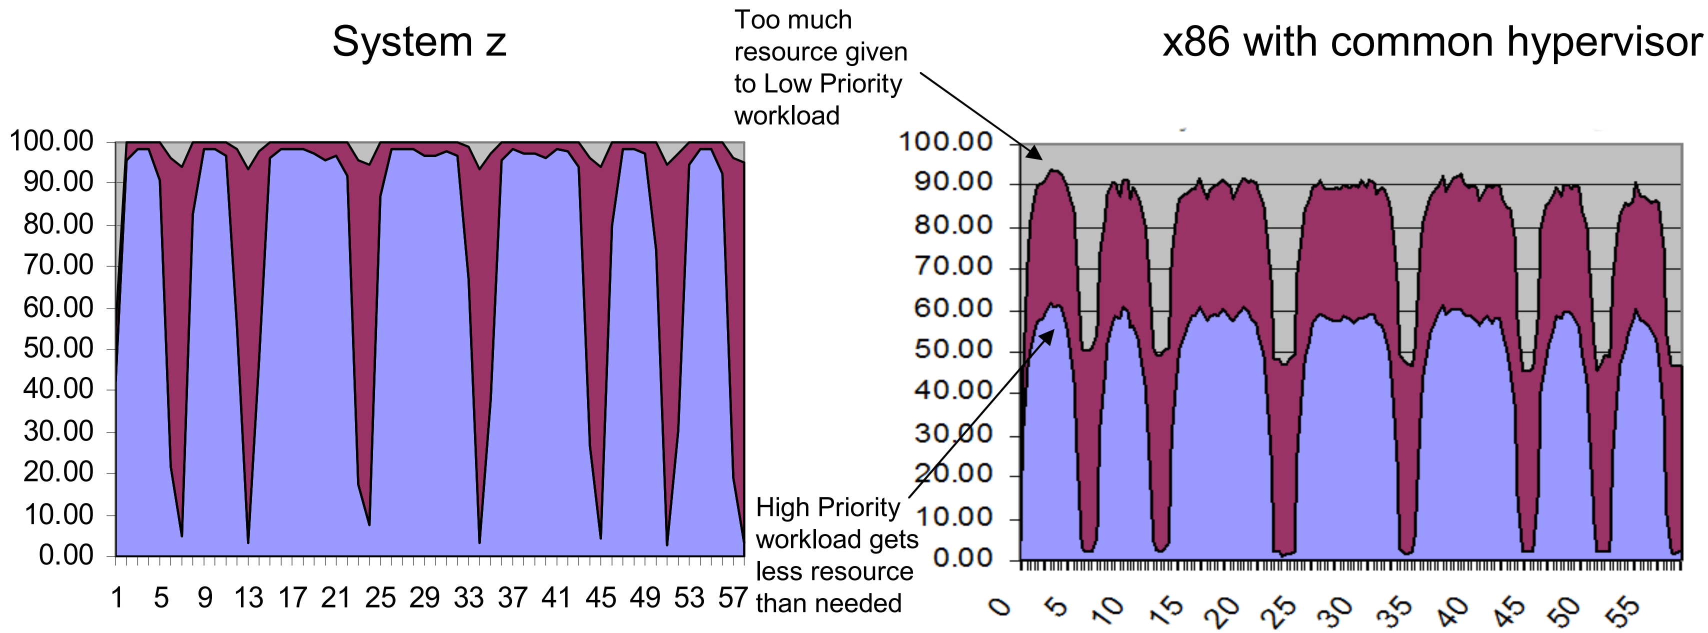

- Perfect workload management
- Consolidate workloads of different priorities on the same platform
- Full use of available processing resource (high utilization)

- Imperfect workload management
- Forces workloads to be segregated on different servers
- More servers are required (low utilization)

Inefficient Workload Management has an Economic Impact

- With Intel virtualization, the only practical solution to maintain SLA of high priority workload is to deploy workloads into separate environments
  - Most Intel virtualized deployments separate Dev environments from Production environments

- The need to maintain multiple environments directly affects total cost
  - Fixed size Intel boxes can force additional boxes to accommodate “spill over” high priority work
    - Spare capacity on additional machine is wasted as nothing else can run on it without impacting primary workload SLA
  - Additional environment needed to deliver lower priority workloads

Deliver High And Low Priority Workloads Together While Maintaining Response Time SLA

Comparison to determine which platform provides the lowest TCA over 3 years

High priority workloads

Low priority workloads

- **IBM WebSphere 8.5 ND**
- **IBM DB2 10 AESE**
- **Monitoring software**

High priority online banking workloads driving a total of **11.89M** transactions per hour and low priority discretionary workloads

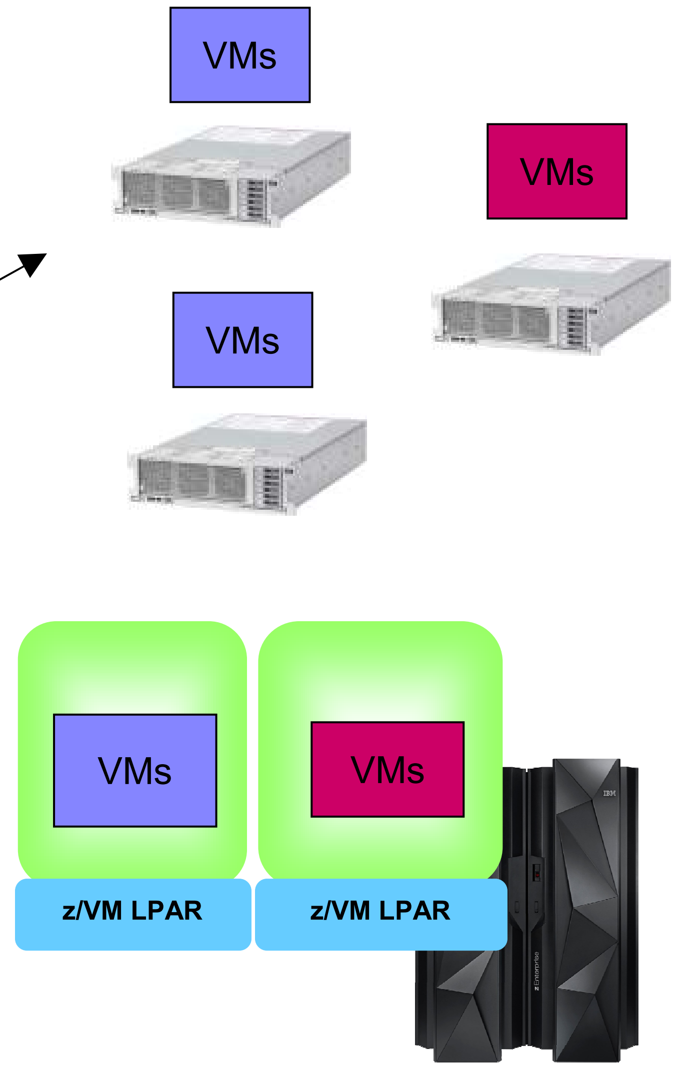

Virtualized on 3 Intel 40 core servers

$13.66M (3 yr. TCA)

z/VM on zEC12
32 IFLs

$5.77M (3 yr. TCA)

58% lower cost!

Consolidation ratios derived from IBM internal studies. Results may vary based on customer workload profiles/characteristics. Prices will vary by country.

Workload Management Summary

- Additional environments needed for Intel can have a cost impact beyond server (hw & sw) acquisition
  - Supporting infrastructure components contribute to additional cost (TCO)
    - Storage, data copies, network, labor, environmental
- Imperfect workload management is one of the factors that leads to core proliferation in an offload scenario
  - Different environments that co-existed on z end up taking dedicated resources, driving up the total number of cores needed
- The perfect workload management capability of System z makes it an ideal Private Cloud platform
  - Support multiple tenants with different priorities
  - Ability to maintain priority enables maximum utilization of the platform, driving down cost per workload

What System z Can Do That Intel-Based Systems Can't

1. Transaction processing at scale
2. Perfect Workload Management
3. Greater Core Density

Servers Based on latest Intel technology (Sandy Bridge)

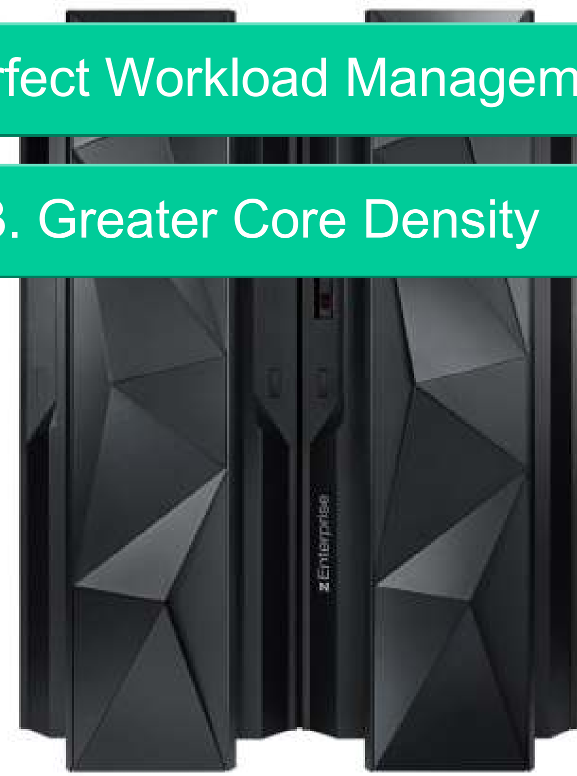

System z

Why Core Proliferation Happens When Moving Workload From System z To Intel

- De-consolidation of applications to dedicated servers – decomposing highly tuned co-located components
- Processing expansion requirements for CICS/COBOL applications
- 3x expansion when converting hierarchical databases to relational
- Functional segregation into production, development and test
- 100% hardware coverage for Disaster Recovery costs double

System z

Core Proliferation For A Large Workload

16x 32-way HP Superdome
App. Production / Dev / Test

8x 48-way HP Superdome
DB Production / Dev /Test

zEC12 41-way Production / Dev / Test

***41** GP processors*
(38,270 MIPS)

***896** processors*
(3,668,600 Perf Units)

***22x** more cores!*

$180M (5 yr. TCA)

$111M (5 yr. TCA)

NOTE: To cover DEV/QA capacity, add 100% servers for distributed servers, add 25% MIPS (8,000) to System z

Core Proliferation For A Mid-sized Workload

6x 8-way Production / Dev
2x 64-way Production / Dev
Application/MQ/DB2/Dev partitions

2x z900 3-way Production / Dev / QA / Test

6 processors
(1,660 MIPS)

176 processors
(800,072 Performance units)

$25.4M (5 yr. TCO)

$17.9M (5 yr. TCO)

***29x** more cores!*

482 Performance Units per MIPS

Core Proliferation For a Small Workload

3x HP DL580 (2ch/20co)
Production / Dev / Test
(2011 x86 technology)

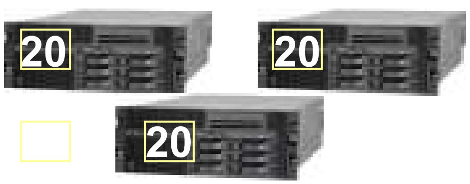

z800 Production /
Dev / Test
(2002 mainframe technology)

2.1 processors
(499 MIPS)

60 Linux processors
(383,022 Perf Units)

29x more cores

(despite the 9 year technology gap!)

1.5 Year Migration

768 Performance Units per MIPS

Core Proliferation For Oracle Workloads

TCO study for a Media and Entertainment Industry customer

107 HP servers

1440 cores total

***30x** more cores!*

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

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 (5 yr. TCO)	**$20.5M**

Intel: Oracle DB + App costs = $13.1M (LIC + maint over 5 yrs.).
IBM: Oracle DB + App costs = $1.92M (LIC + maint over 5 yrs.)

What System z Can Do That Intel-Based Systems Can't

1. Transaction processing at scale
2. Perfect Workload Management
3. Greater Core Density
4. Spare Capacity for Growth

Servers Based on latest Intel technology (Sandy Bridge)

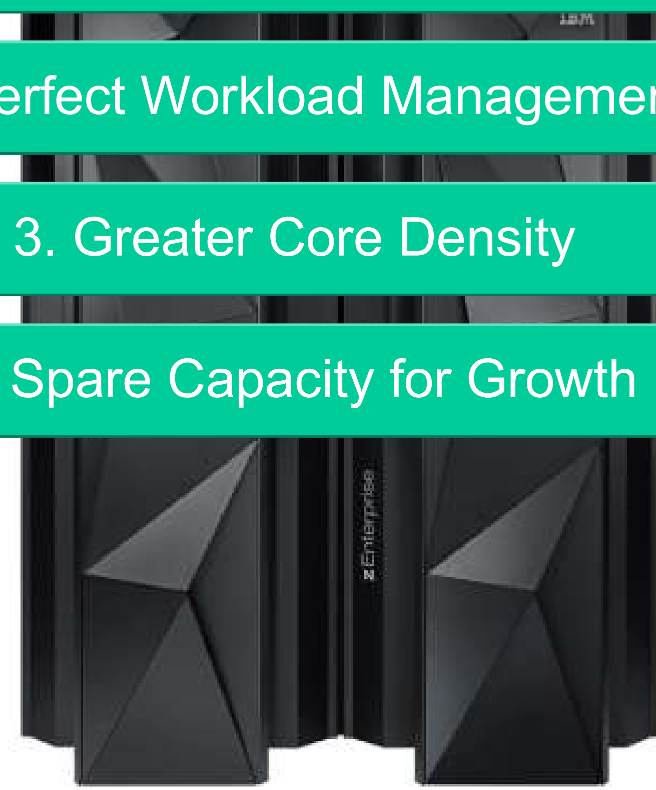

System z

System z Capacity On Demand Provides Elasticity To Handle Unexpected Peaks

- Capacity on Demand
  - “Books” are shipped fully populated
  - Activate dormant processors as needed
  - Use for temporary or permanent capacity
  - Self-managed on/off
- New capacity is immediately available for work without service disruption

One Book with 36 Processors

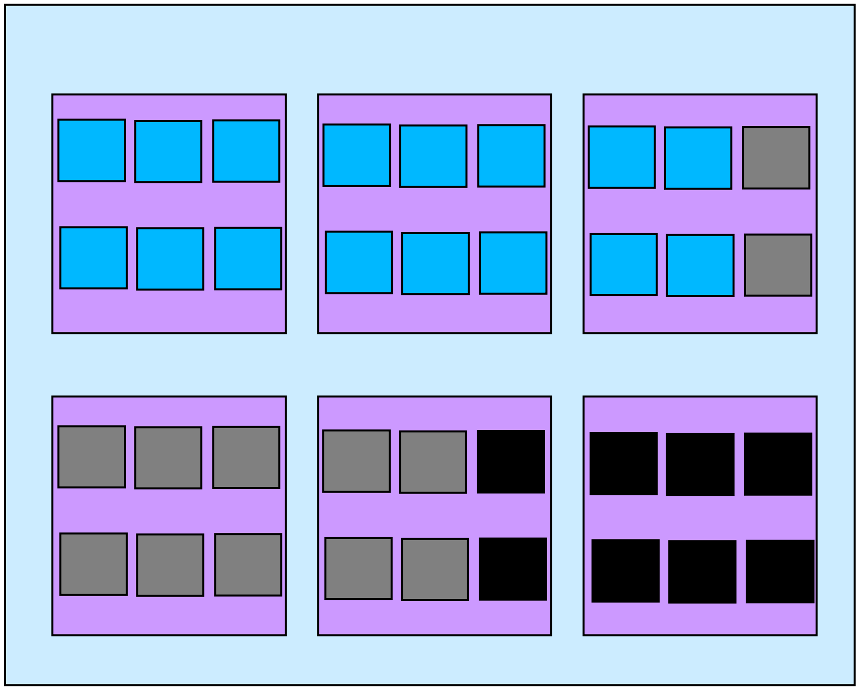

Active processors – pay full price

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

Dark processors (unused) – no charge

What System z Can Do That Intel-Based Systems Can't

Servers Based on latest Intel technology (Sandy Bridge)

1. Transaction processing at scale
2. Perfect Workload Management
3. Greater Core Density
4. Spare Capacity for Growth
5. Comprehensive Disaster Recovery

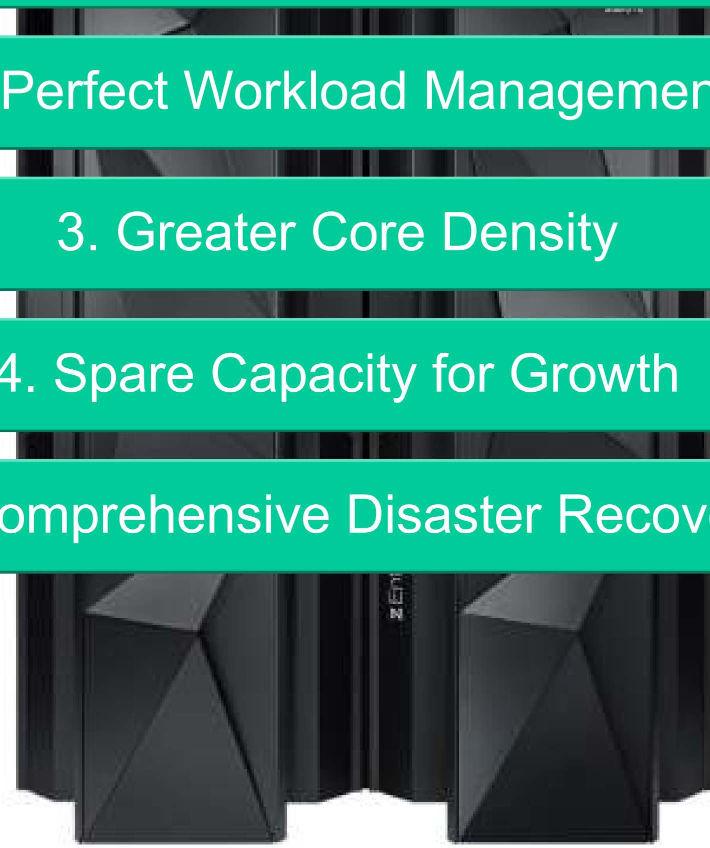

System z

System z Disaster Recovery Is Systematic And Comprehensive

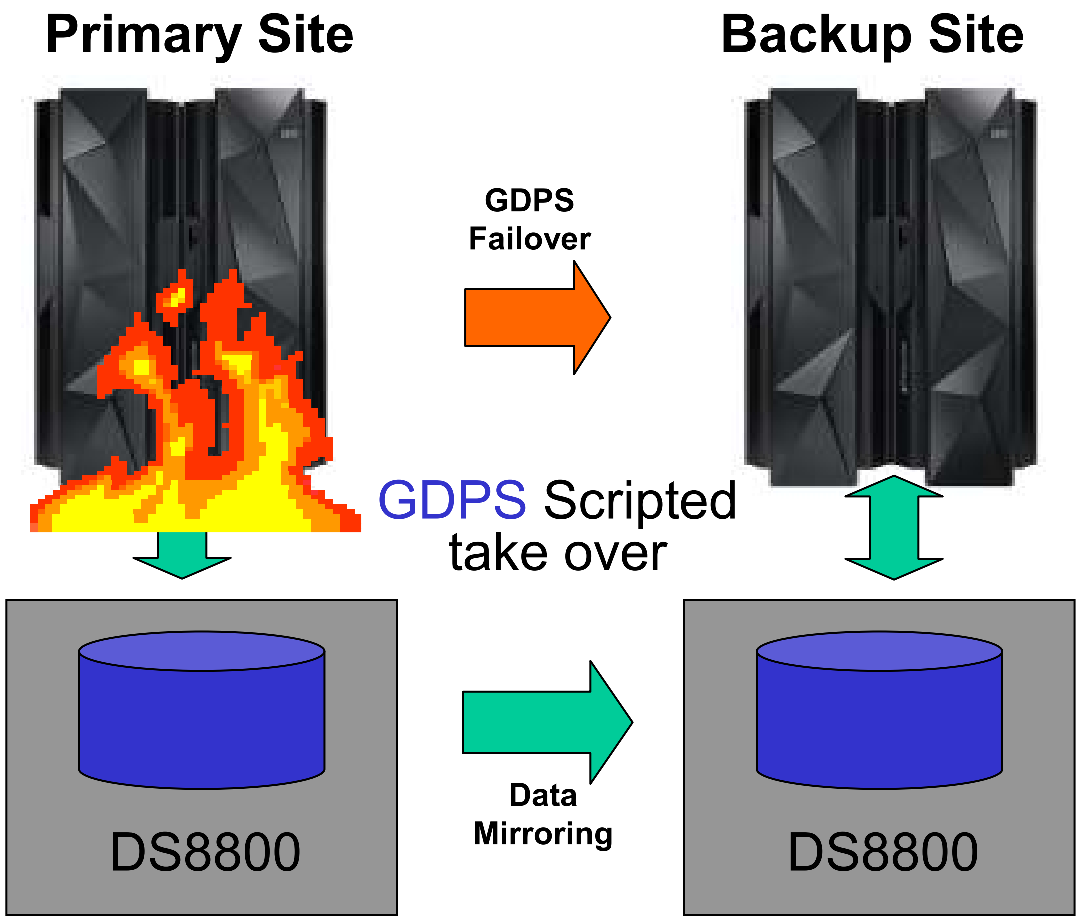

- Site Failover
  - Failover to secondary site in case of complete site failure
- Data Mirroring
  - Protect data in the event of a disk system failure

Supports systematic Disaster Recovery for virtualized Linux environments also

Complexity Of Intel Disaster Recovery Solutions Prohibits Wide Spread Use

- Workloads on standalone Intel servers require a disaster recovery solution for each server
  - Data mirroring
  - Failover and restart
- Embedded storage is difficult to mirror
- Comprehensive workload failover is not feasible for hundreds of servers

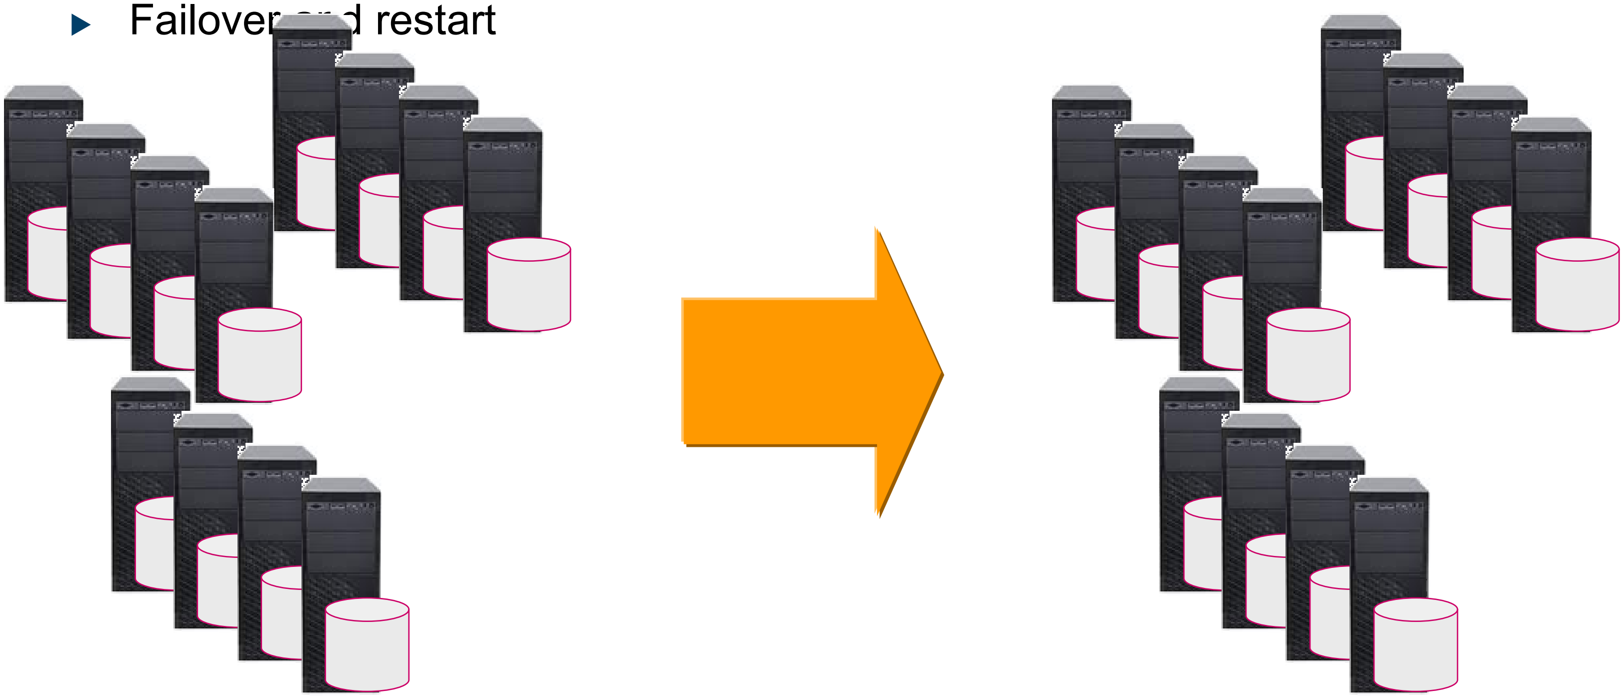

Consolidation Of Workloads On System z Simplifies Disaster Recovery

- Workloads are consolidated onto z/VM partitions as Linux guests
- Linux on System z can be failed over as part of GDPS

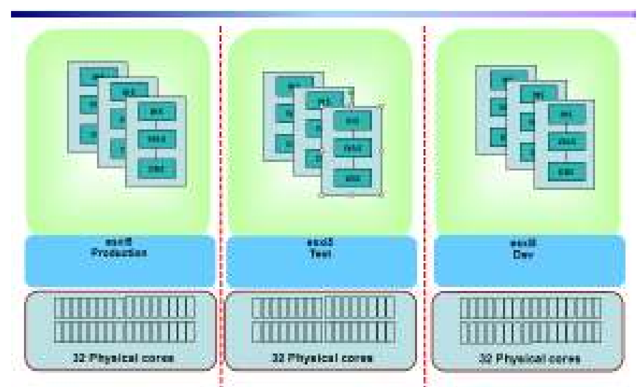

What System z Can Do That Intel-Based Systems Can't

1. Transaction processing at scale
2. Perfect Workload Management
3. Greater Core Density
4. Spare Capacity for Growth
5. Comprehensive Disaster Recovery
6. Runs Longer without Stopping

System z

Servers Based on latest Intel technology (Sandy Bridge)

System z Has More Comprehensive Protection To Ensure Better Availability Than Intel

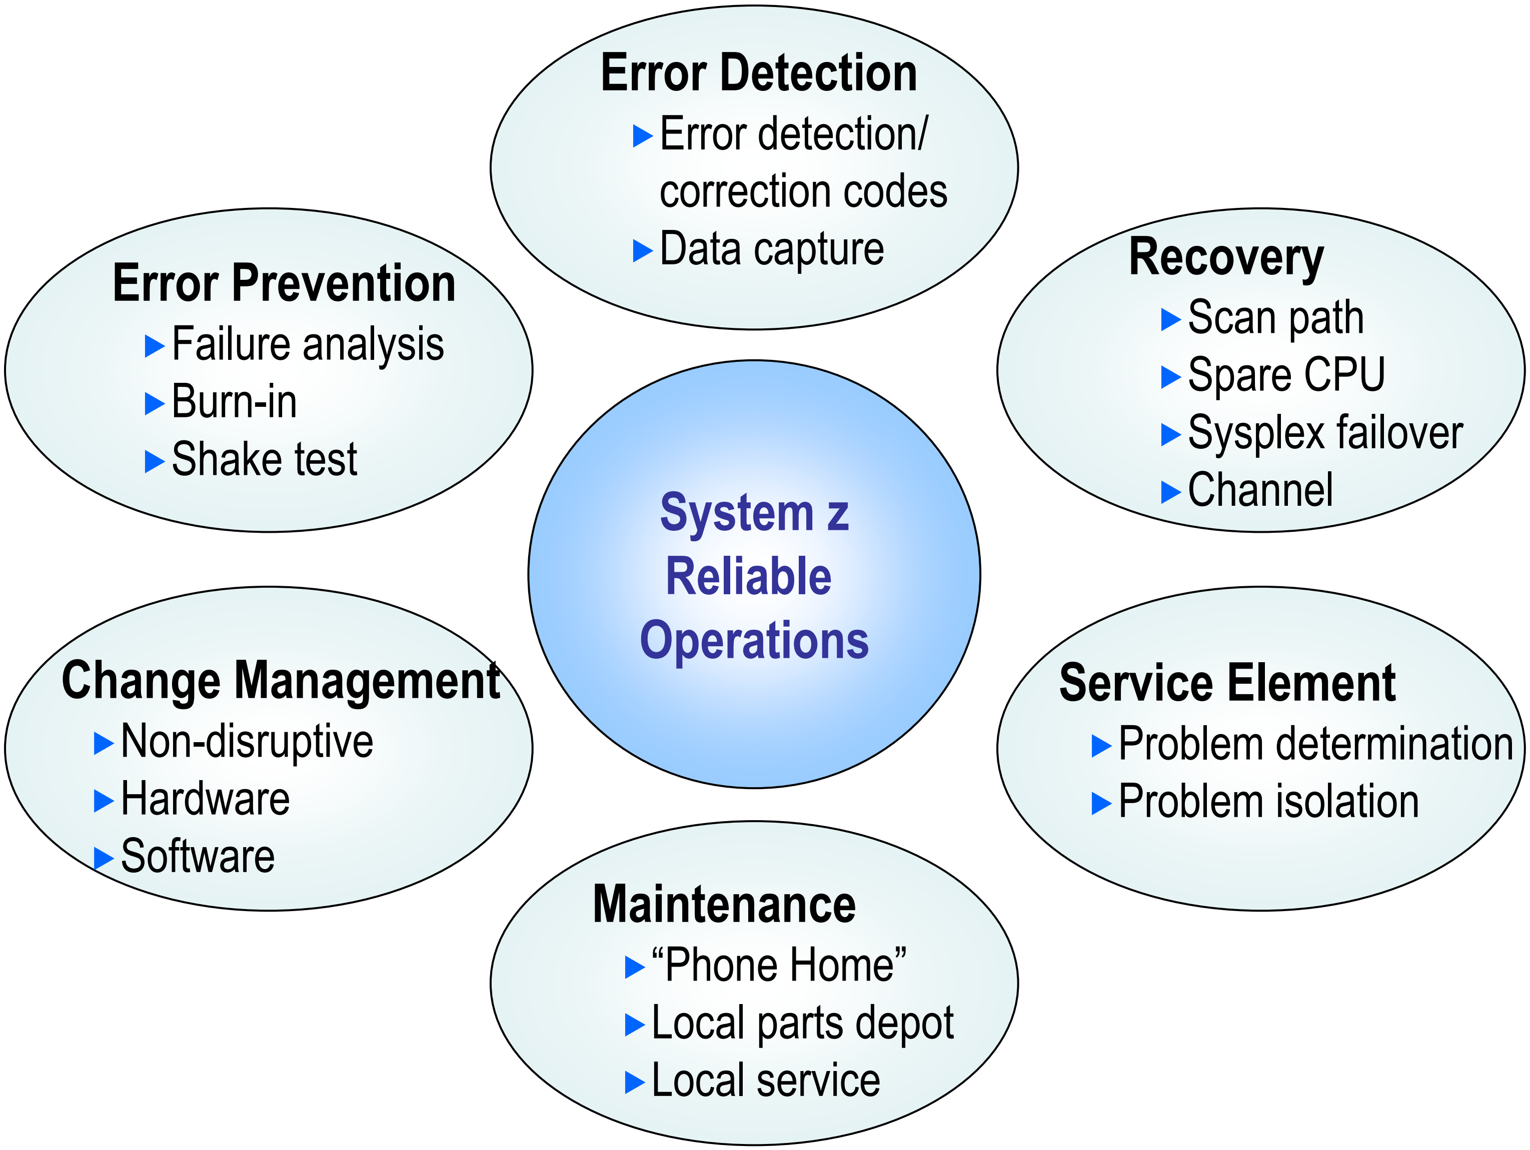

Example: CPU Sparing

- zEC12 has 2 spare CPUs per server
- System controllers can detect a failing processor chip
- Status of the unit of workload running on the failing CPU can be saved
- Failing CPU can be switched with the spare with NO interruption with the workload
- Alternatively, spare processors can be enabled at certain times during unexpected peak workloads
  - Another aspect of Capacity on Demand (COD)

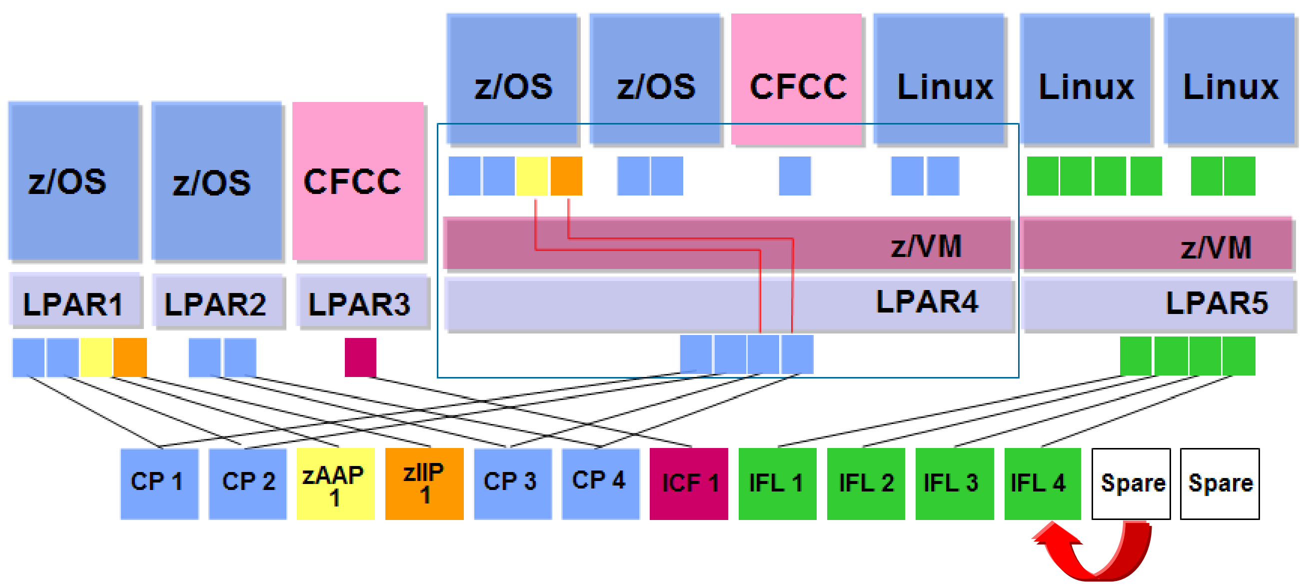

Example: I/O Channel Failover

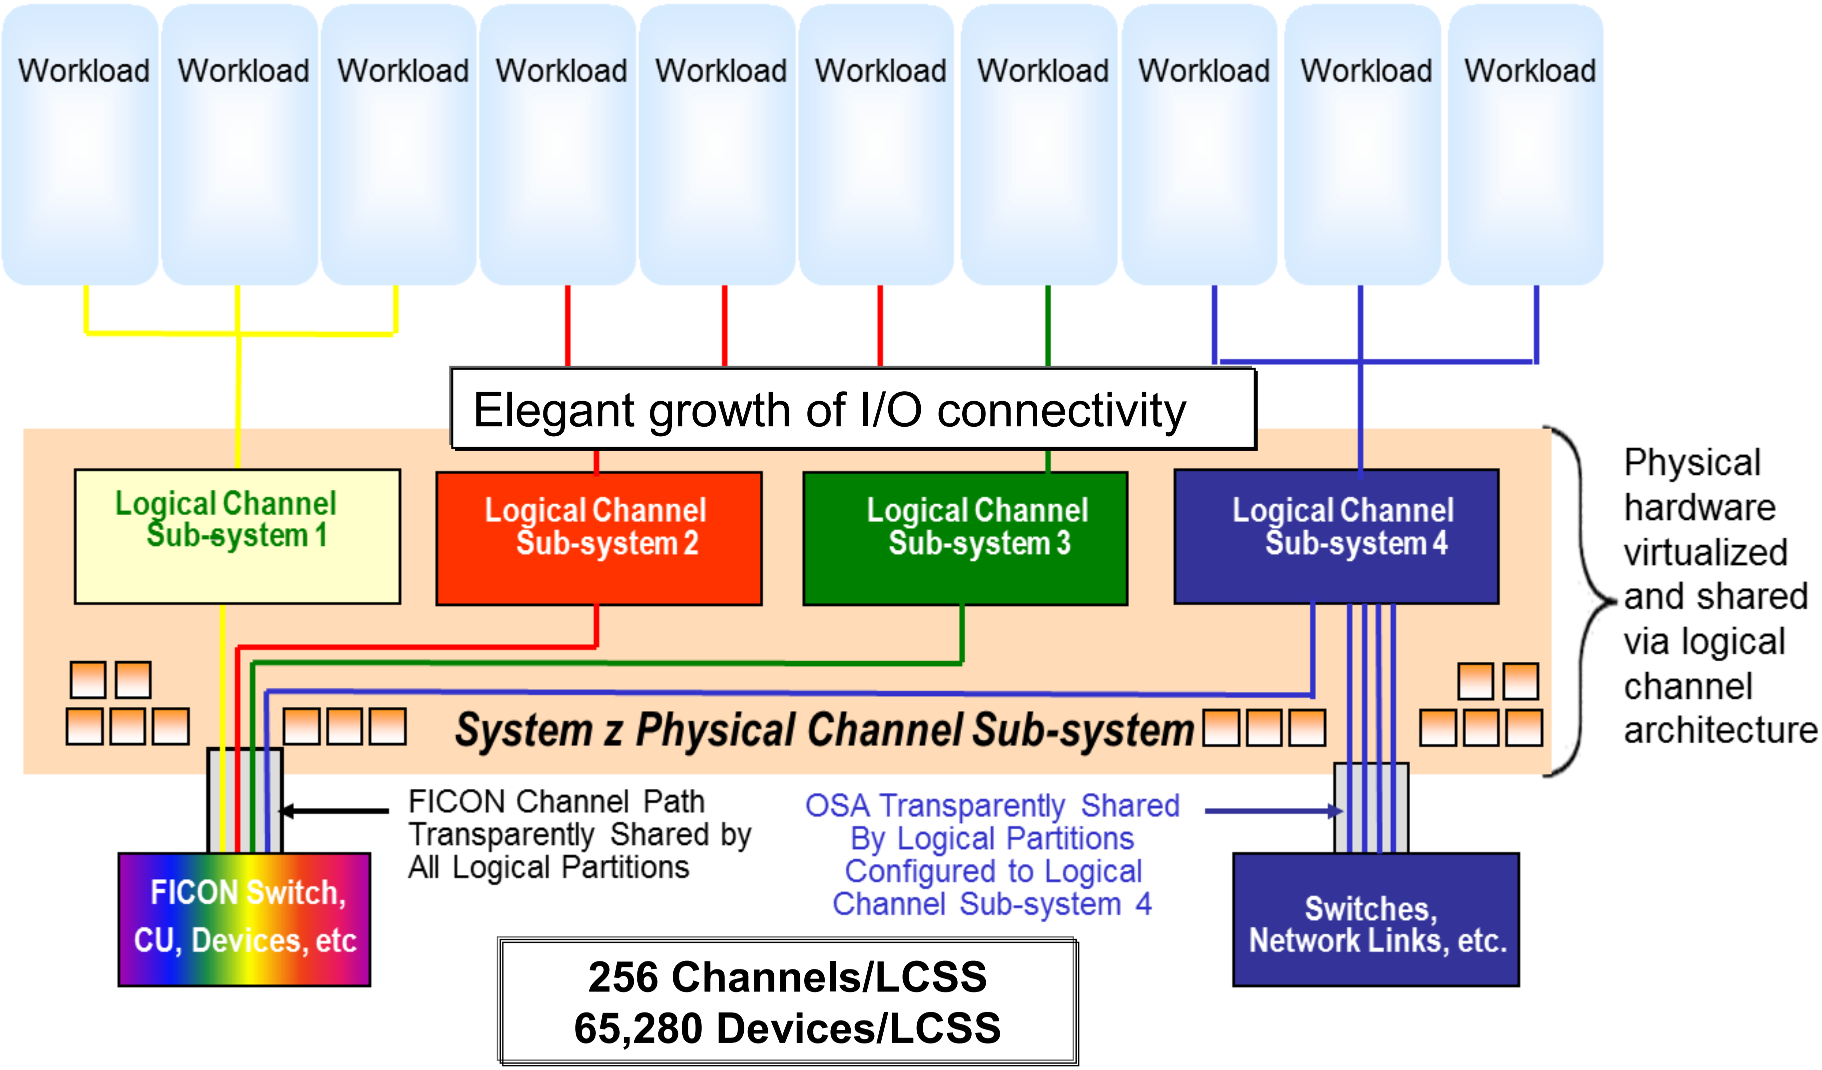

System z Supports Concurrent Operations During Hardware Repair – *Intel Can't*

Capability	zEC12	x86
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

Single book systems may not support concurrent memory upgrades

What System z Can Do That Intel-Based Systems Can’t

Servers Based on latest Intel technology (Sandy Bridge)

1. Run Bigger and More Workloads
2. Perfect Workload Management
3. Greater Core Density
4. Spare Capacity for Growth
5. Comprehensive Disaster Recovery
6. Runs Longer without Stopping
7. Tougher Security

System z

Enterprise Security Requires Many Elements: Infrastructure

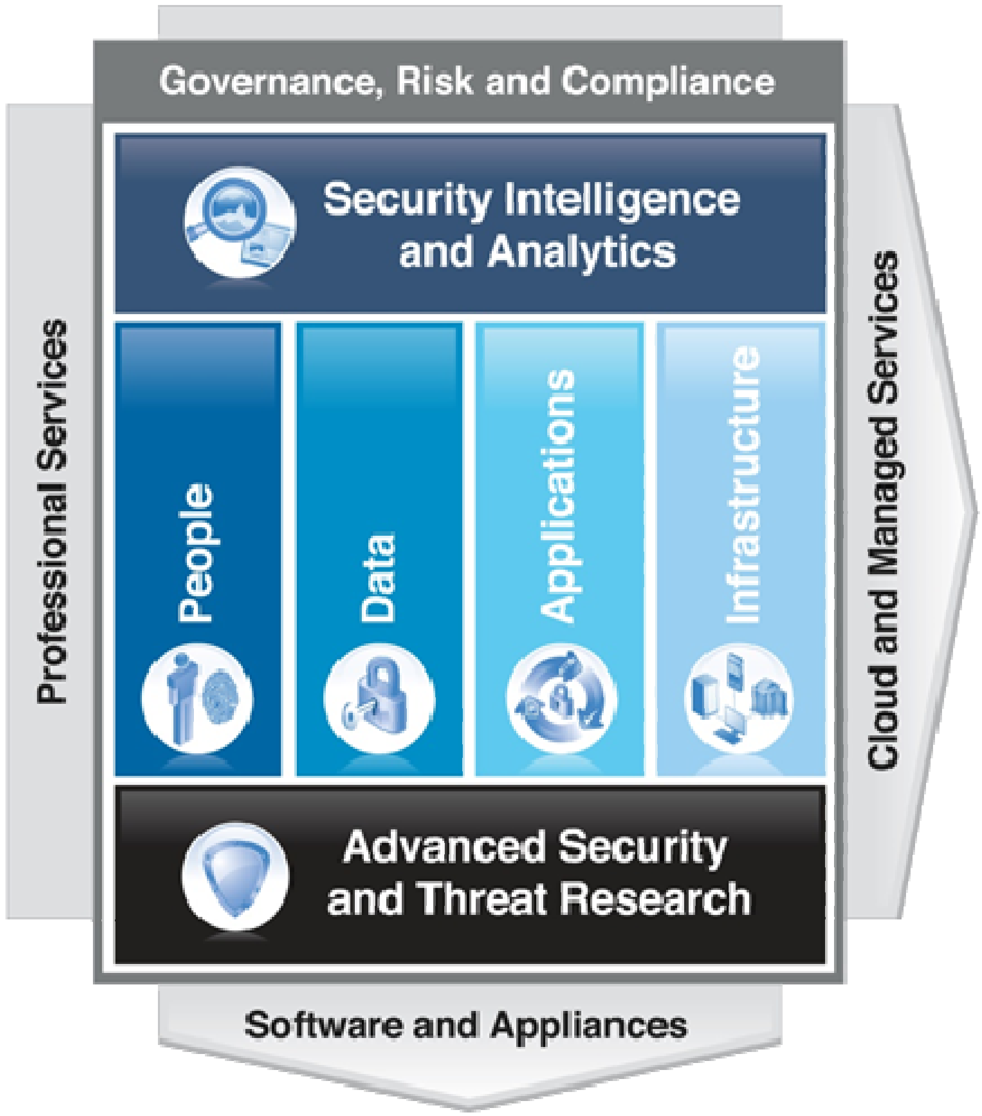

Common Criteria Certifications Show System z Platform Security Leads the Industry

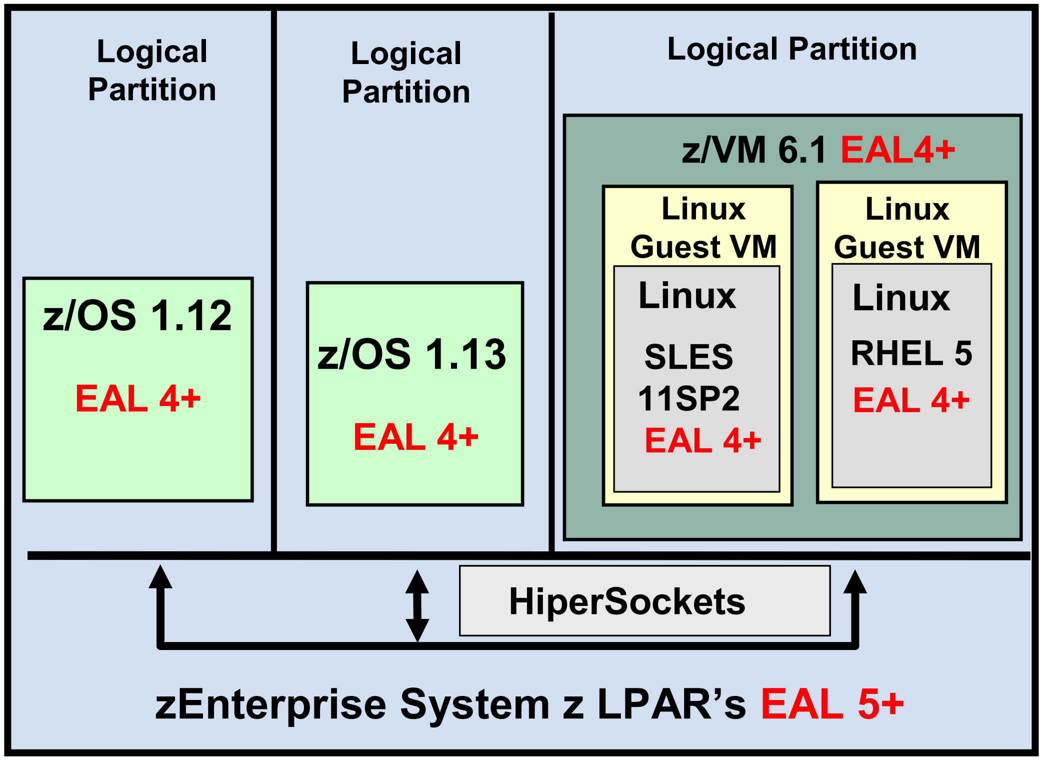

System z

- What is Common Criteria?
  - Common Criteria is an accepted standard for evaluating the inherent security of a computing system
  - Common Criteria is based on a set of functional and assurance requirements
  - A higher Enterprise Assurance Level (**EAL)** rating is more secure
  - The security requirements in Common Criteria have gained support as "best practices"
  - **IBM System z holds the highest EAL grades in Common Criteria!**

EAL Ratings: https://www.bsi.bund.de/SharedDocs/Zertifikate/CC/Betriebssysteme/0852.html

DB2 Maturity Delivers A Proven Track Record For Data Security

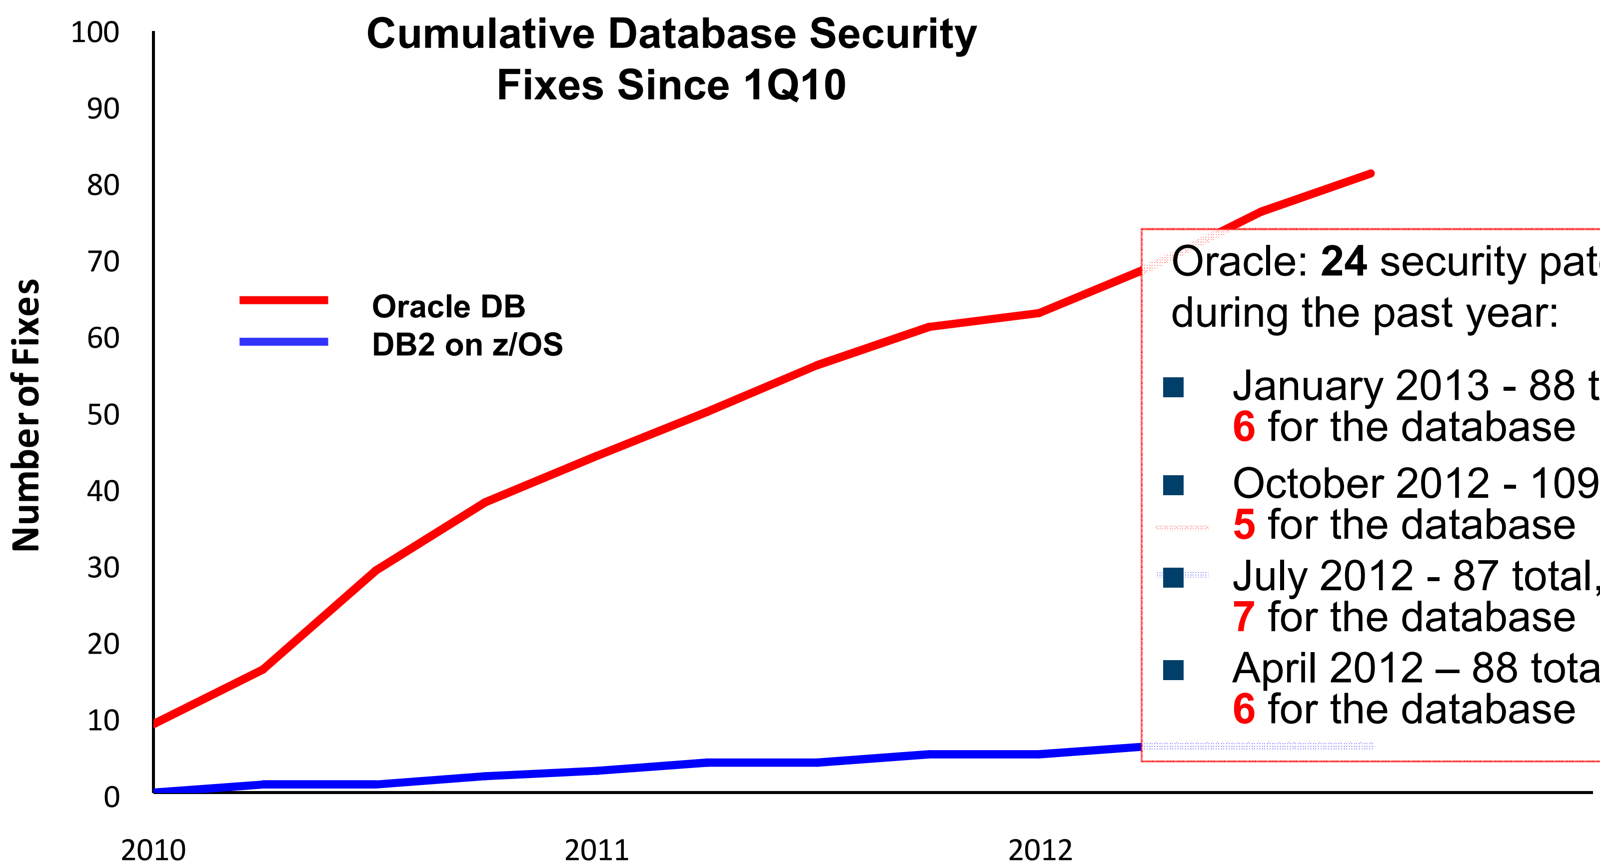

Oracle: **24** security patches during the past year:

- January 2013 - 88 total, **6** for the database
- October 2012 - 109 total, **5** for the database
- July 2012 - 87 total, **7** for the database
- April 2012 – 88 total, **6** for the database

Source: http://www.oracle.com/technetwork/topics/security

Elements Of System z Security

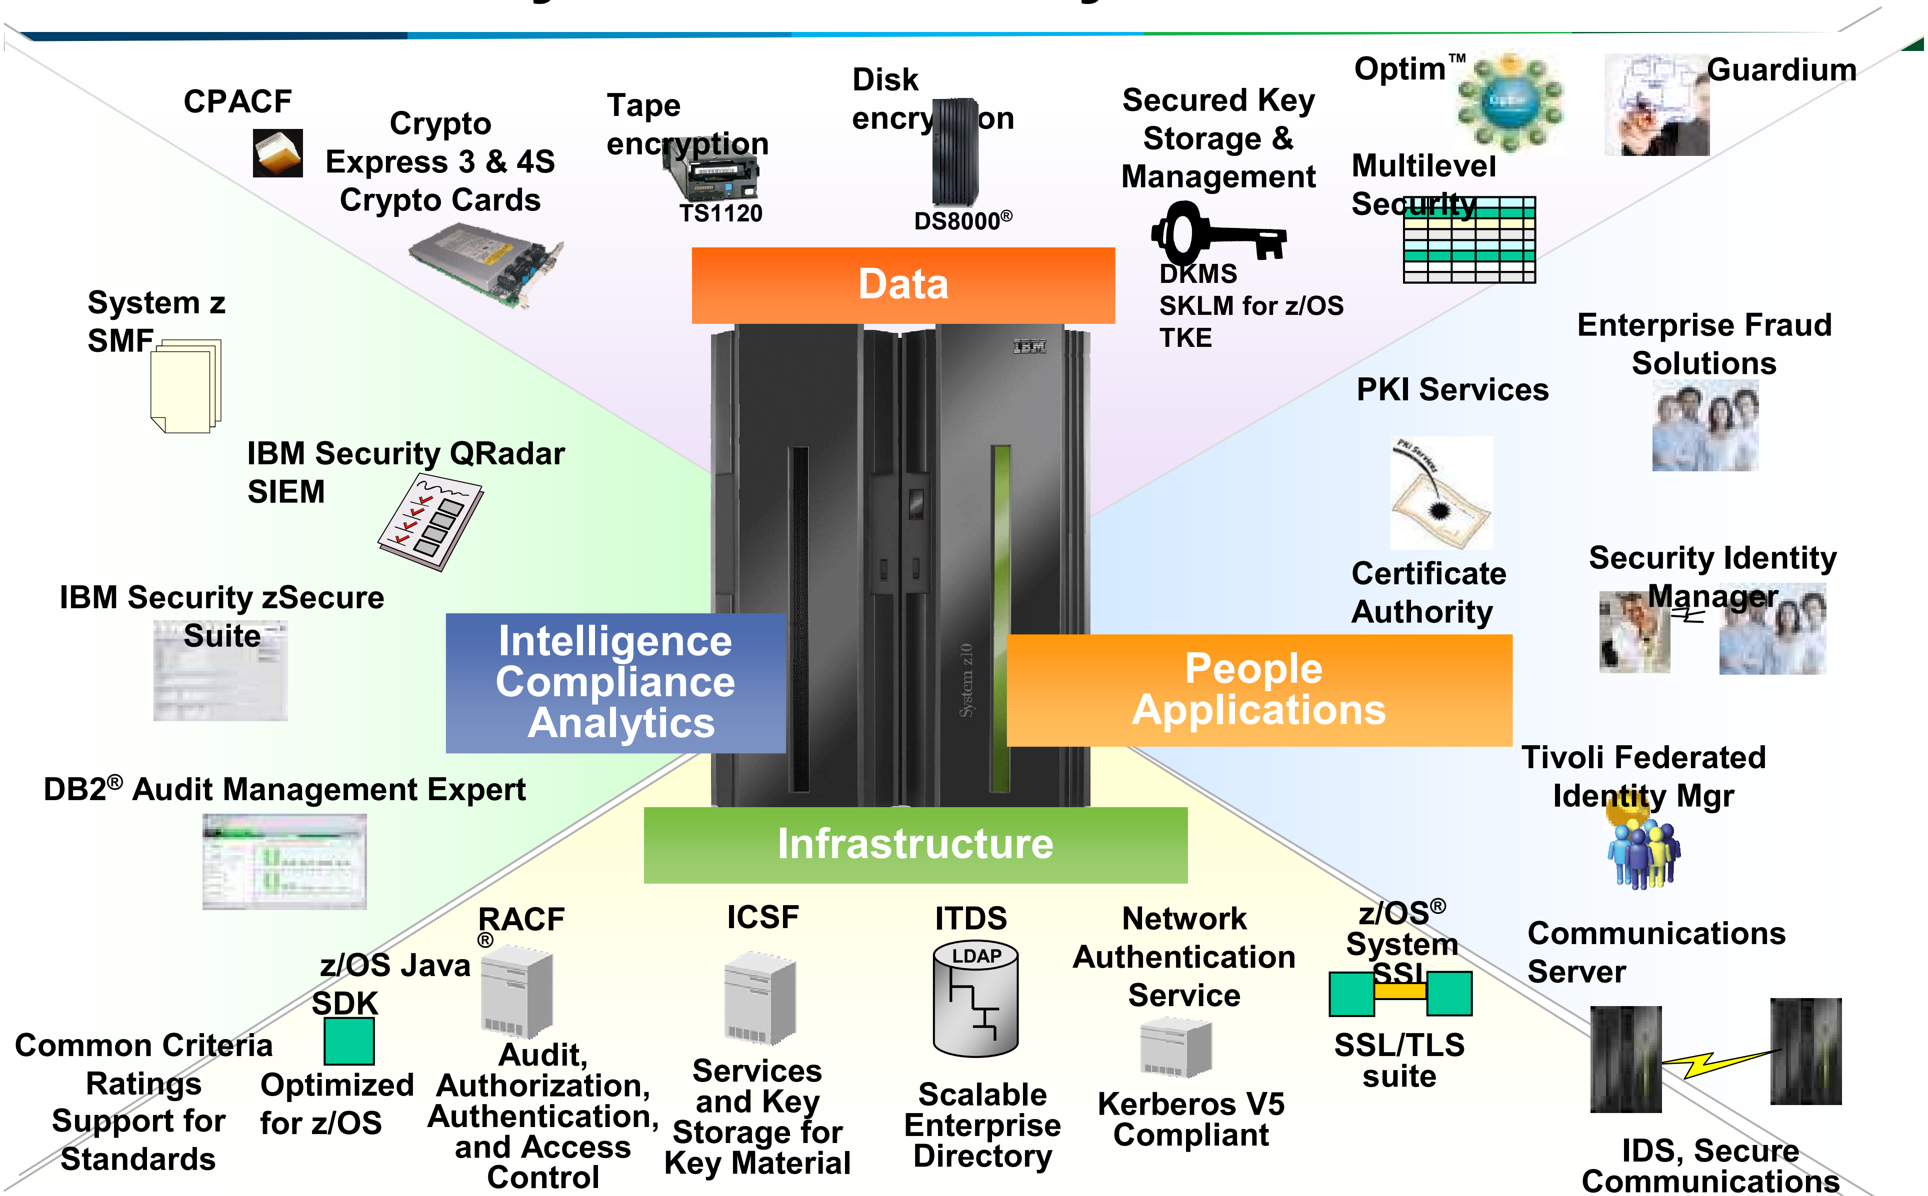

Virtualized System z security is superior to other platforms and augmentation costs less

Security Natively Covered by Platform

Security Level Description	MF	x86	UNIX
Normal corporate	100.00%	21.00%	35.00%
Credit card processing involved	100.00%	14.00%	26.00%
Banking	100.00%	8.00%	14.00%
Healthcare	100.00%	7.50%	11.00%
Research	78.00%	3.00%	8.00%
Defense	64.00%	1.00%	3.00%

Major security deficiencies on distributed platforms

Distributed platforms require ***considerable additional expense***

On System z most security requirements are standard

Incremental Cost to Achieve Required Security

Security Level Description	MF	x86	UNIX
Normal corporate	0.00%	25.20%	12.10%
Credit card processing involved	0.00%	38.40%	16.90%
Banking	0.00%	63.70%	22.40%
Healthcare	0.00%	81.60%	30.70%
Research	2.10%	134.80%	56.90%
Defense	4.30%	187.90%	97.50%

Little additional augmentation required on System z

Source: "Comparing Virtualization Alternatives – What's best for your business?"
© 2012, Solitaire Interglobal Ltd. http://public.dhe.ibm.com/common/ssi/ecm/en/zsl03192usen/ZSL03192USEN.PDF

The Choice Is Clear!

System z
is better than Intel-based systems
for
Systems of Record

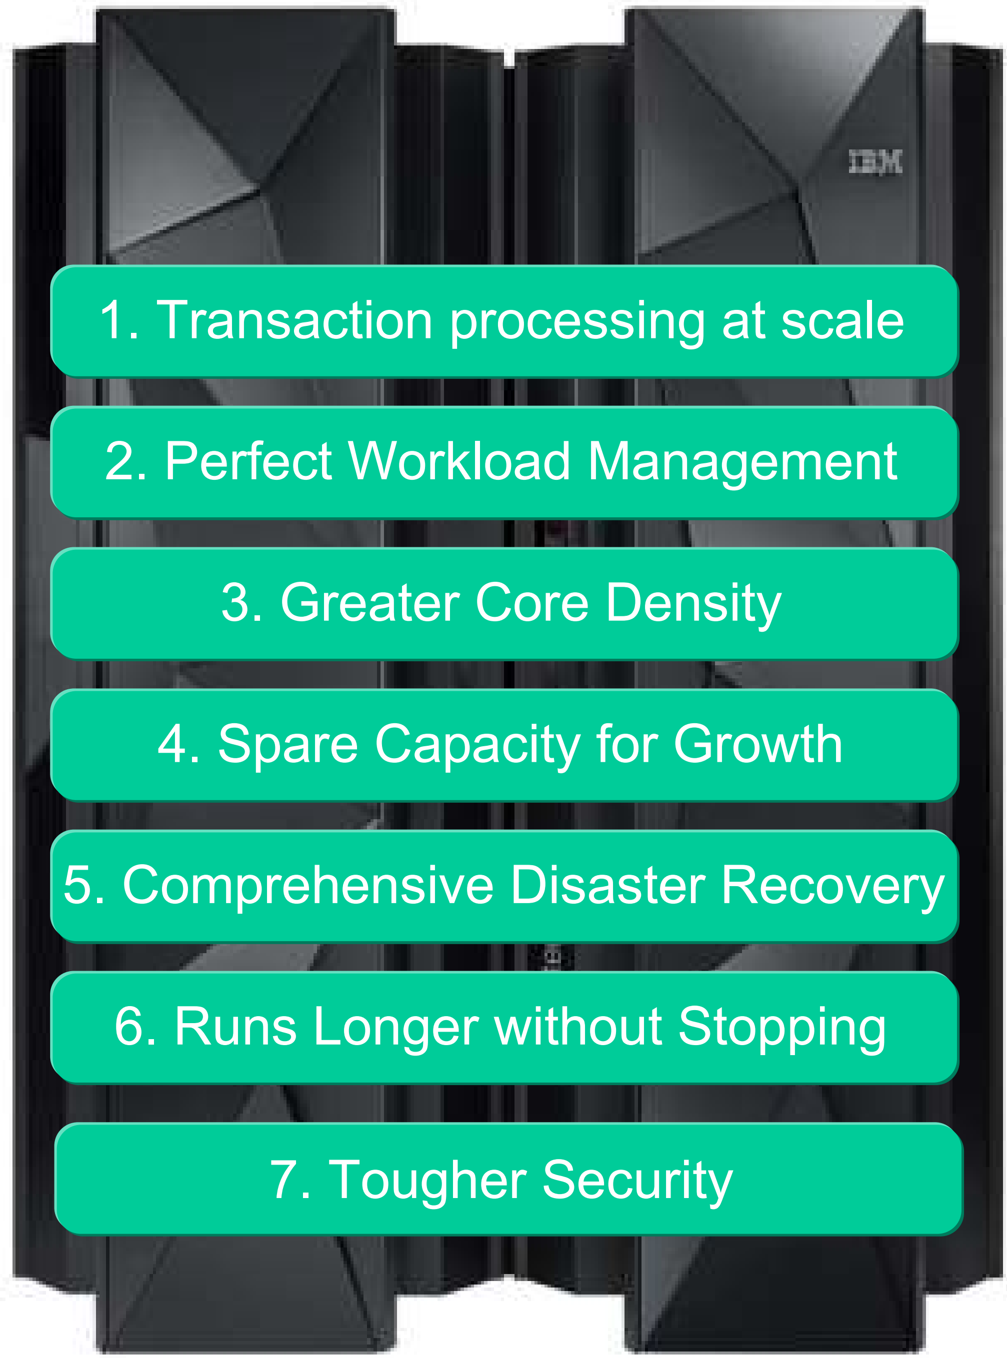

The IBM Eagle team helps customers understand mainframe costs and value

- **Worldwide** team of senior technical IT staff
- **Free of Charge** Total Cost of Ownership (TCO) studies
  - Help customers evaluate the lowest cost option among alternative approaches
  - Includes a one day on-site visit and is **specifically tailored to a customer's enterprise**
- Studies cover POWER, PureSystems and Storage accounts in addition to System z
  - For both IBM customer and Business Partner customer accounts
- Over 300 customer studies since formation in 2007
- Contact: eagletco@us.ibm.com

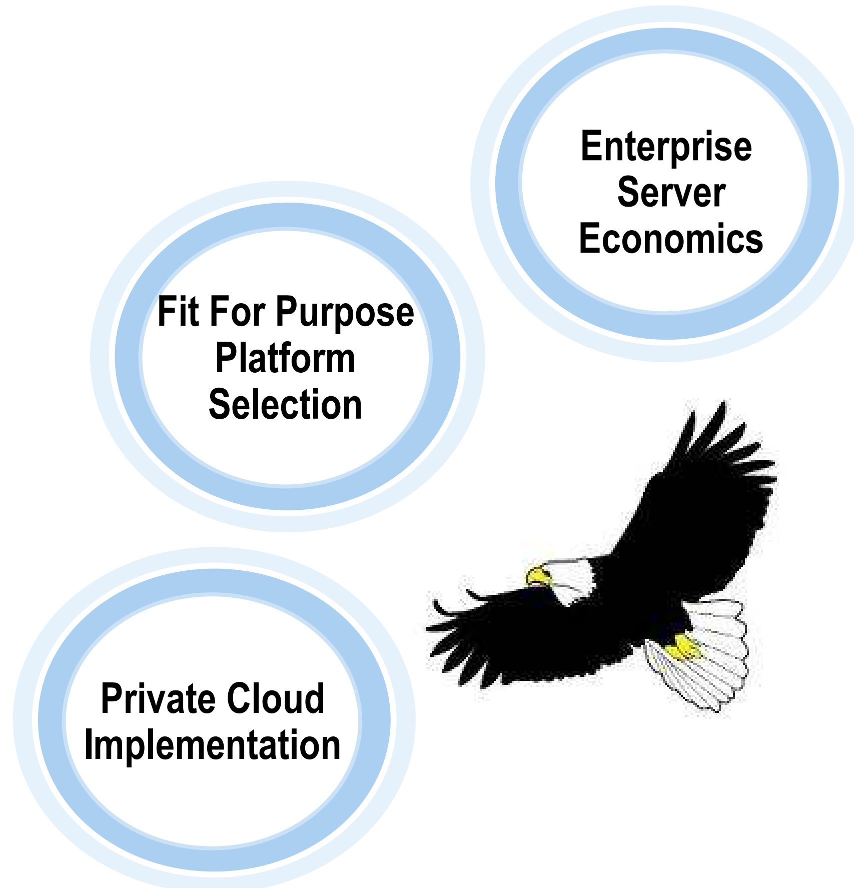

For More Information please contact…

Len Santalucia, CTO & Business Development Manager
Vicom Infinity, Inc.
One Penn Plaza – Suite 2010
New York, NY 10119
212-799-9375 office
917-856-4493 mobile
lsantalucia@vicominfinity.com

About Vicom Infinity

Account Presence Since Late 1990's
IBM Premier Business Partner
Reseller of IBM Hardware, Software, and Maintenance
Vendor Source for the Last 4 Generations of Mainframes/IBM Storage
Professional and IT Architectural Services
Vicom Family of Companies Also Offer Leasing & Financing, Computer Services, and IT Staffing & IT Project Management