

# IBM System z Technology Summit



zEnterprise – The First System Of Systems

**Virtualization & Consolidation  
On zEnterprise**

**Victor Leith**

27<sup>th</sup> Jan 2011




# A Deeper Look At Some Topics

---

- How was “fit for purpose” determined?
- Why was Linux on z/VM best for the heavy I/O workloads?
- Network simplification with zEnterprise
- Storage simplification with zEnterprise

# Virtualization Concept



## Virtual Resources


- Proxies for real resources: **same interfaces/functions, different attributes.**
- May be part of a physical resource or multiple physical resources.

## Virtualization

- Creates virtual resources and "maps" them to real resources.
- Primarily accomplished with software and/or firmware.

## Resources

- Components with **architected interfaces/functions.**
- May be centralized or distributed. Usually physical.
- Examples: memory, disk drives, networks, servers.




- Separates presentation of resources to users from actual resources
- Aggregates pools of resources for allocation to users as virtual resources

# Server Virtualization Business Value

## Roles:

- Consolidations
- Dynamic provisioning/hosting
- Workload management
- Workload isolation
- Software release migration
- Mixed production and test
- Mixed OS types/releases
- Reconfigurable clusters
- Low-cost backup servers



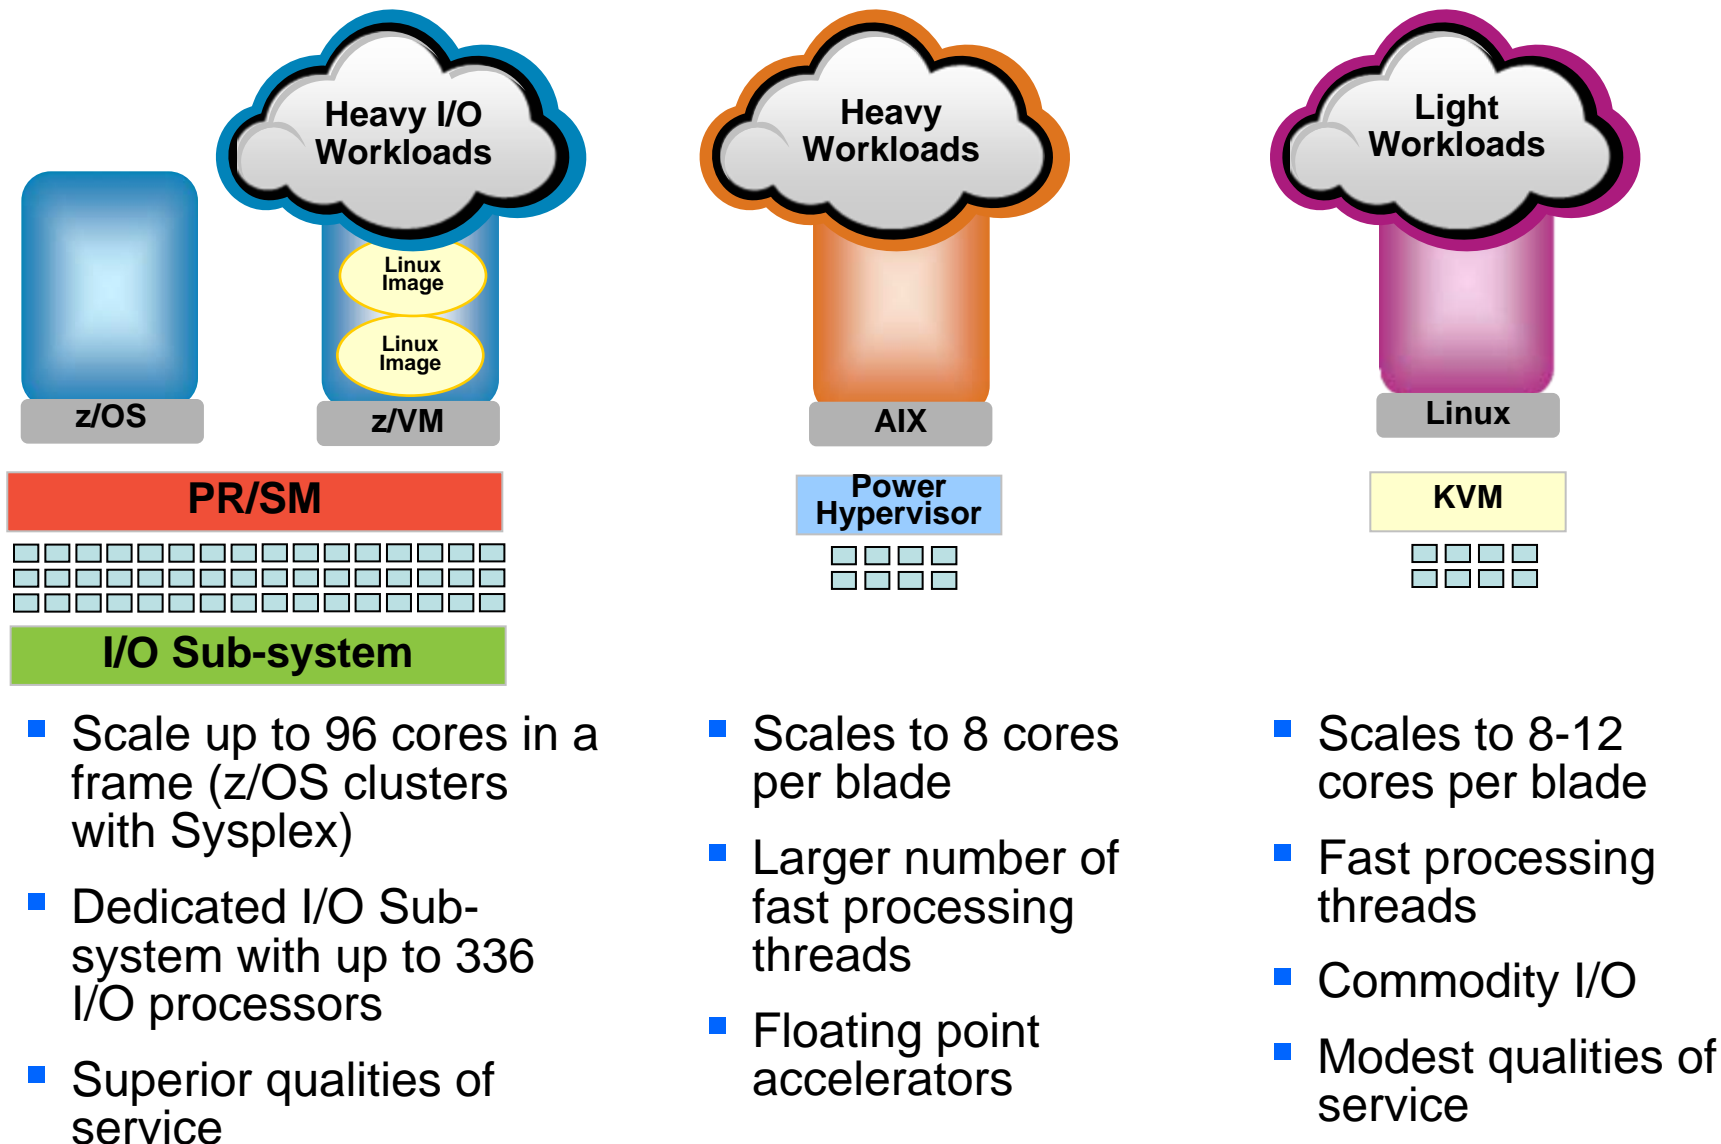
## Benefits:

- Higher resource utilization
- Greater usage flexibility
- Improved workload QoS
- Higher availability / security
- Lower cost of availability
- Lower management costs
- Improved interoperability
- Legacy compatibility
- Investment protection

In the final analysis, the virtualization benefits take three forms:

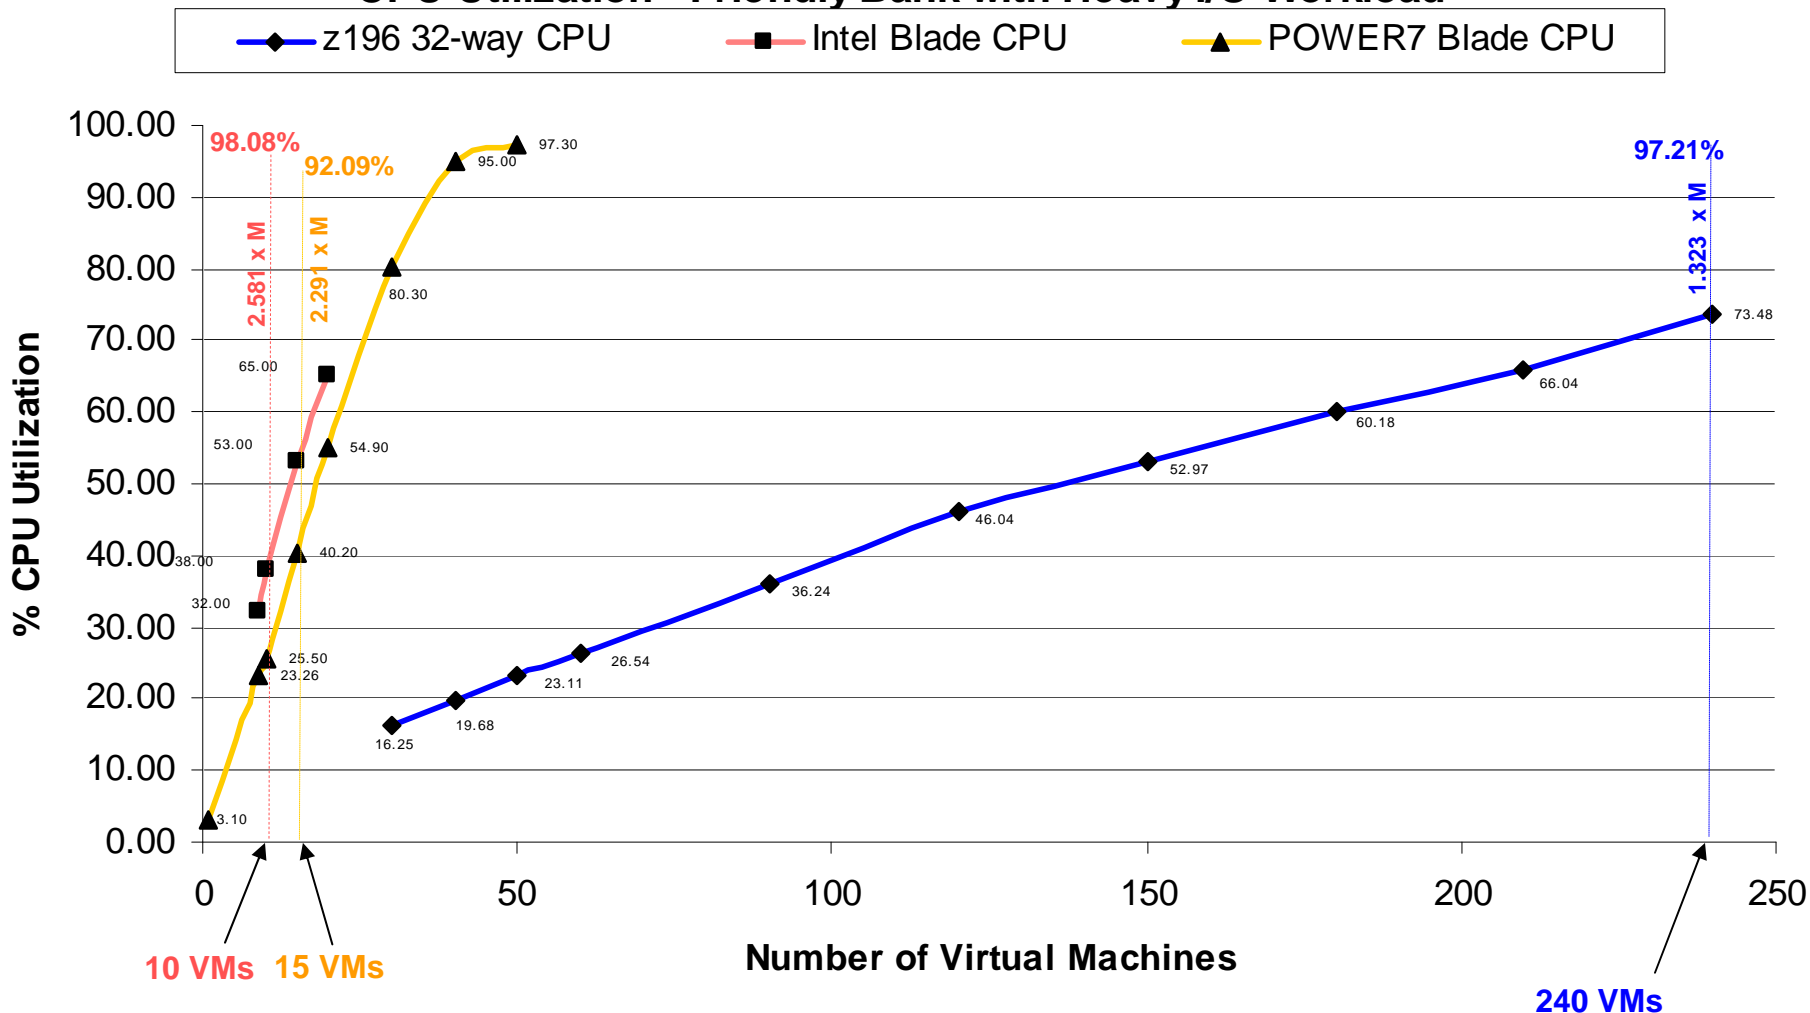
- **Reduced hardware costs**
  - ▶ Higher physical resource utilization
  - ▶ Smaller footprints
- **Improved flexibility and responsiveness**
  - ▶ Virtual resources can be adjusted dynamically to meet new or changing needs and to optimize service level achievement
  - ▶ Virtualization is a key enabler of on demand operating environments
- **Reduced management costs**
  - ▶ Fewer physical servers to manage
  - ▶ Many common management tasks become much easier

# zEnterprise Extends Cost Advantages To A Broader Range Of Workloads



# Consolidation Ratios For Distributed Workloads With Heavy I/O


## CPU Utilization - Friendly Bank with Heavy I/O Workload



Consolidation ratios derived from IBM internal studies. z196 32-way performance projected from z196 8-way and z10 32-way measurements. zBX with x86 blades is a statement of direction only. Results may vary based on customer workload profiles/characteristics.

# Deploying Workloads With Heavy I/O Requirements

*Benchmark to determine which platform provides the lowest TCA over 3 years*



- IBM WebSphere ND
- Monitoring software
- On 4 core "Older" Intel

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

10 workloads per Intel blade




Virtualized on Intel  
8 core blade  
**\$21,413** per workload

15 workloads per POWER7 blade



PowerVM on PS701  
8 core blade  
**\$14,325** per workload

240 workloads per 32-way z/VM




**I/O bandwidth large scale pool**

z/VM on zEnterprise CPC  
32 IFLs  
**\$14,052** per workload

Consolidation ratios derived from IBM internal studies. z196 32-way performance projected from z196 8-way and z10 32-way measurements. zBX with x86 blades is a statement of direction only. Results may vary based on customer workload profiles/characteristics. Prices will vary by country.

# Deploying Heavy Workloads

*Benchmark to determine which platform provides the lowest TCA over 3 years*



- IBM WebSphere ND
- Monitoring software
- On 8 core Nehalem servers

Online banking workloads, each driving **460** transactions per second with light I/O

1 workload per Intel blade



Virtualized on Intel  
8 core blade  
**\$214,133** per workload

2 workloads per POWER7 blade



PowerVM on PS701  
8 core blade  
**\$107,437** per workload

more parallel threads

23 workloads per 32-way z/VM




z/VM on zEnterprise CPC  
32 IFLs  
**\$146,631** per workload

Consolidation ratios derived from IBM internal studies. z196 32-way performance projected from z196 8-way and z10 32-way measurements. zBX with x86 blades is a statement of direction only. Results may vary based on customer workload profiles/characteristics. Prices will vary by country.



# Deploying Light Workloads

*Benchmark to determine which platform provides the lowest TCA over 3 years*



- IBM WebSphere ND
- Monitoring software
- On 4 core "Older" Intel

Online banking workloads, each driving **22** transactions per second with light I/O

36 workloads per Intel blade



Fast low cost threads

Virtualized on Intel 8 core blade  
**\$5,948** per workload

34 workloads per POWER7 blade



PowerVM on PS701 8 core blade  
**\$6,320** per workload


270 workloads per 32-way z/VM



z/VM on zEnterprise CPC 32 IFLs  
**\$12,491** per workload

Consolidation ratios derived from IBM internal studies. z196 32-way performance projected from z196 8-way and z10 32-way measurements. zBX with x86 blades is a statement of direction only. Results may vary based on customer workload profiles/characteristics. Prices will vary by country.

# Options For Deploying Distributed Workloads – Best Fit Strategy On zEnterprise Produces Lowest Cost



Deploy all distributed workloads on x blades

**\$202M**



Deploy all distributed workloads on p blades

**\$134M**



Deploy all distributed workloads on Linux on System z

**\$196M**



Best Fit deployment on zEnterprise (Linux on System z, x blade, p blade)

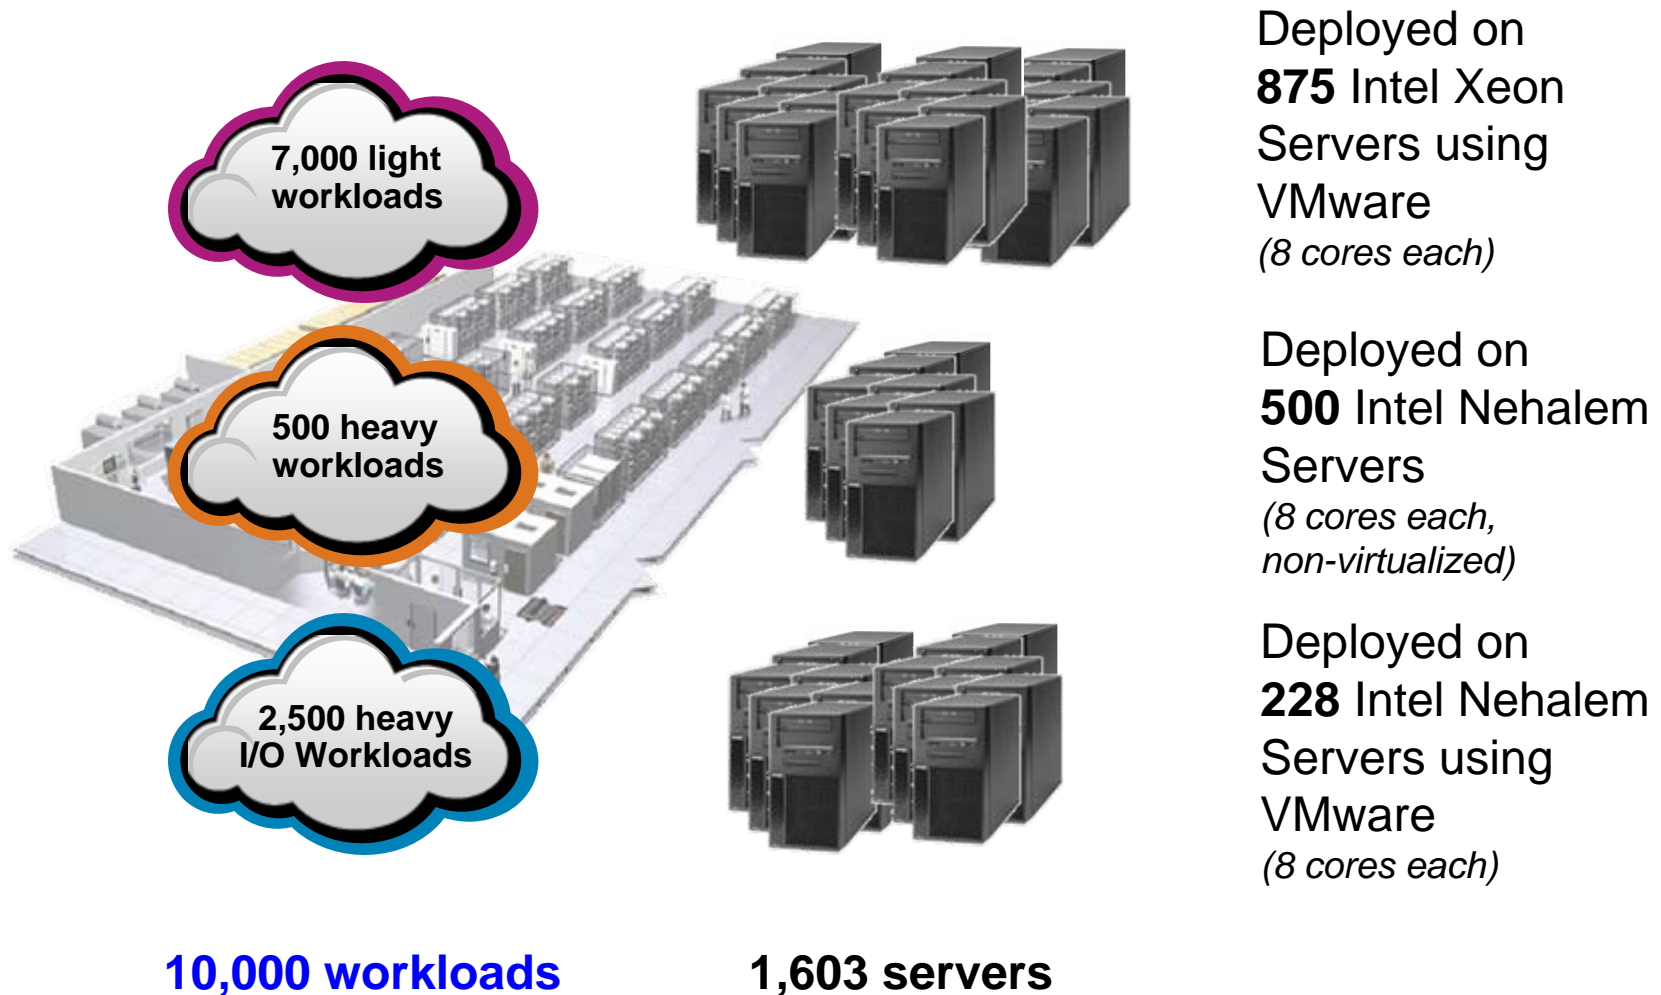
**\$130 M**



**36% less**

Consolidation ratios derived from IBM internal studies. z196 32-way performance projected from z196 8-way and z10 32-way measurements. zBX with x86 blades is a statement of direction only. Results may vary based on customer workload profiles/characteristics. Prices will vary by country.

# Large Data Center – What Did It Cost To Deploy 10,000 Workloads On Virtualized Intel Servers?



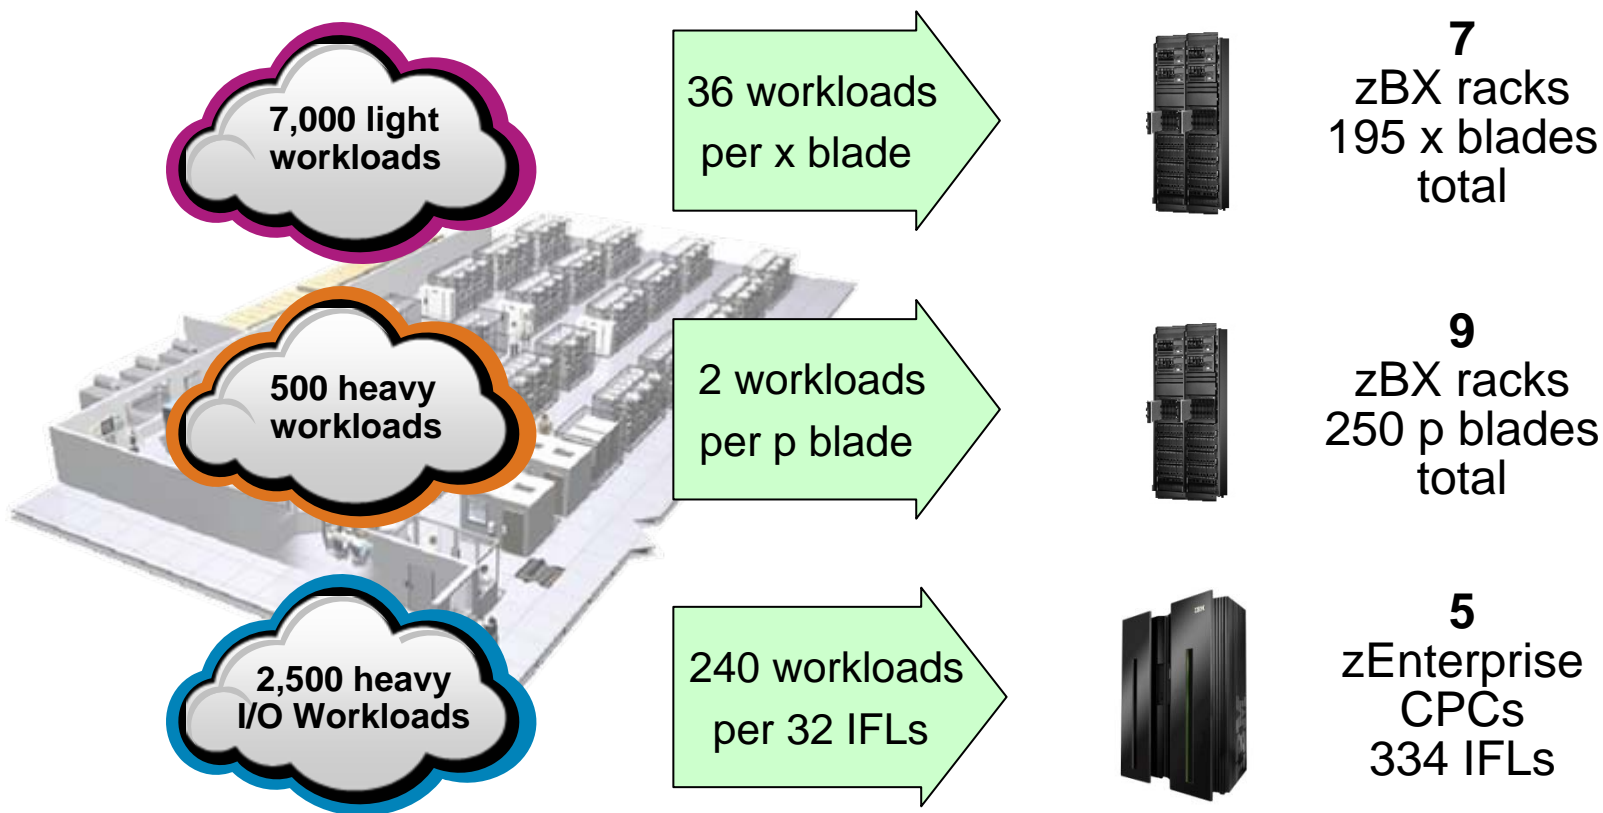
Deployed on  
**875** Intel Xeon  
Servers using  
VMware  
*(8 cores each)*

Deployed on  
**500** Intel Nehalem  
Servers  
*(8 cores each,  
non-virtualized)*

Deployed on  
**228** Intel Nehalem  
Servers using  
VMware  
*(8 cores each)*

IBM analysis of a customer scenario with 10,000 distributed workloads. Deployment configuration is based on consolidation ratios derived from IBM internal studies.


# Large Data Center – What Does It Cost To Deploy 10,000 Workloads On zEnterprise?



**Best fit assignments**

Configuration is based on consolidation ratios derived from IBM internal studies. z196 32-way performance projected from z196 8-way and z10 32-way measurements. The zBX with x86 blades is a statement of direction only. Results may vary based on customer workload profiles/characteristics.

# Compare Server Cost Of Acquisition



Server configurations are based on consolidation ratios derived from IBM internal studies. Prices are in US currency, prices will vary by country

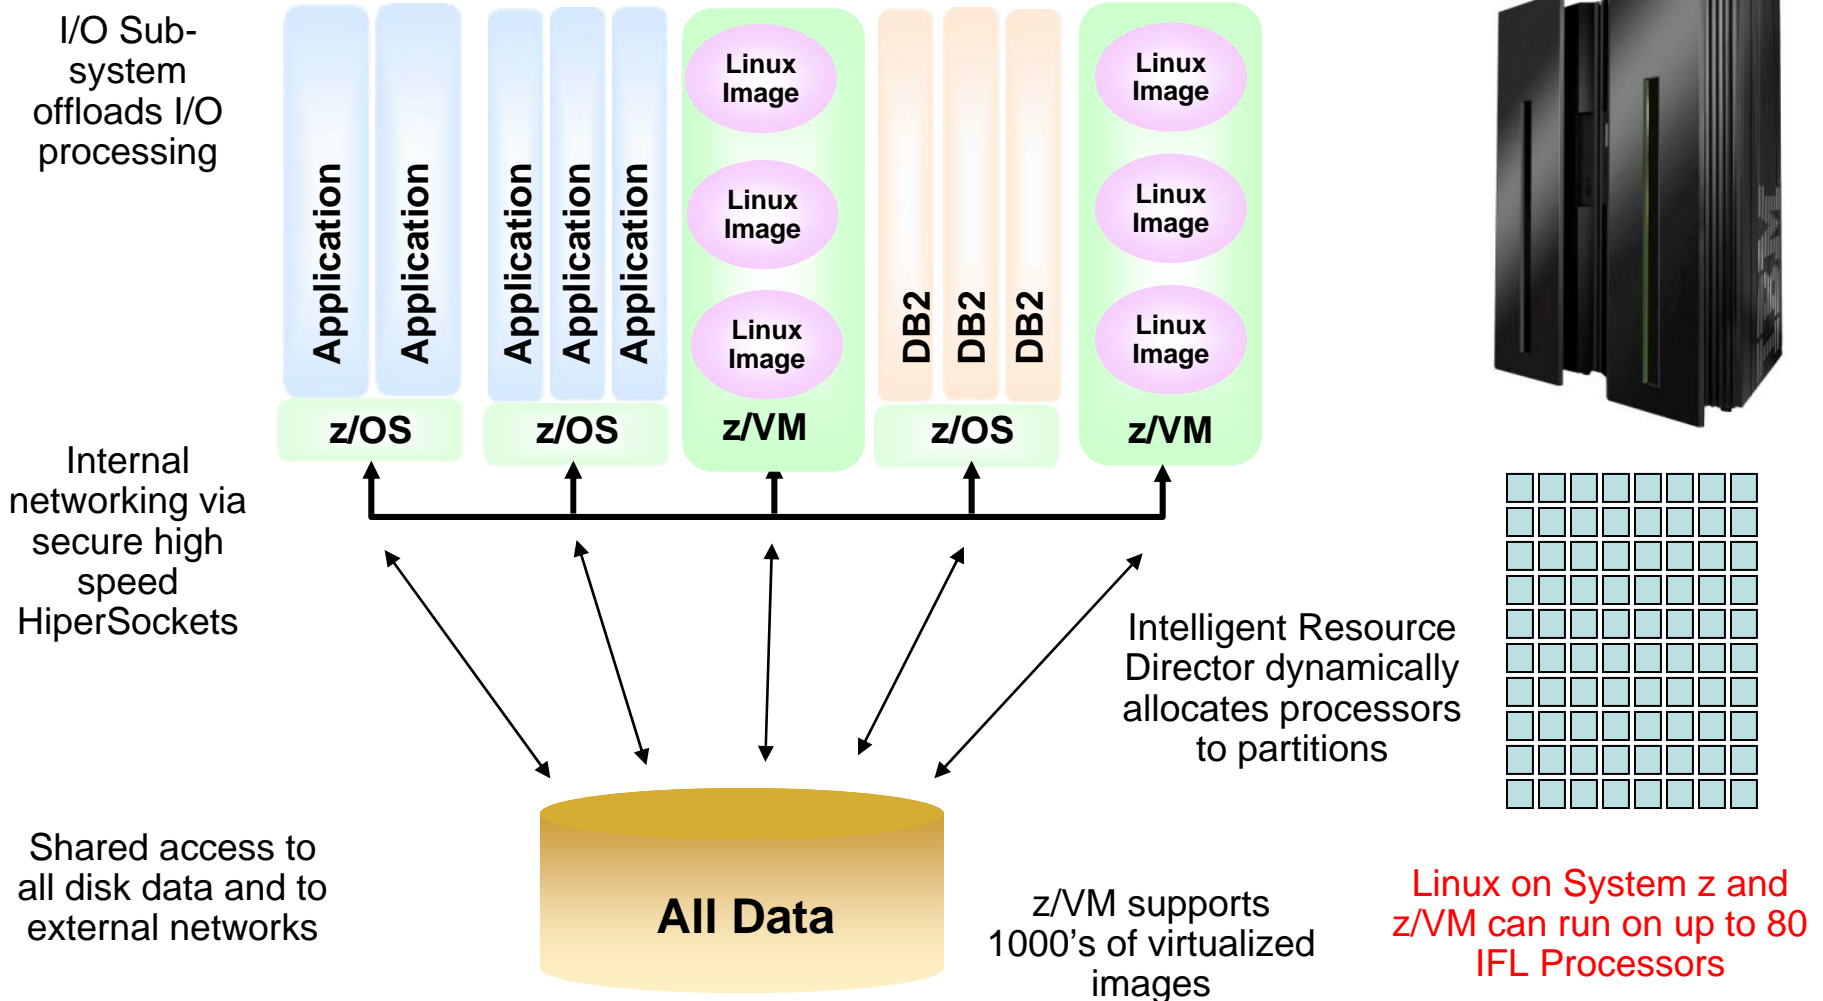
**59% less**

# Linux On z196 Achieves Lowest TCA For Heavy Processing And I/O Workloads

- Larger scale of shared processor pools (32 cores vs. 8 cores)
- Statistical benefit of sharing a larger pool of processors
- Software priced per core
- Cost benefit of Enterprise Linux Server Solution Edition pricing
- Dedicated I/O Sub-system offloads I/O processing
- Greater I/O bandwidth
- Virtualization of I/O processing resources
- Built-in storage virtualization and switching

# z196 Is Designed For Large Scale Virtualization And Consolidation

Logical Partitions Share Processors, Common Cache Structures, and I/O



# z/VM On System z – Optimized For Large Scale Virtualization

- Large scale virtualization yields pooling benefits
  - ▶ Shared processor pool
  - ▶ Lower headroom requirement to accommodate variations in workload demand
- On System z, up to 32 IFL processor cores can be supported by a single z/VM LPAR
  - ▶ Large scale virtualization platform can support hundreds of virtual machines
- zBX blades are limited to 8-12 cores (currently)




# System z Solution Edition For Enterprise Linux And The Enterprise Linux Server

*Transforming the economics of large scale integration at a special packaged price!*

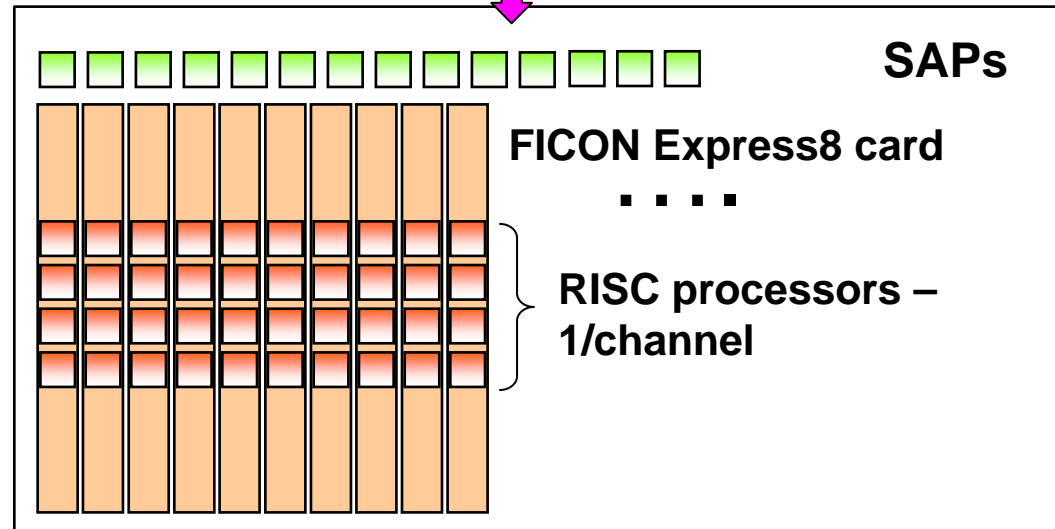
- System z Solution Edition for Enterprise Linux
  - ▶ Integrated Facility for Linux (IFL) processors, memory and z/VM added to an existing mainframe
  - ▶ Hardware and software maintenance for three or five years
- Enterprise Linux Server
  - ▶ Standalone System zEnterprise server with IFLs, memory, I/O connectivity, and z/VM
  - ▶ Hardware and software maintenance for three or five years
- Linux on System z available from distribution partners
  - ▶ (Novell SUSE and Red Hat)




# z196 - Optimized For High I/O Bandwidth



- Up to 80 General Purpose (GP) or Specialty Engine processors
  - ▶ Execute business logic



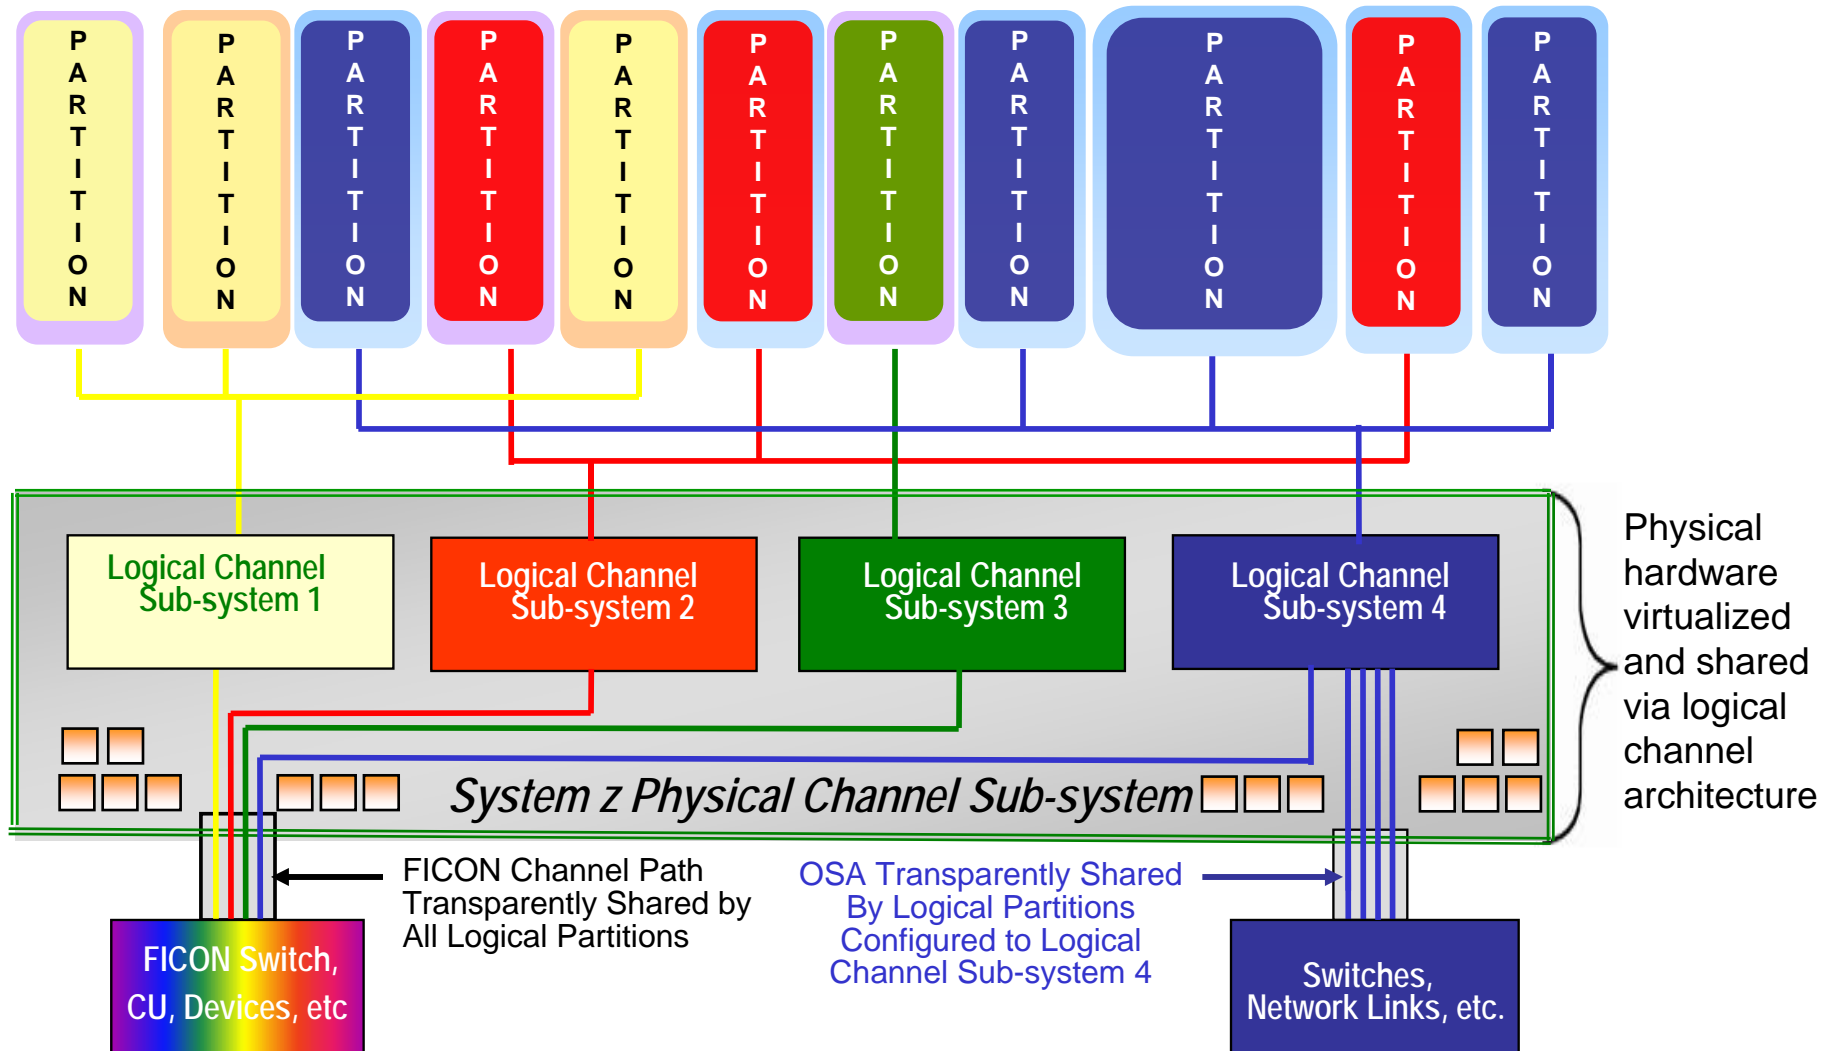
- Up to 14 System Assist Processors (SAP) to manage I/O requests
  - ▶ Can sustain up to **2.2M IOPS<sup>1</sup>**
- Logical Channel Sub-system virtualizes I/O
  - ▶ Up to 1024 logical channels
- Up to 84 physical FICON cards for I/O transfers
  - ▶ Up to **336 RISC channel I/O processors**
  - ▶ High Performance FICON connections (zHPF)



- IBM DS8800 Storage System
  - ▶ Up to **440K IOPS capability** with zHPF
- Benefits both z/OS and z/VM workloads


<sup>1</sup>Recommend 70% max SAP Utilization – 1.5M IOPS

# Physical I/O Adapters And Channels Are Virtualized And Shared By The Consolidated Workloads



# z/VM Security For Virtualization

- Operates without interference/harm from guest virtual machines
- Virtual machines cannot circumvent system security features
- Protects virtual machines from each other
- Ensures that a user only has access to resources specifically permitted
- Tracks who is accessing all system resources
- LPAR certified Common Criteria EAL5
- z/VM certified at Common Criteria EAL4+
- HiperSockets for highly secure internal networking
- Access to System z Crypto features
  - ▶ CPACF, CryptoExpress3

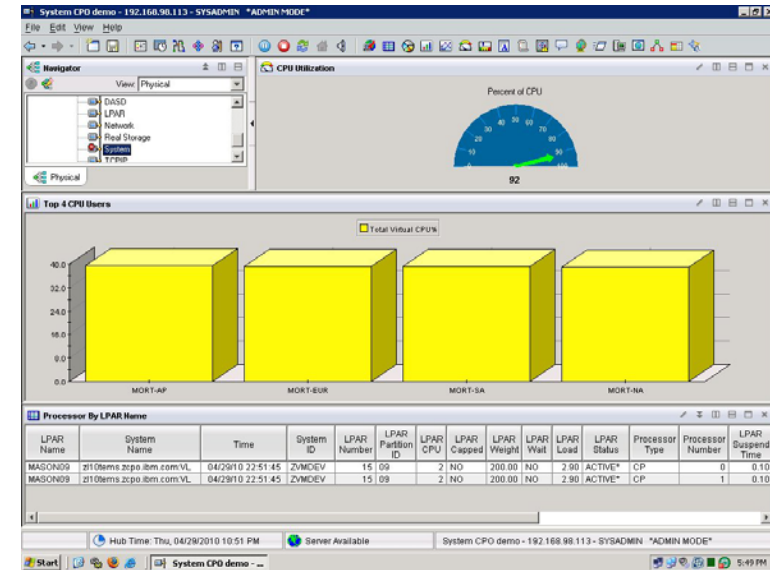


# Linux On System z Workloads Inherit System z Qualities Of Service

- Reliability, availability, serviceability characteristics of System z
- Site failover for disaster recovery
- Capacity on demand upgrades
- Add physical processors to Linux environment without disruption

# DEMO: Dynamically Add New Processor To z/VM LPAR To Handle Increased Risk Analysis Workload

1. A customer has in-house Risk Analysis program running on Linux on System z
2. Increased workload to all 4 Linux guests is causing z/VM LPAR utilization of 90%+
3. Customer determines this is a long term trend - additional physical capacity needed
4. New capacity made available to LPAR as new Logical CPU, available for work
  - ▶ Without disruption in service



**VMware can't recognize and take advantage of additional physical processors without bringing down and rebooting the system**

**Note:** Assumes available processors on installed books

# Key Functional Distinctions Of z/VM vs VMware ESX

| Attribute  | z/VM V6.1  | VMware ESX 4.1  |
|--|--|---|
| Real CPU sharing   | Architecturally limitless; more than 60 VMs per CPU (workload dependent)           | Up to 25 VMs per CPU (workload dependent)                       |
| Architectural maximum number of VMs                          | Thousands per copy of z/VM   | 320 per copy of VMware  |
| Real capacity on demand                                      | Yes, non-disruptively  | No  |
| In-memory support  | Minidisk cache; Virtual Disks in Storage; DCSS (shared program executables)        | Shared virtual memory pages (detected via background operation) |
| Virtual Machine (VM) scalability                             | Up to 64 CPUs, 1 TB of memory  | Up to 8 CPUs, 255 GB of memory                                  |
| Run multiple hypervisor copies per server                    | Yes; up to 60 copies of z/VM on one mainframe                                      | No  |
| Command and control, monitoring, automation infrastructure   | Pervasive, robust, time-tested   | Modest, yet easy to use   |
| Resource over-commitment support (memory, CPU, network, I/O) | Extensive  | Modest  |
| Virtualization   | z/VM can run as a guest of itself to multiple levels (z/VM on z/VM on z/VM on ...) | No  |

VMware ESX is trying to catch up with the unrivaled capabilities of z/VM

## *Linux on the mainframe*

IBM first announced Linux software and services for the mainframe on 17 May 2000. At the 10th anniversary of this announcement, more than 70 of the top 100 worldwide mainframe customers were running Linux.

# IBM's Approach To Virtualization Is Superior to VMware For Enterprise-wide Consolidation

## *VMware is a viable solution for Windows-based smaller scale projects, but:*

- Has limitations on those attributes important to large-scale consolidations
- x86 platforms lack the sophisticated I/O subsystem needed to efficiently handle heavy I/O workloads
- Only supports one copy/physical server
- vCenter will only allow monitoring, reporting and provisioning of Virtual Machines on x86 physical servers

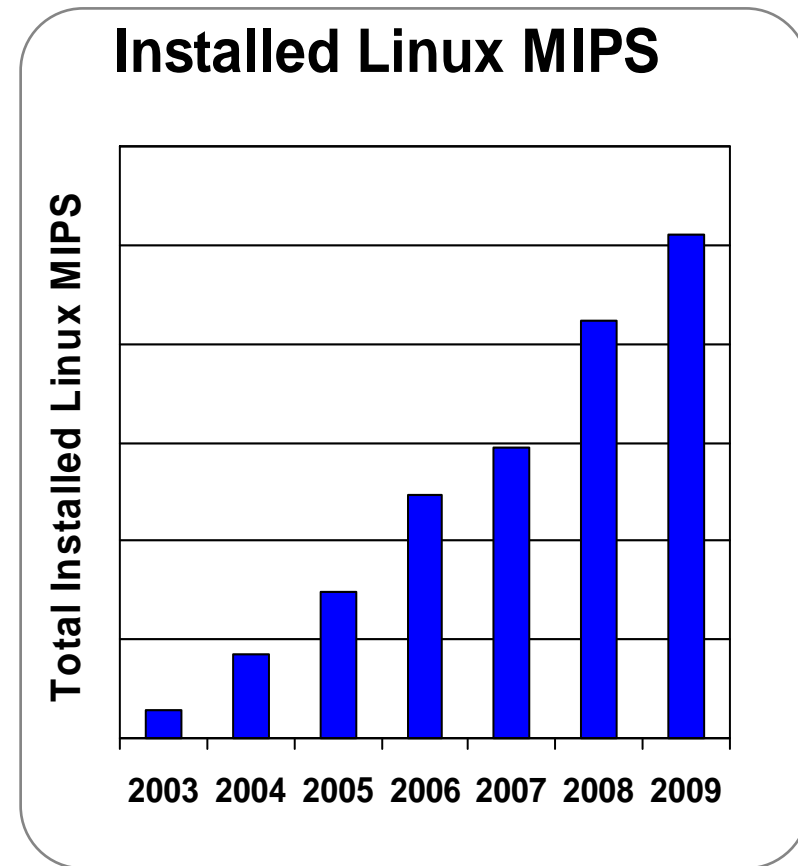
## *z/VM is the more scalable, flexible, resilient solution:*

- Architecturally superior in terms of real CPU sharing, virtual machine scaling, ability to dynamically add real capacity
- A single physical server can run up to 60 copies of z/VM, enabling failover, workload isolation, and scalability without duplicating hardware
- z/VM consolidation on a single footprint makes systematic disaster recovery easier vs. multiple x86 servers
- With zEnterprise and Tivoli software, can manage an entire Data Center, involving multiple platforms, from one central hub
- zEnterprise with the Unified Resource Manager provides mainframe-like governance and qualities of service across workloads that extend beyond the boundaries of IBM System z to multiple platform environments



# Client Adoption Drives Linux Success Installed Linux MIPS At 43% CAGR<sup>1</sup>

- The momentum continues:
  - ▶ **Shipped IFL MIPS increased 65% from YE07 to YE09**
- Linux is 16% of the System z customer install base (MIPS)
- 70% of the top 100 System z clients are running Linux on the mainframe
- >3,100 applications available for Linux on System z



<sup>1</sup>Based on YE 2004 to YE 2009

# Compare Network Cost Of Acquisition



Additional network parts

313 switches

7,038 cables

6,412 adapters

**13,763** total network parts

**\$3.8M** TCA

Additional network parts

7 switches

142 cables

74 adapters

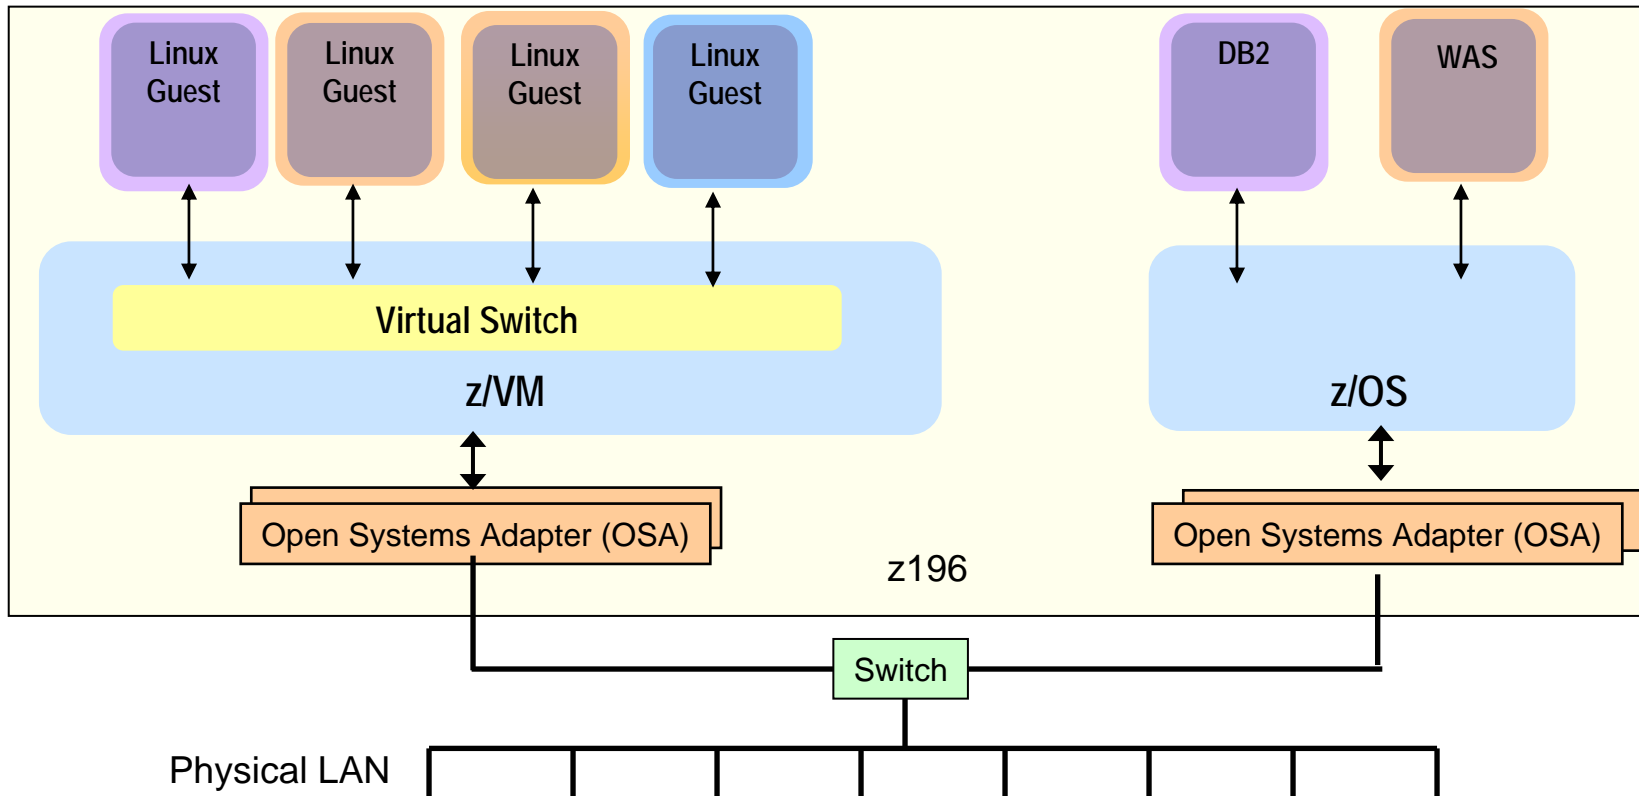
**223** total network parts

**\$197K** TCA

**95% less**

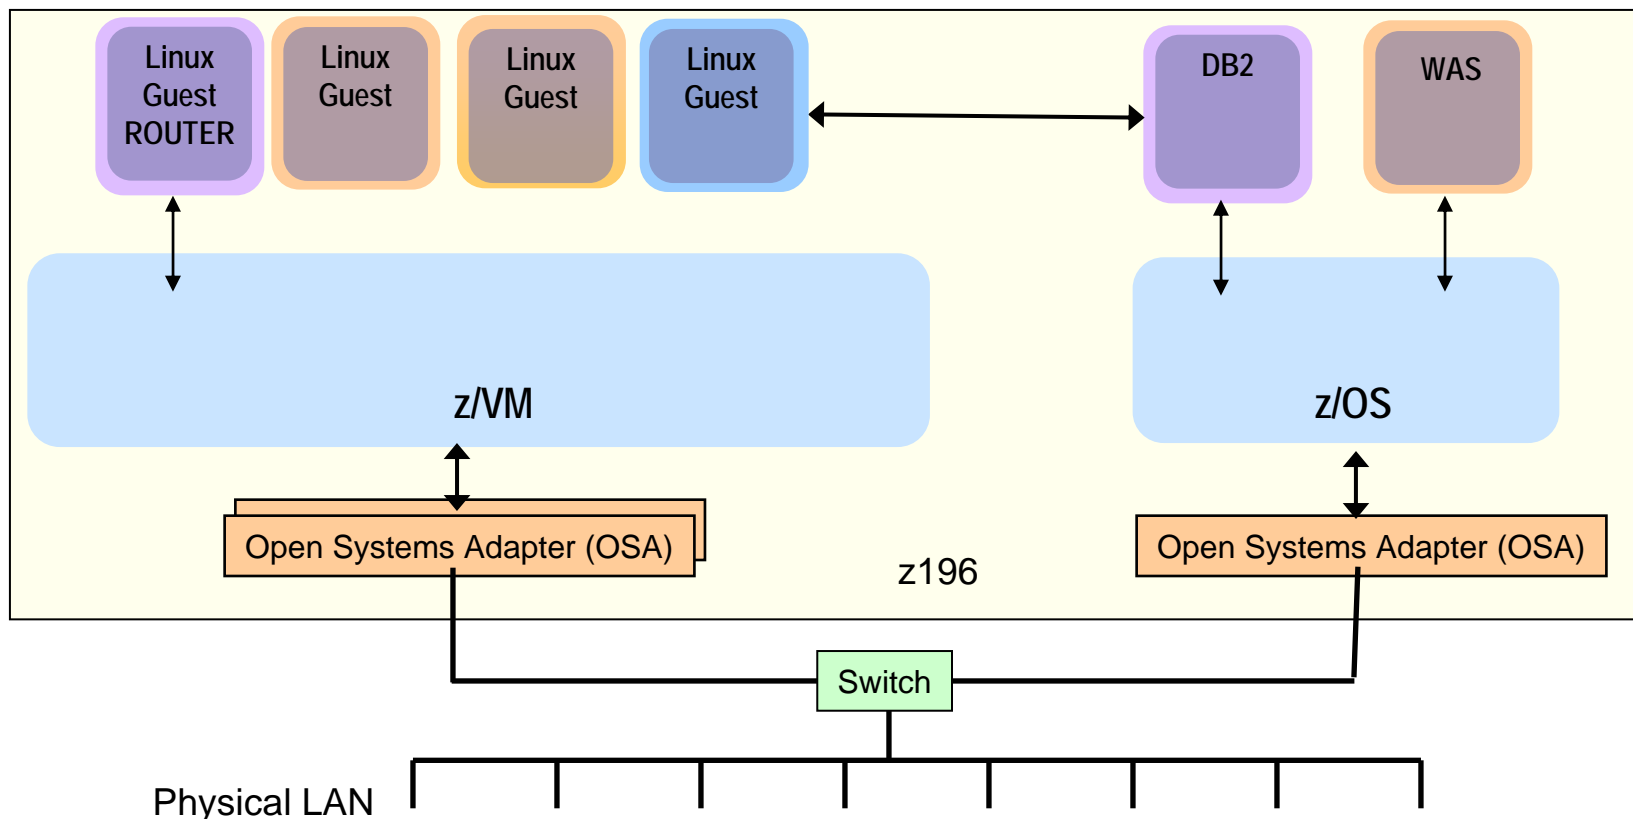
Network configuration is based on IBM internal studies.  
Prices are in US currency, prices will vary by country

# System z Features Enable Network Simplification – z/VM Virtual Switch



- Linux guests can talk to each other via z/VM virtual switch – memory speed
- Linux guests can talk to outside world via z/VM virtual switch connected to shared OSA adapter
- Attach up to 8 physical OSA ports - redundancy, balancing
- Dynamically add new physical OSA to support Linux workload growth

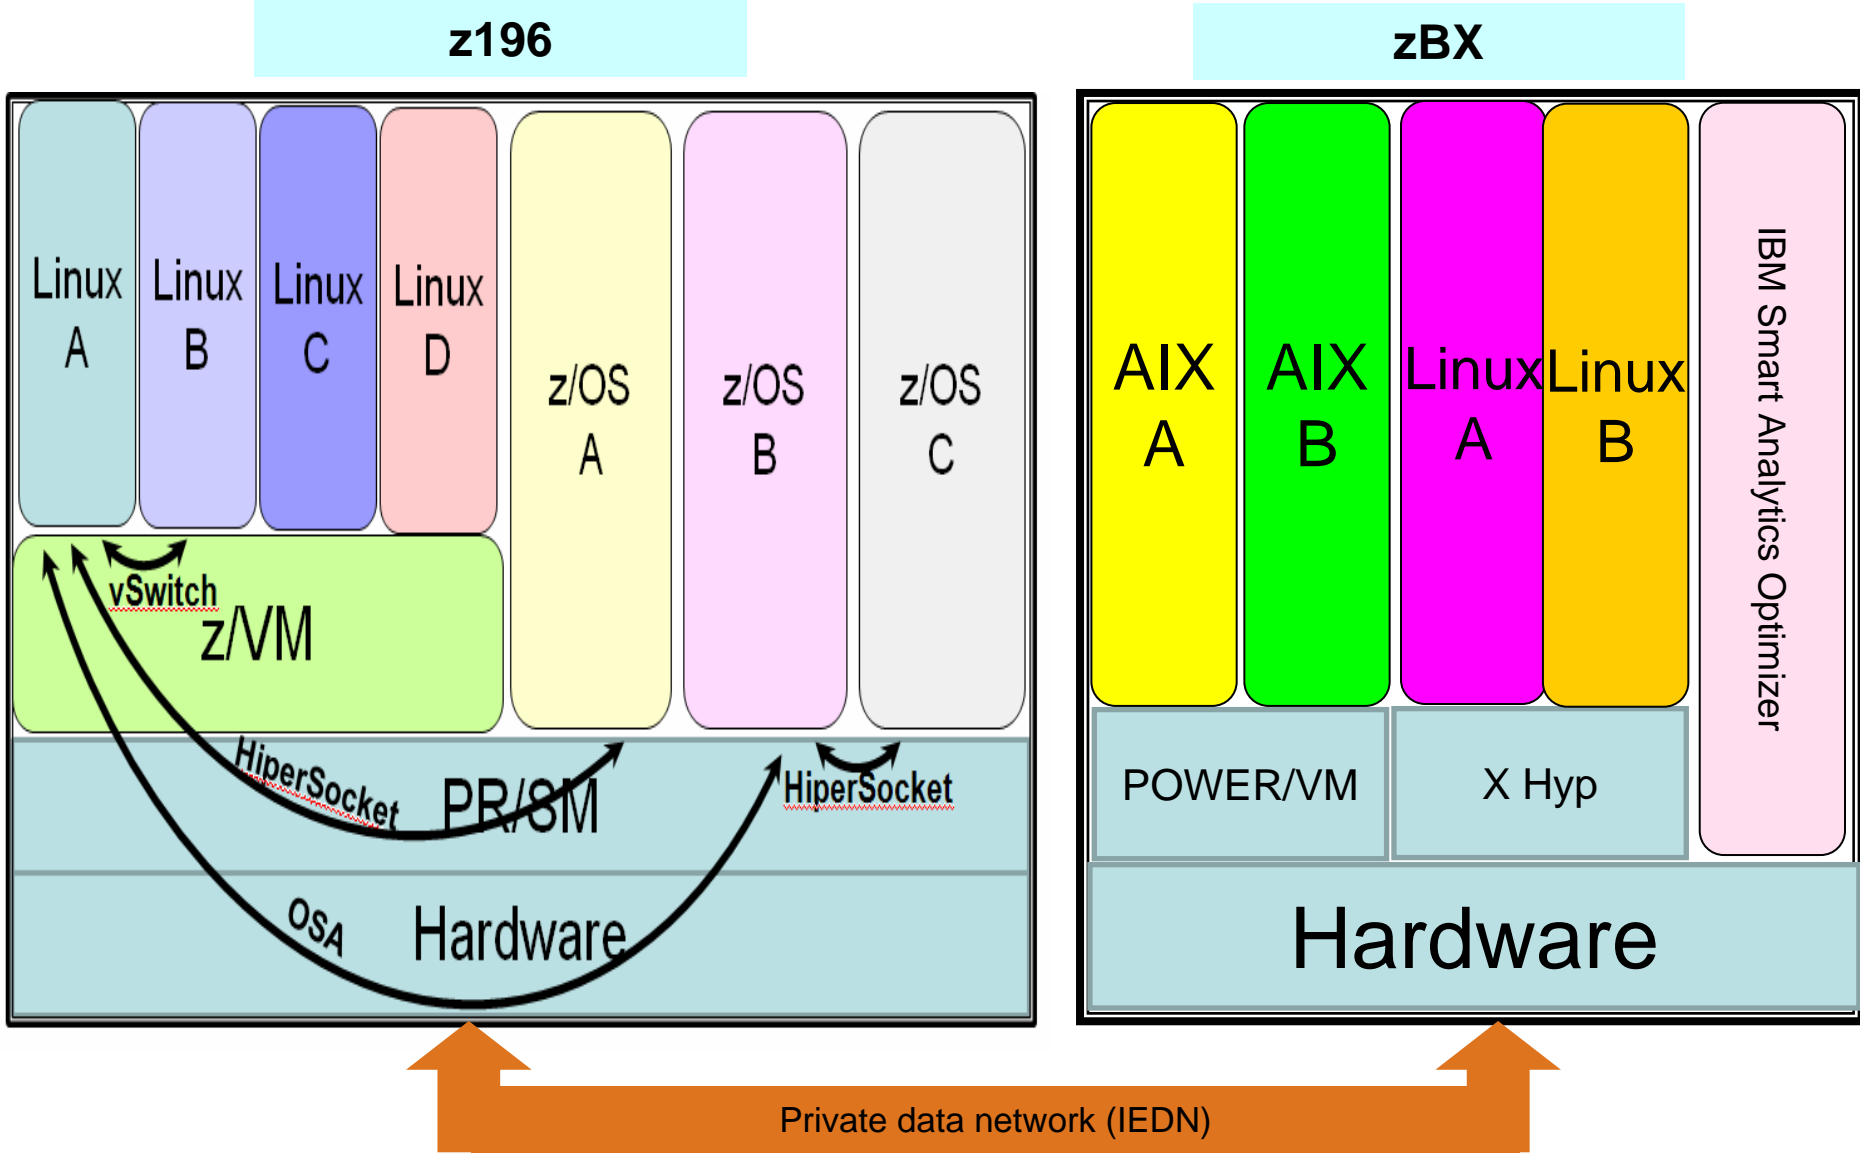
# System z Features Enable Network Simplification – HiperSockets



- Linux guests can talk to z/OS applications
- **Secure** IP communication at memory speed

- Close integration of data-intensive applications with database
- Reduces network management and physical assets

# Network Simplification Extends To The zBX



# Compare Storage Cost



**7.7 PB** embedded storage

31% utilization

1,603 points of admin

**\$211M** TCO (3 years)

240GB active storage required per workload (2.4PB total)

**4.5 PB** provisioned storage

53% utilization

10 points of admin

**\$108M** TCO (3 years)

**49% less**

# IBM System Storage – Optimized For Different Requirements



DS8700

- Mix of random and sequential I/O
- Highest availability and performance with High Performance FICON, large cache, and Easy Tier for SSDs



XIV

- Mostly random block I/O
- Ideal for distributed apps
- Exceptional ease of use and management productivity




SONAS

- Mostly sequential file server I/O
- Scalable network storage
- Ideal for consolidating distributed filers

# Best Fit Storage

Distributed light workload -  
240GB active storage  
55% block/45% file



7 zBX racks  
with x blades

+




XIV  
6 via SAN



SONAS  
1

Distributed heavy workload -  
240 GB active storage  
55% block/45% file





9 zBX racks  
with p blades

+



XIV  
1 via SAN

Distributed light workloads with heavy I/O -  
240 GB active storage  
100% block



5 zEnterprise CPCs

+




DS8700  
3

Storage configuration is based on IBM internal studies.  
Individual customer configuration will vary




# Consolidation Also Reduces Storage Costs

Storage Costs in a 10,000 Workload Environment



Storage numbers based on IBM study.  
Individual customer scenarios will vary.  
Prices are in US currency, prices will vary by country

# zEnterprise Is A Roadmap To The Data Center Of The Future



- Lower cost per unit of work for large scale workloads
- Revolutionary cost reductions for smaller scale workloads
- Data center simplification
- Improve quality of service
- No other platform can match!

**Mainframe workloads  
+  
distributed workloads  
best fit for cost**

