IBM System z Technology Summit

zEnterprise – The First System Of Systems

Virtualization & Consolidation On zEnterprise

David Rhoderick

25th May 2011

A Deeper Look At Some Topics

- Why virtualization matters?
- What is an ensemble?
- 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

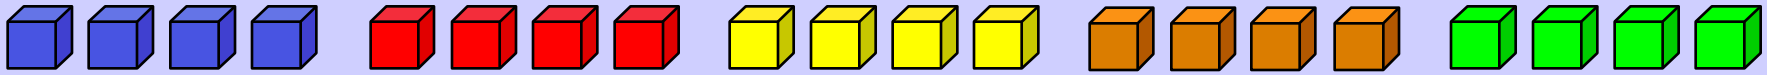

Virtual Resources

- 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.

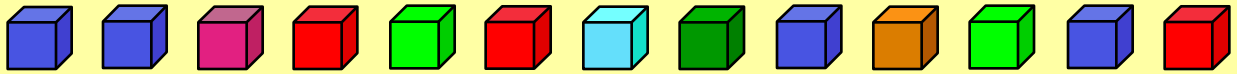

- **Separates presentation of resources to users from actual resources**
- **Aggregates pools of resources for allocation to users as virtual resources**

Server Virtualization Business Value

Roles

- Consolidation
- Dynamic provisioning
- Workload management
- Workload isolation
- Mixed production/test/OS
- Low cost backup servers

Benefits

- Higher resource utilization
- Greater usage flexibility
- Improved workload QoS
- Higher availability/security
- Lower management costs
- Investment protection

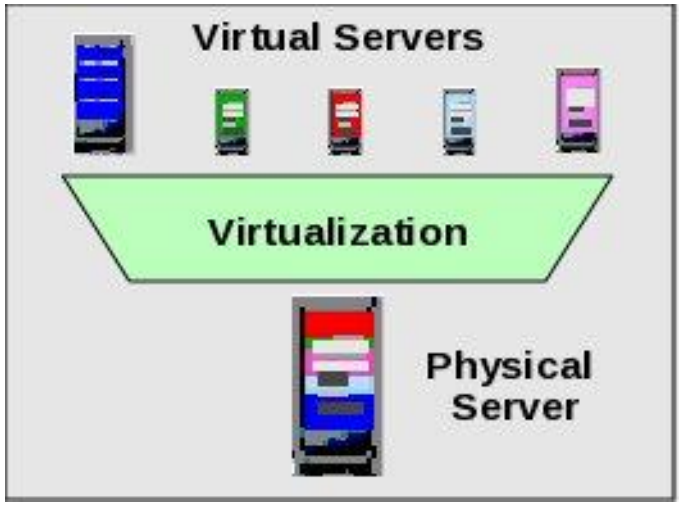

Server Virtualization Benefits

- 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

BUT, reliability (RAS) becomes even more essential

What Is A zEnterprise Node?

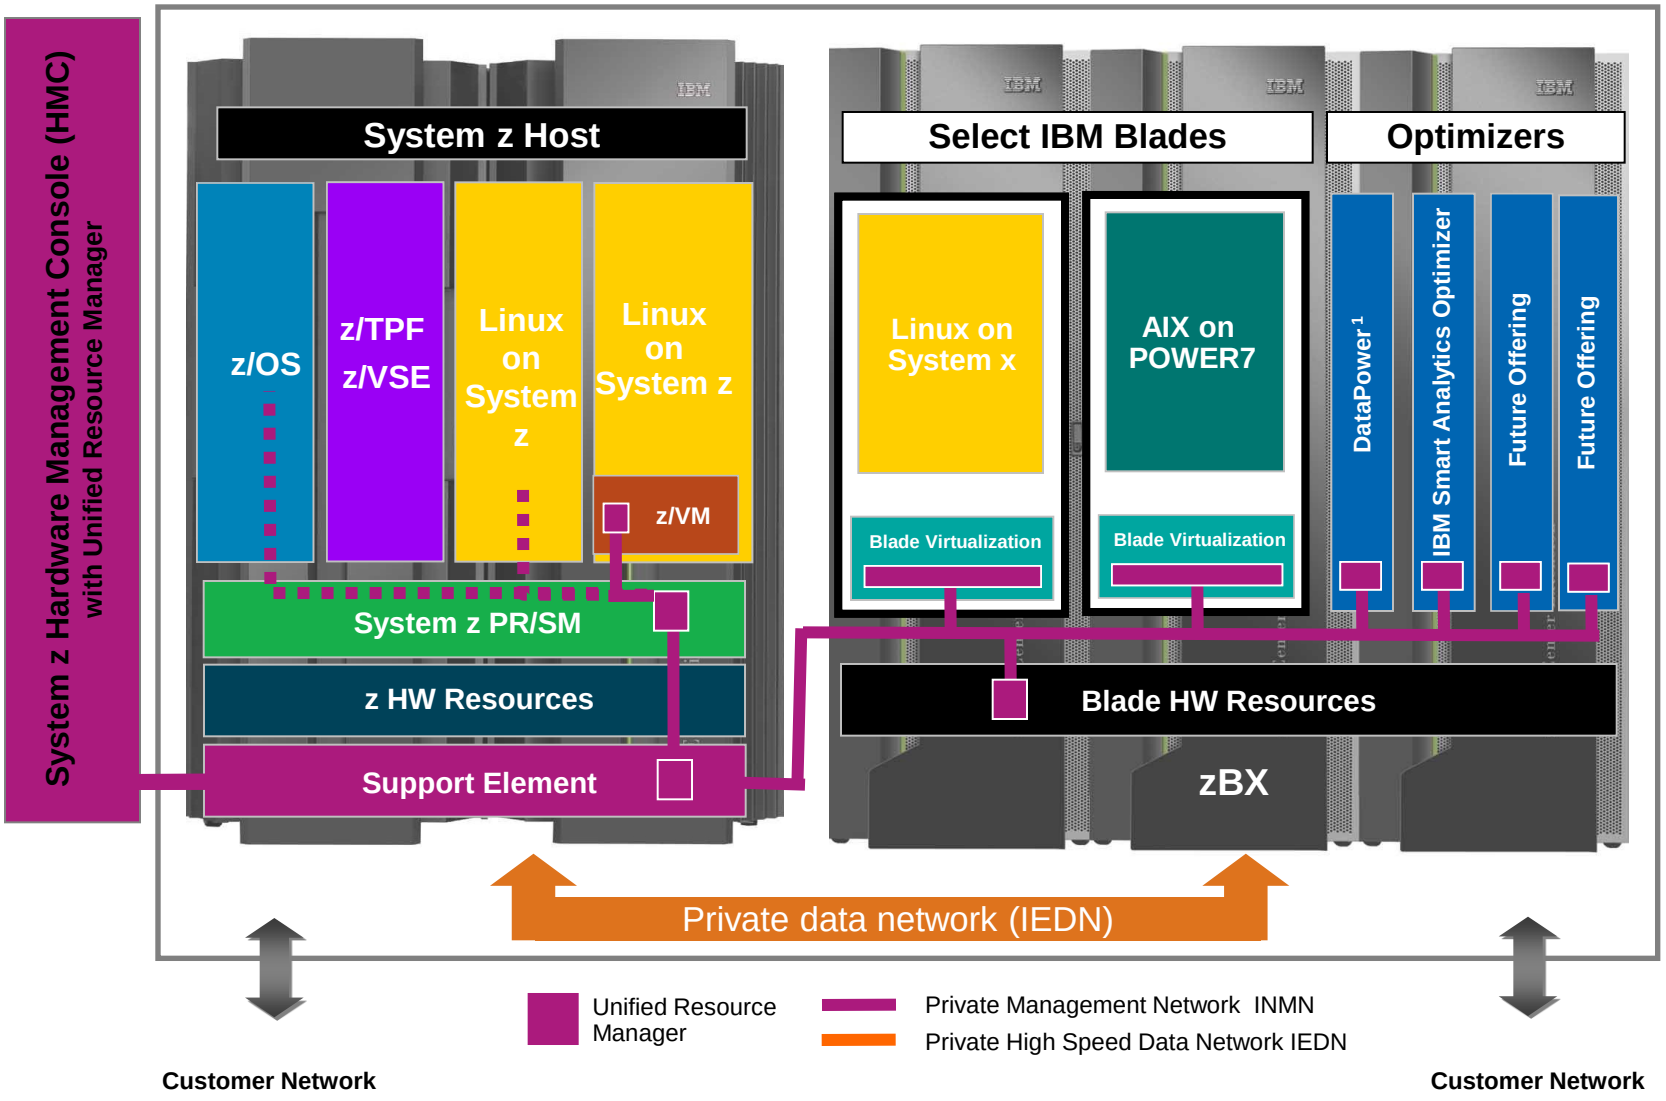

What's A zEnterprise Ensemble?

- ***A collection of up to 8 zEnterprise nodes that are managed collectively by the Unified Resource Manager as a single logical virtualized system***
- A zEnterprise node is a z196 CPC with 0 or 1 zBX
- The zBX contains 1 – 4 racks each containing 1 – 2 blade centers
- An ensemble can consist of a single z196 with no zBX attached, or two to eight z196s where at least one of the z196s has a zBX attached
- Blade-based fit-for-purpose solutions
- Unified Resource Manager manages the ensemble

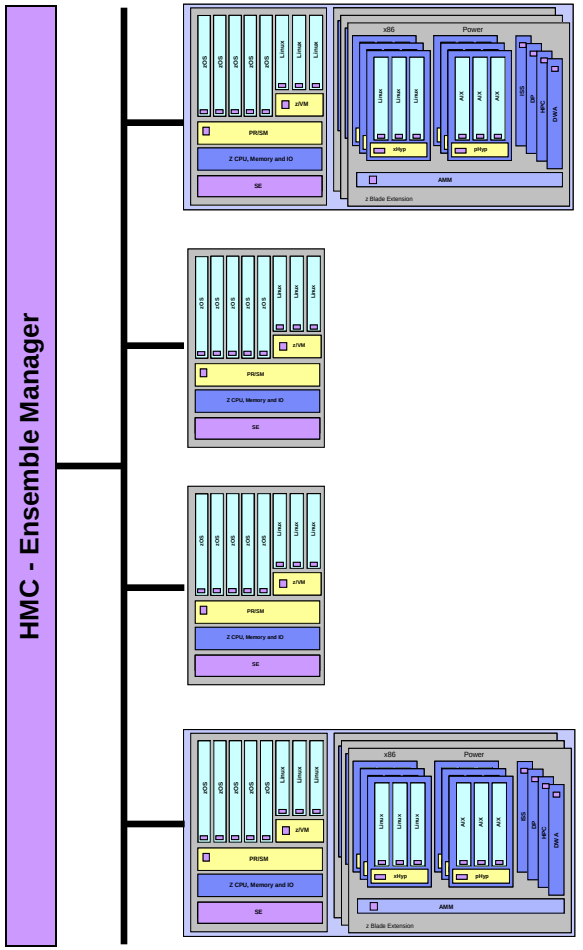

zEnterprise Extends Cost Advantages To A Broader Range Of Workloads

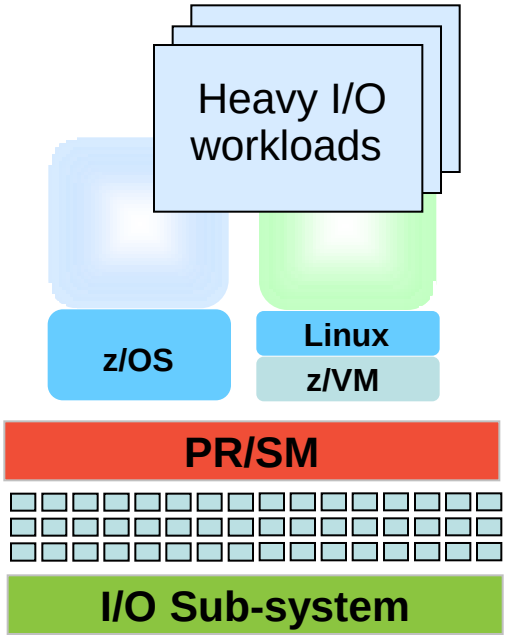

- Scale up to 96 cores in a frame (z/OS clusters with Sysplex)
- Dedicated I/O Sub-system with up to 336 I/O processors
- Superior qualities of service

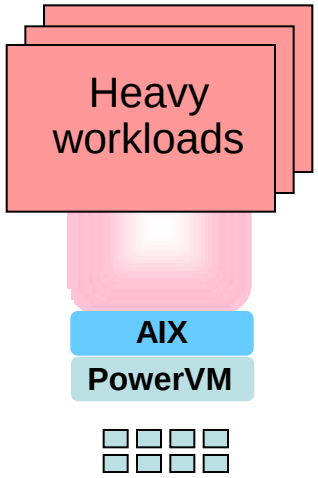

- Scales to 8 cores per blade
- Larger number of fast processing threads
- Floating point accelerators

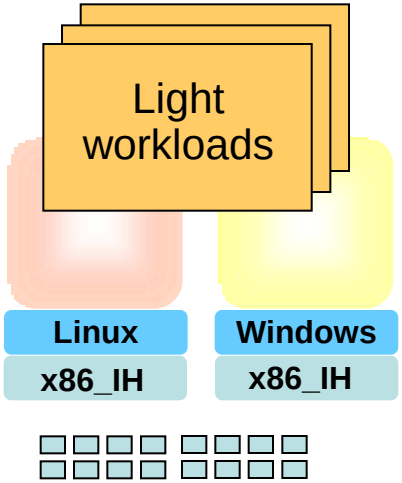

- Scales to 8-16 cores per blade
- Fast processing threads
- Commodity I/O
- Modest qualities of service

Consolidation Ratios For Distributed Workloads With Heavy I/O

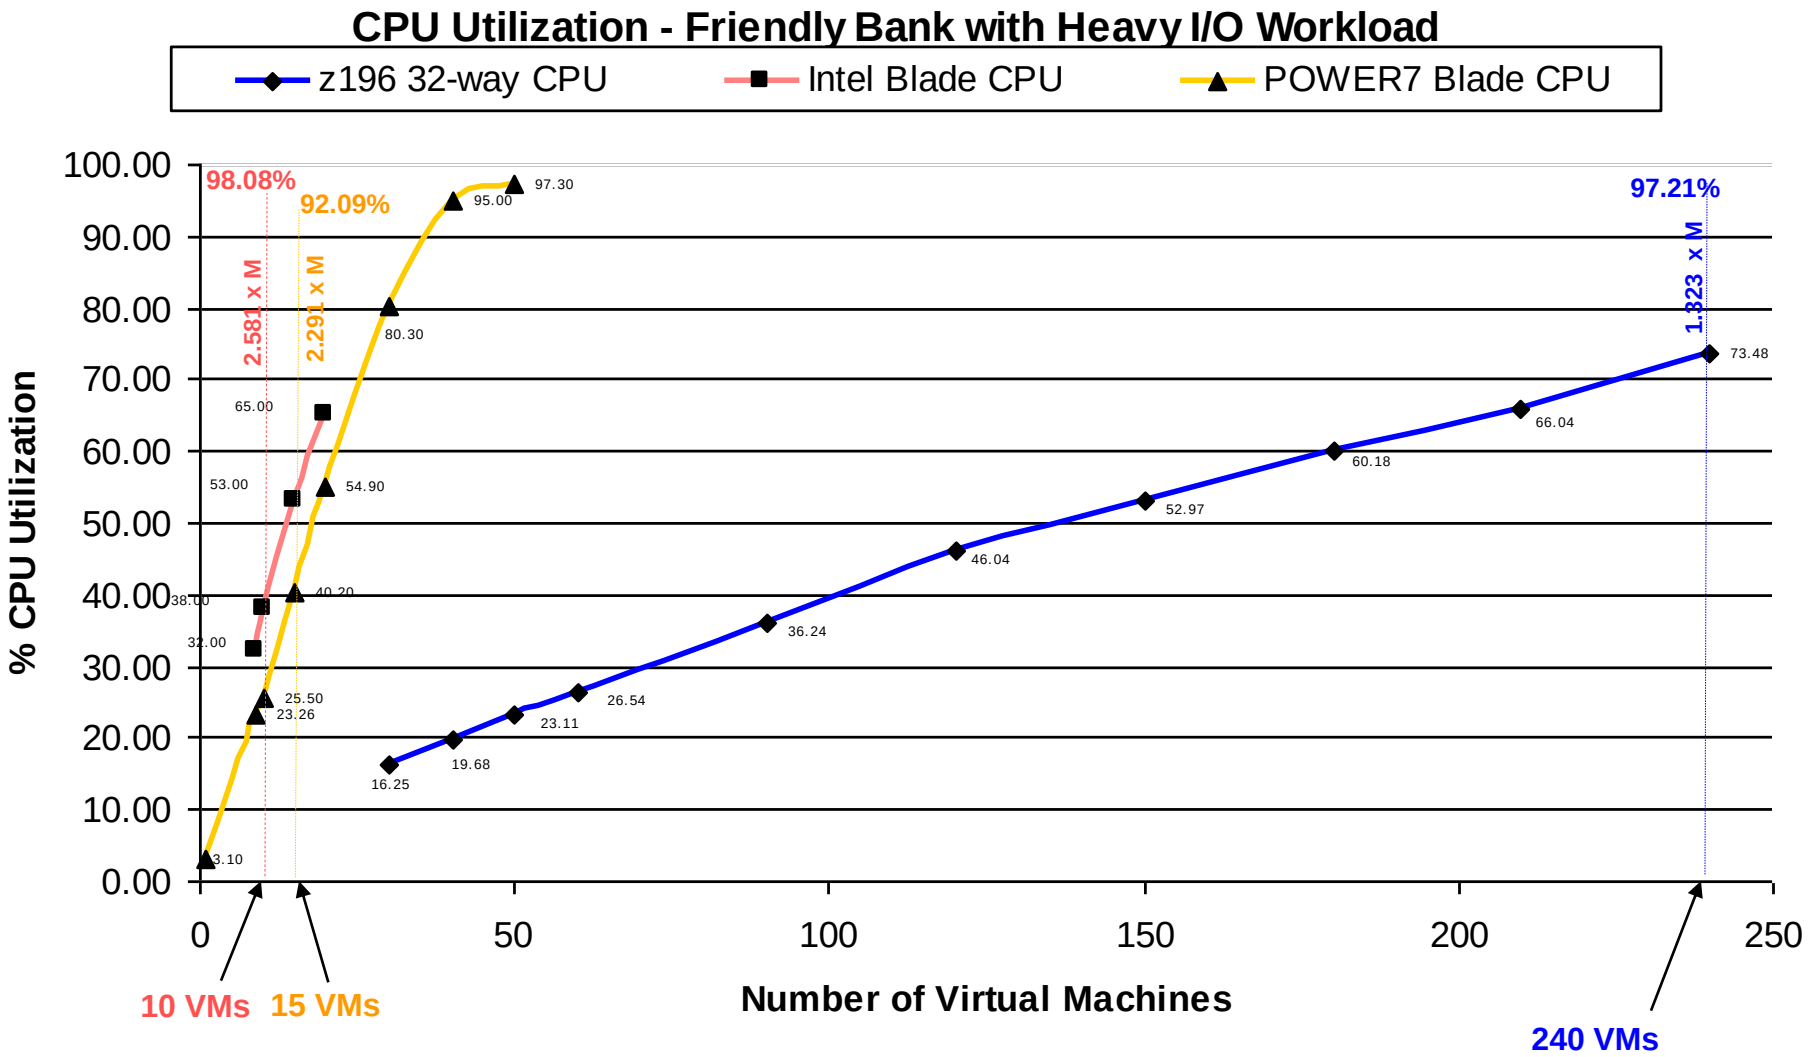

Consolidation ratios derived from IBM internal studies. z196 32-way performance projected from z196 8-way and z10 32-way measurements. zBX with x86 blades is a statement of direction only. Results may vary based on customer workload profiles/characteristics.

Deploying Heavy I/O Workloads

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

workloads

- **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 workload 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

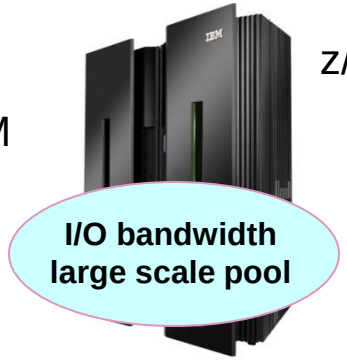

z/VM on zEnterprise CPC
32 IFLs
$14,052 per workload

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

workloads

- **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

more parallel threads

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

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

workloads

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

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

36 workload 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

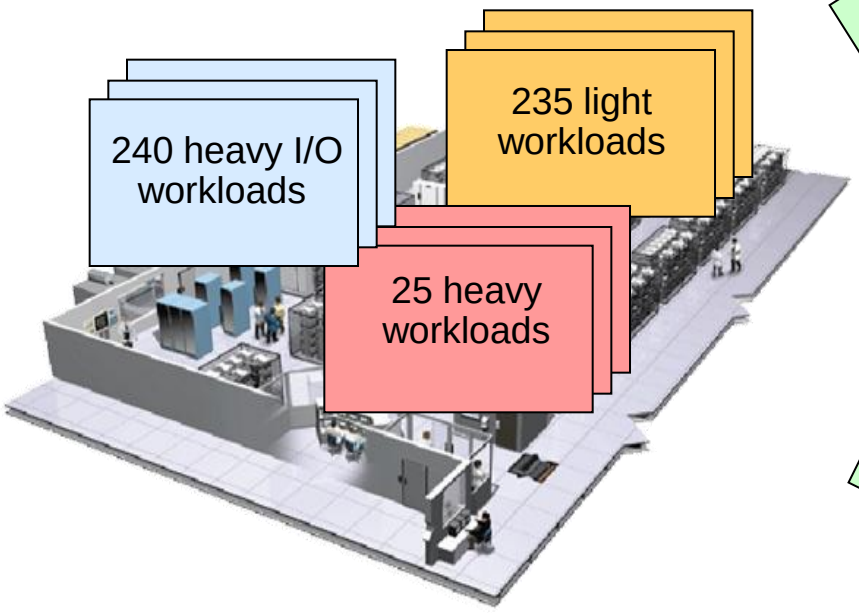

Deploy all distributed workloads on x blades

$11.6 M

Deploy all distributed workloads on p blades

$7.7 M

Deploy all distributed workloads on Linux on System z

$10.0 M

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

$7.5M

35% less

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

What Does It Cost To Deploy 500 Workloads On Virtualized Intel Servers?

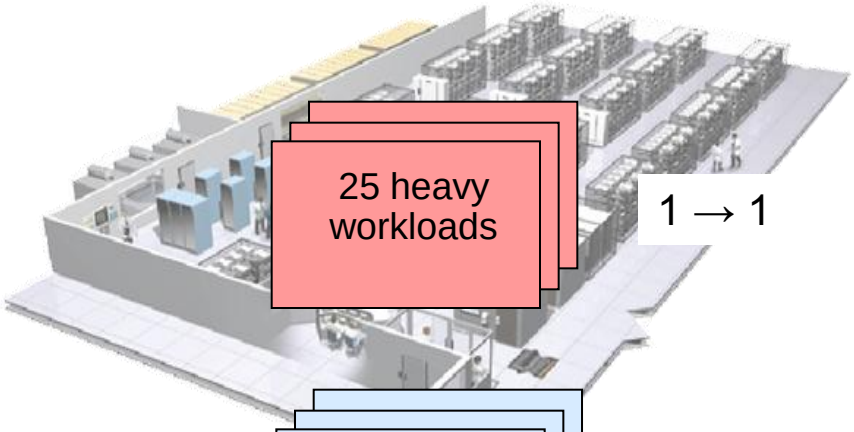

235 light workloads

8 → 1

Deployed on ***30*** Intel Xeon Servers using VMware
(8 cores each)

25 heavy workloads

1 → 1

Deployed on ***25*** Intel Nehalem Servers
(8 cores each, non-virtualized)

240 heavy I/O workloads

11 → 1

Deployed on ***22*** Intel Nehalem Servers using VMware
(8 cores each)

500 workloads

77 servers

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

What Does It Cost To Deploy 500 Workloads On zEnterprise?

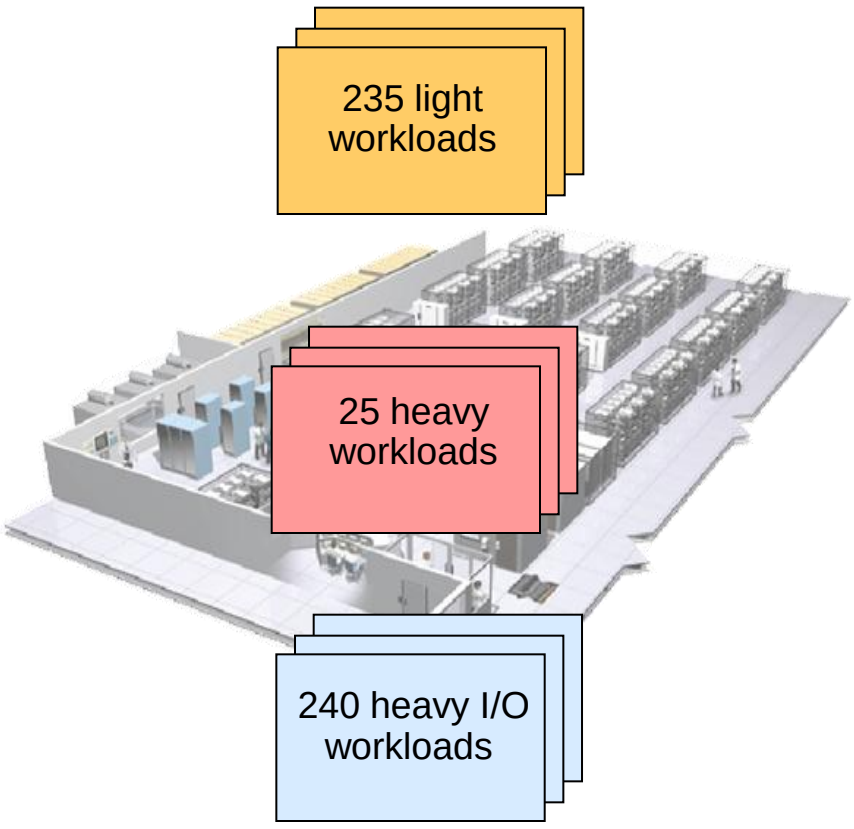

36 workloads per x blade

7
IBM x86 blades in 1 zBX

2 workloads per p blade

13
POWER7 Blades in same zBX

240 workloads per 32 IFLs

1
zEnterprise CPCs 32 IFLs

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

240 heavy I/O workloads

25 heavy workloads

235 light workloads

As deployed on Intel

77 Intel Servers

616 cores

$15.2M TCA (3 years)

Best fit on zEnterprise

2 Frames

192 cores

$7.5M TCA (3 years)

51%

less

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

Linux On z196 Achieves Lowest TCA For Heavy 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

. . . AND, much higher reliability (RAS)

z196 Is Designed For Large Scale Virtualization And Consolidation

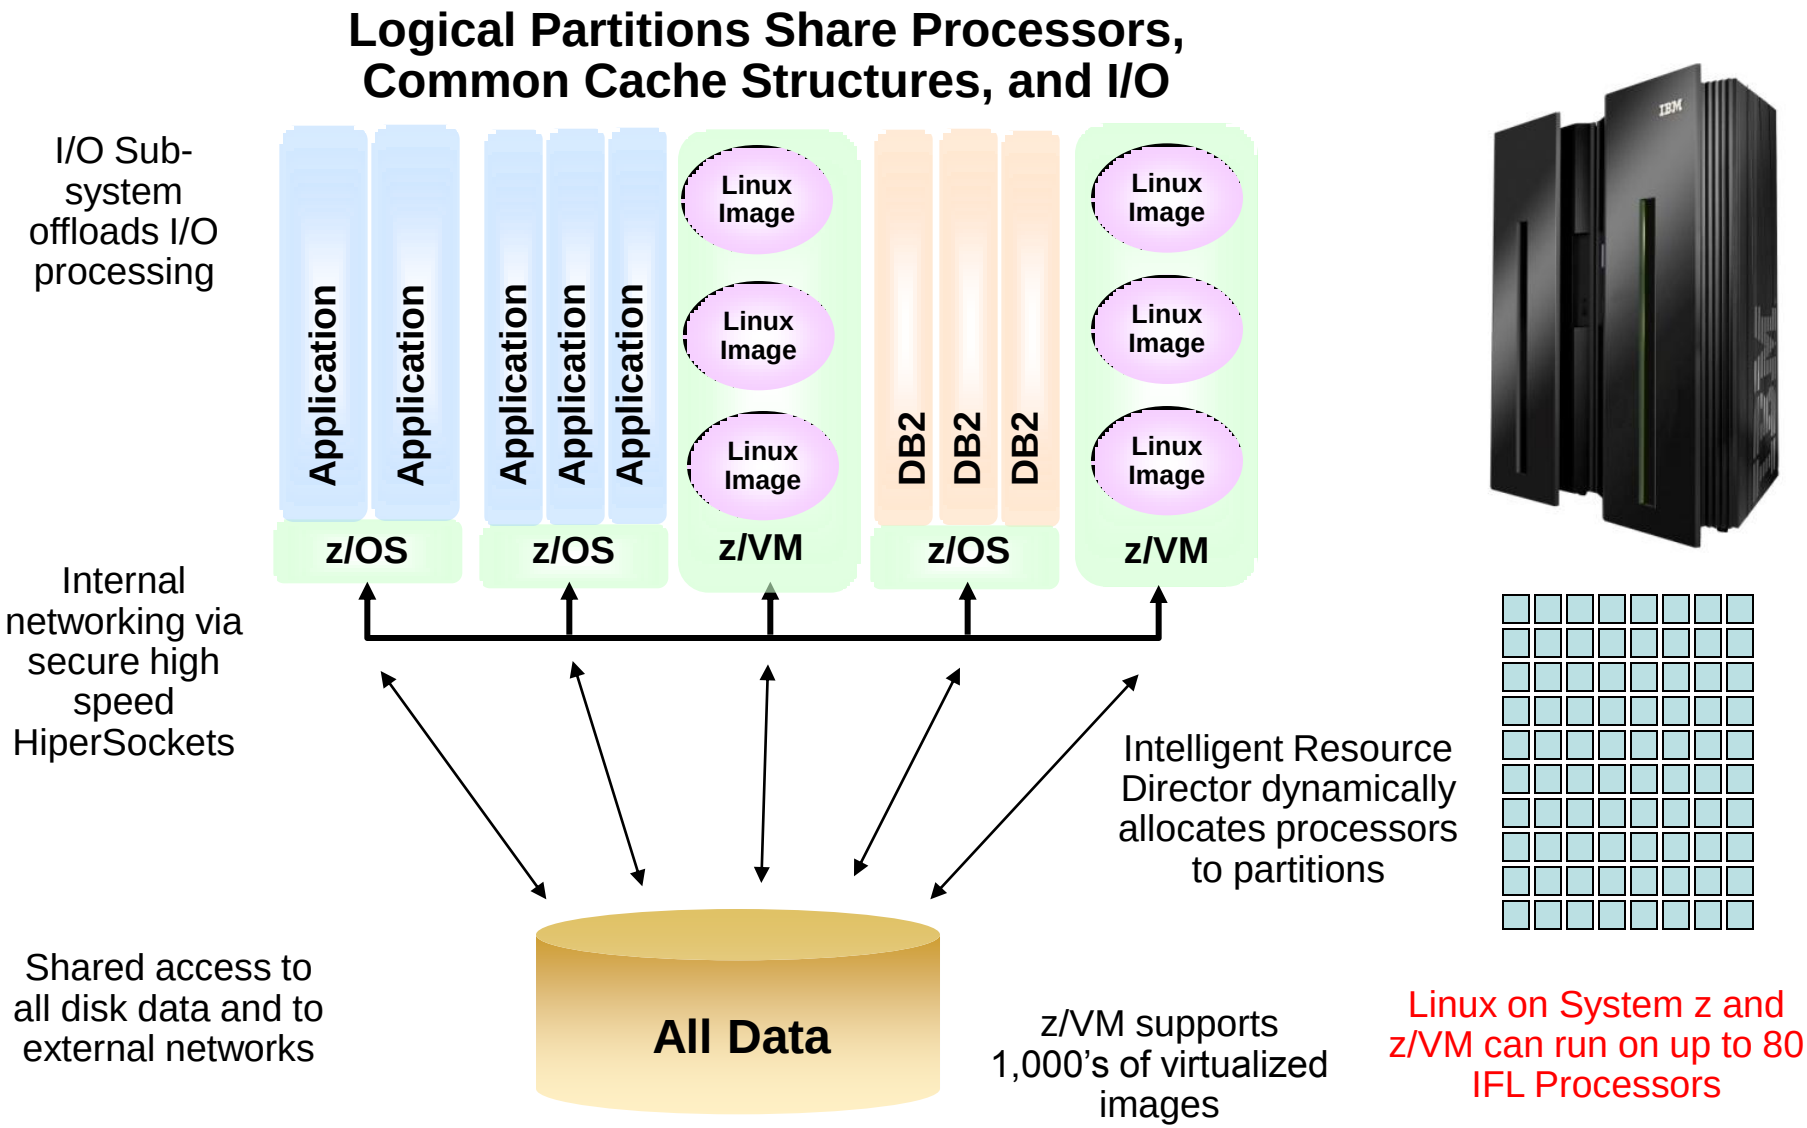

z/VM On System z – Optimized For Large Scale Virtualization

- 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 8-16 cores (currently)

System z Solution Editions For Linux

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

- System z Solution Edition for Enterprise Linux
  - IFLs, memory and z/VM added to an existing mainframe
  - Hardware and software maintenance for 3 or 5 years
- Enterprise Linux Server
  - zEnterprise server with IFLs, memory, I/O connectivity, and z/VM
  - Hardware and software maintenance for 3 or 5 years
- Linux on System z available from distribution partners
  - Novell SUSE and Red Hat

z196 - Optimized For High I/O Bandwidth

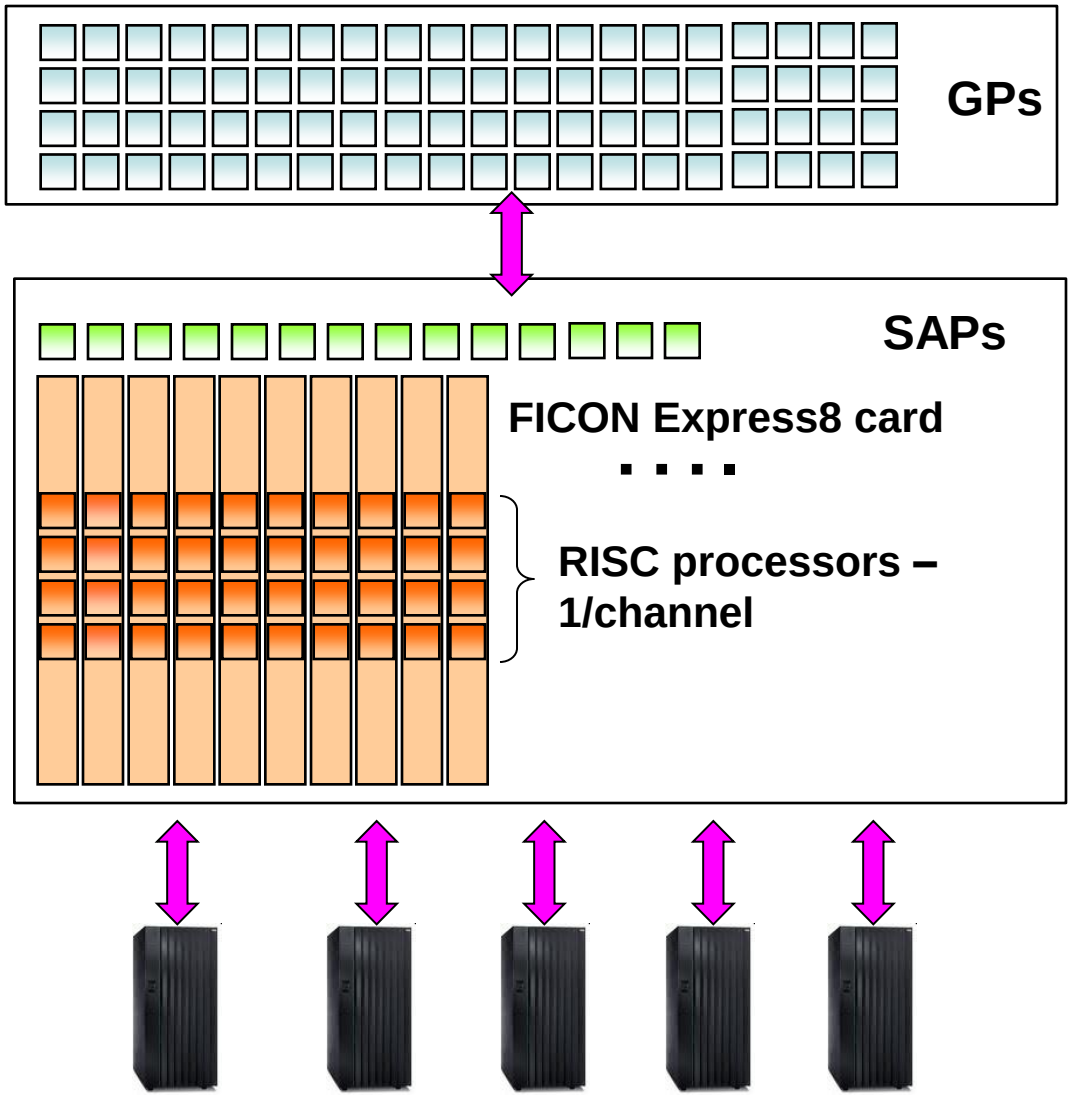

- Up to 80 General Purpose (GP) or Specialty Engine processors
  - Execute business logic
- Up to 14 System Assist Processors (SAP) to manage I/O requests
  - Can sustain up to **2.2M IOPS**[1]
- Logical Channel Sub-system virtualizes I/O
  - Up to 1024 logical channels
- Up to 84 physical FICON cards 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

[1]Recommend 70% max SAP Utilization – 1.5M IOPS

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 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

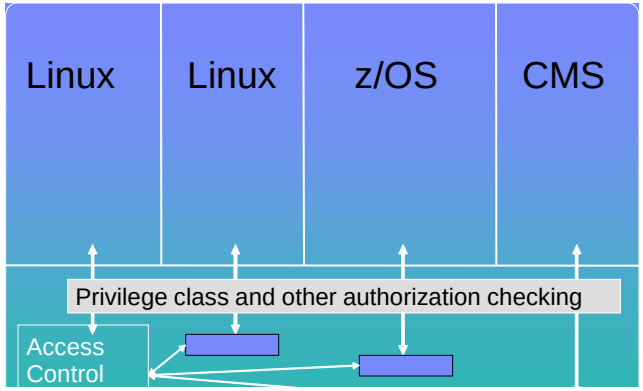

Linux On System z Workloads Inherit System z Qualities Of Service

- 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

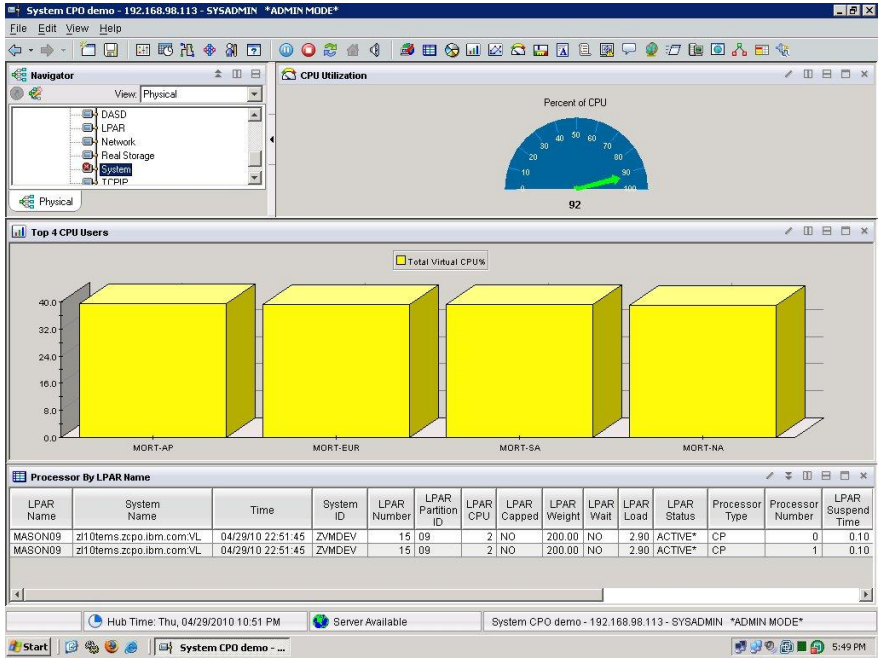

VMware can't recognize and take advantage of additional physical processors without bringing down and rebooting the system

Note: Assumes available processors on installed books

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

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

- Has limitations on those attributes important to large-scale consolidations
- x86 platforms lack the I/O subsystem to handle heavy I/O workloads
- Only supports one copy per physical server
- vCenter only allows monitoring, reporting and provisioning of VMs 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 with multiple platforms from one central hub
- zEnterprise Unified Resource Manager provides 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[1]

- 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

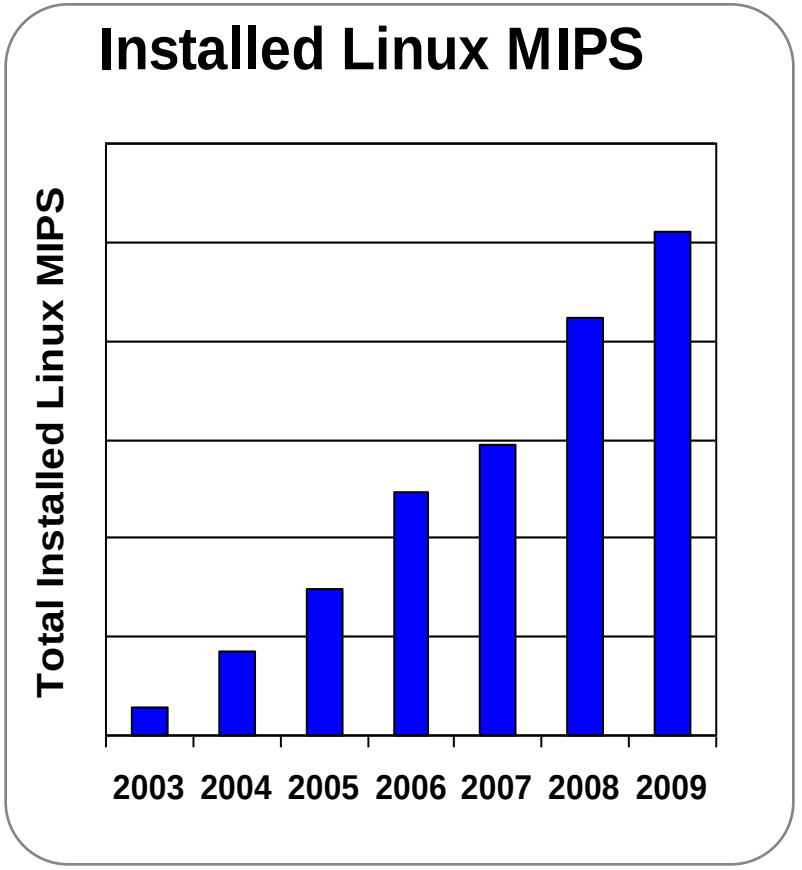

[1]Based on YE 2004 to YE 2009

Compare Network Cost Of Acquisition

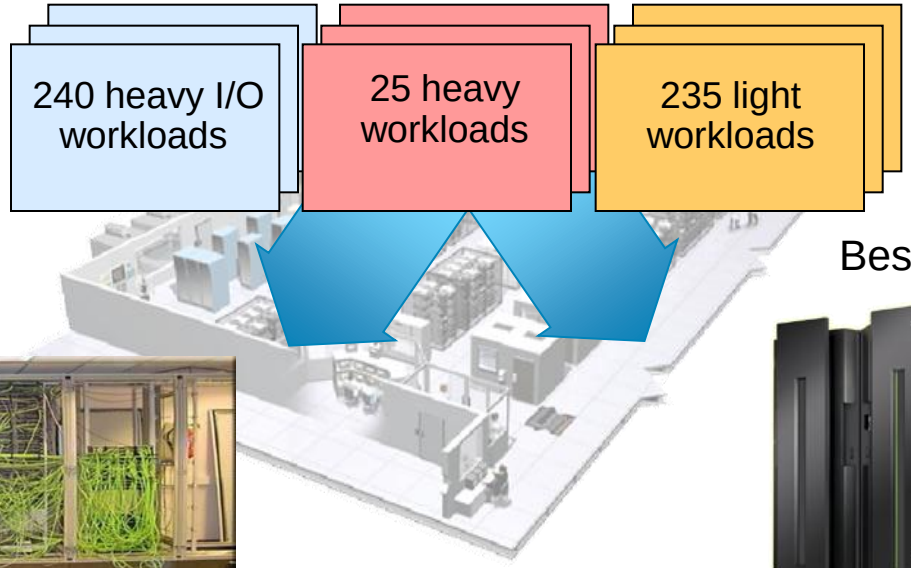

As deployed on Intel

Additional network parts

16 switches	340 cables	308 adapters

664 total network parts

$0.20M TCA

Best fit on zEnterprise

Additional network parts

1 switches	10 cables	10 adapters

21 total network parts

$0.03M TCA

86%
less

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

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

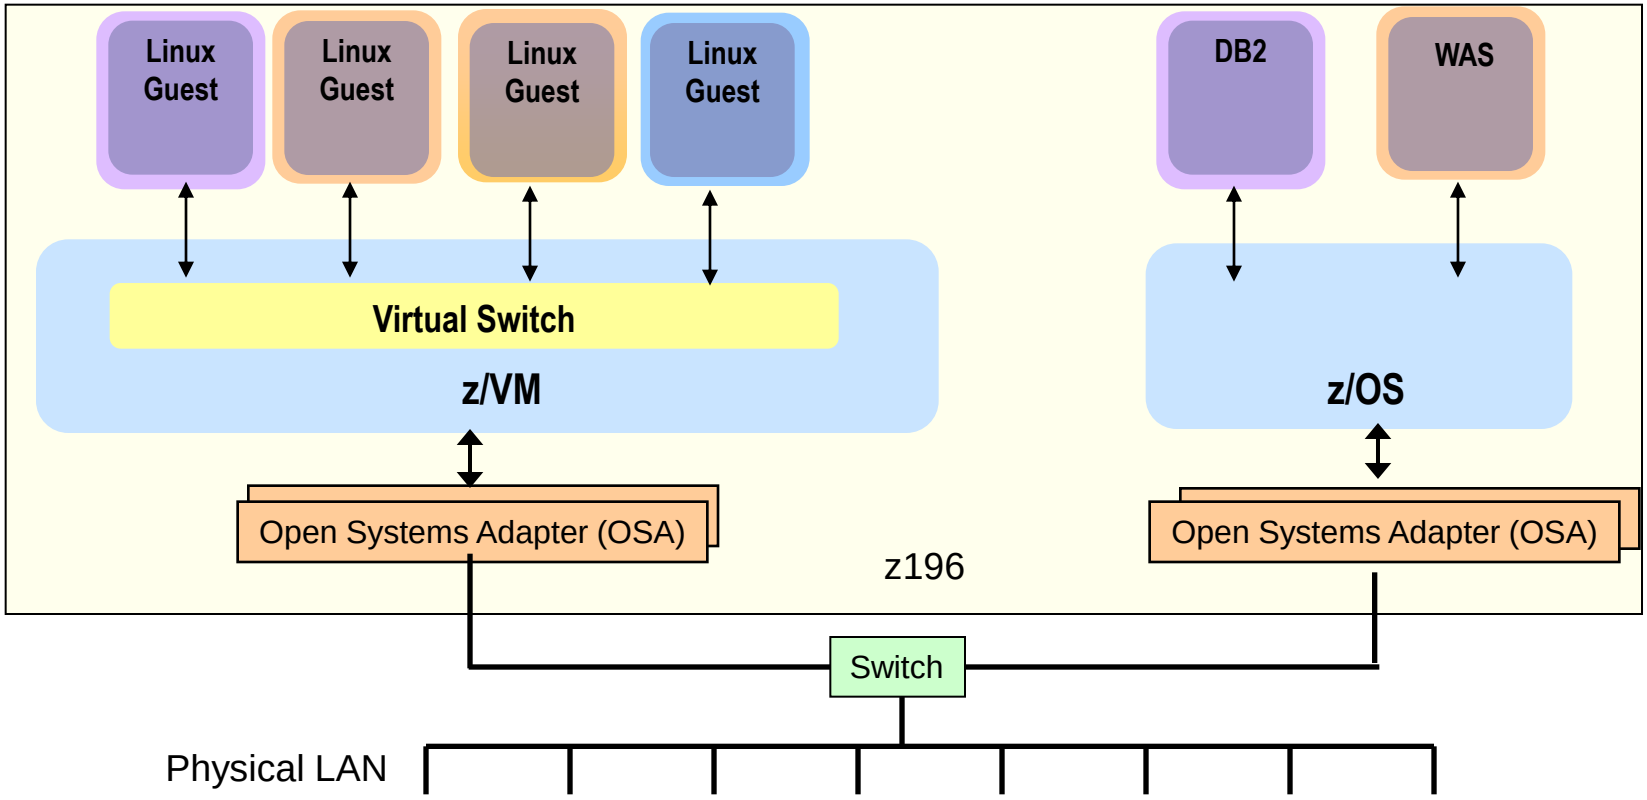

- Linux guests can talk to each other via zVM virtual switch – memory speed
- Linux guests can talk to outside world via z/VM virtual switch connected to shared OSA adapter
- Attach up to 8 physical OSA ports - redundancy, balancing
- Dynamically add new physical OSA to support Linux workload growth

System z Features Enable Network Simplification – HiperSockets

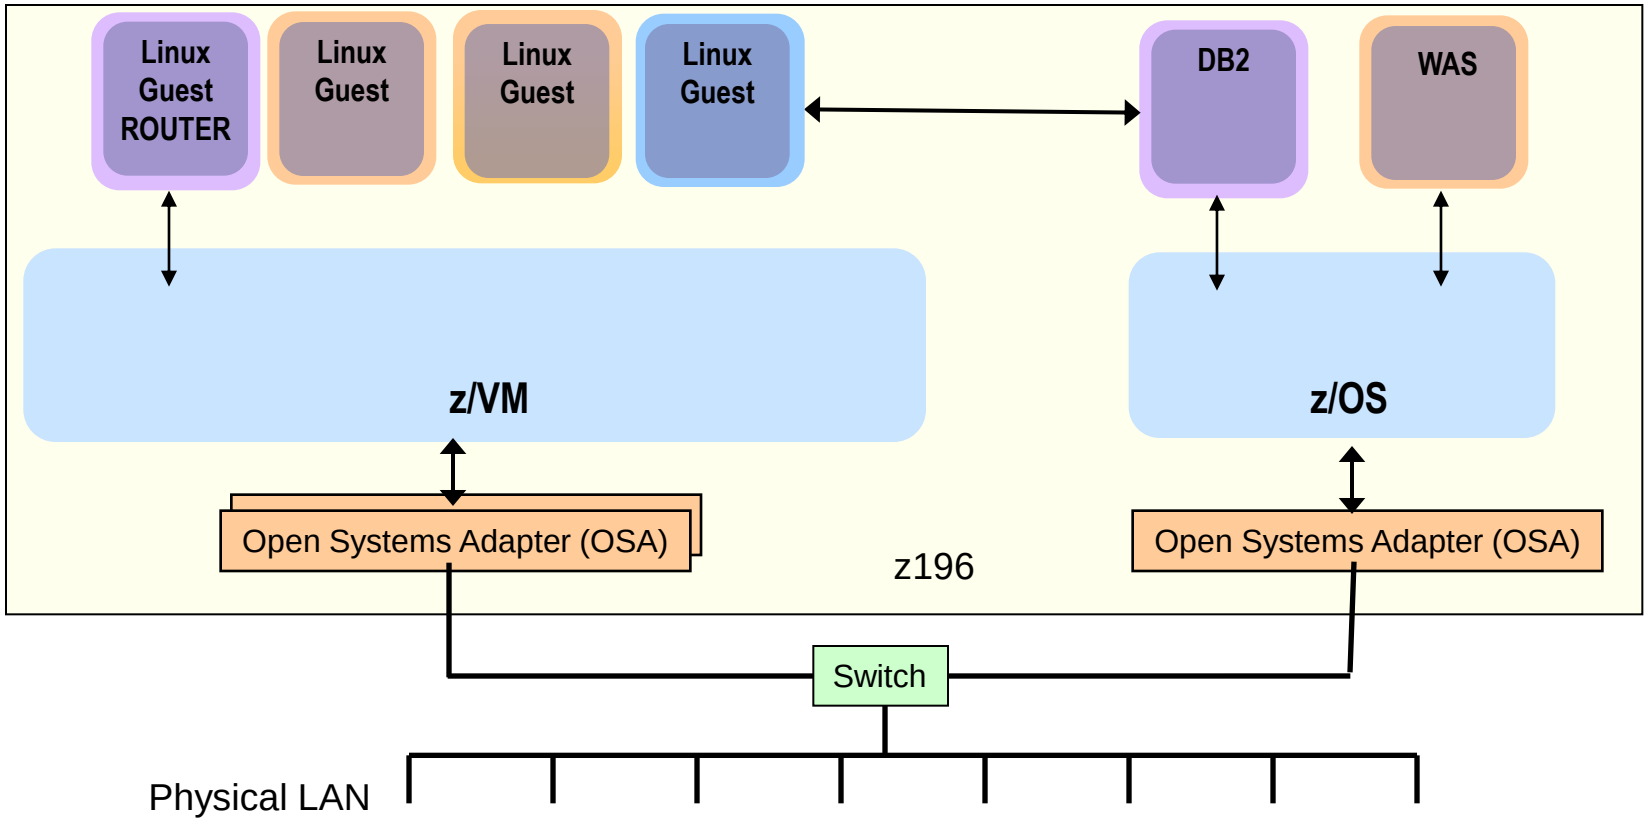

- Linux guests can talk to z/OS applications
- **Secure** IP communication at memory speed
- Close integration of data-intensive applications with database
- Reduces network management and physical assets

Network Simplification Extends To The zBX

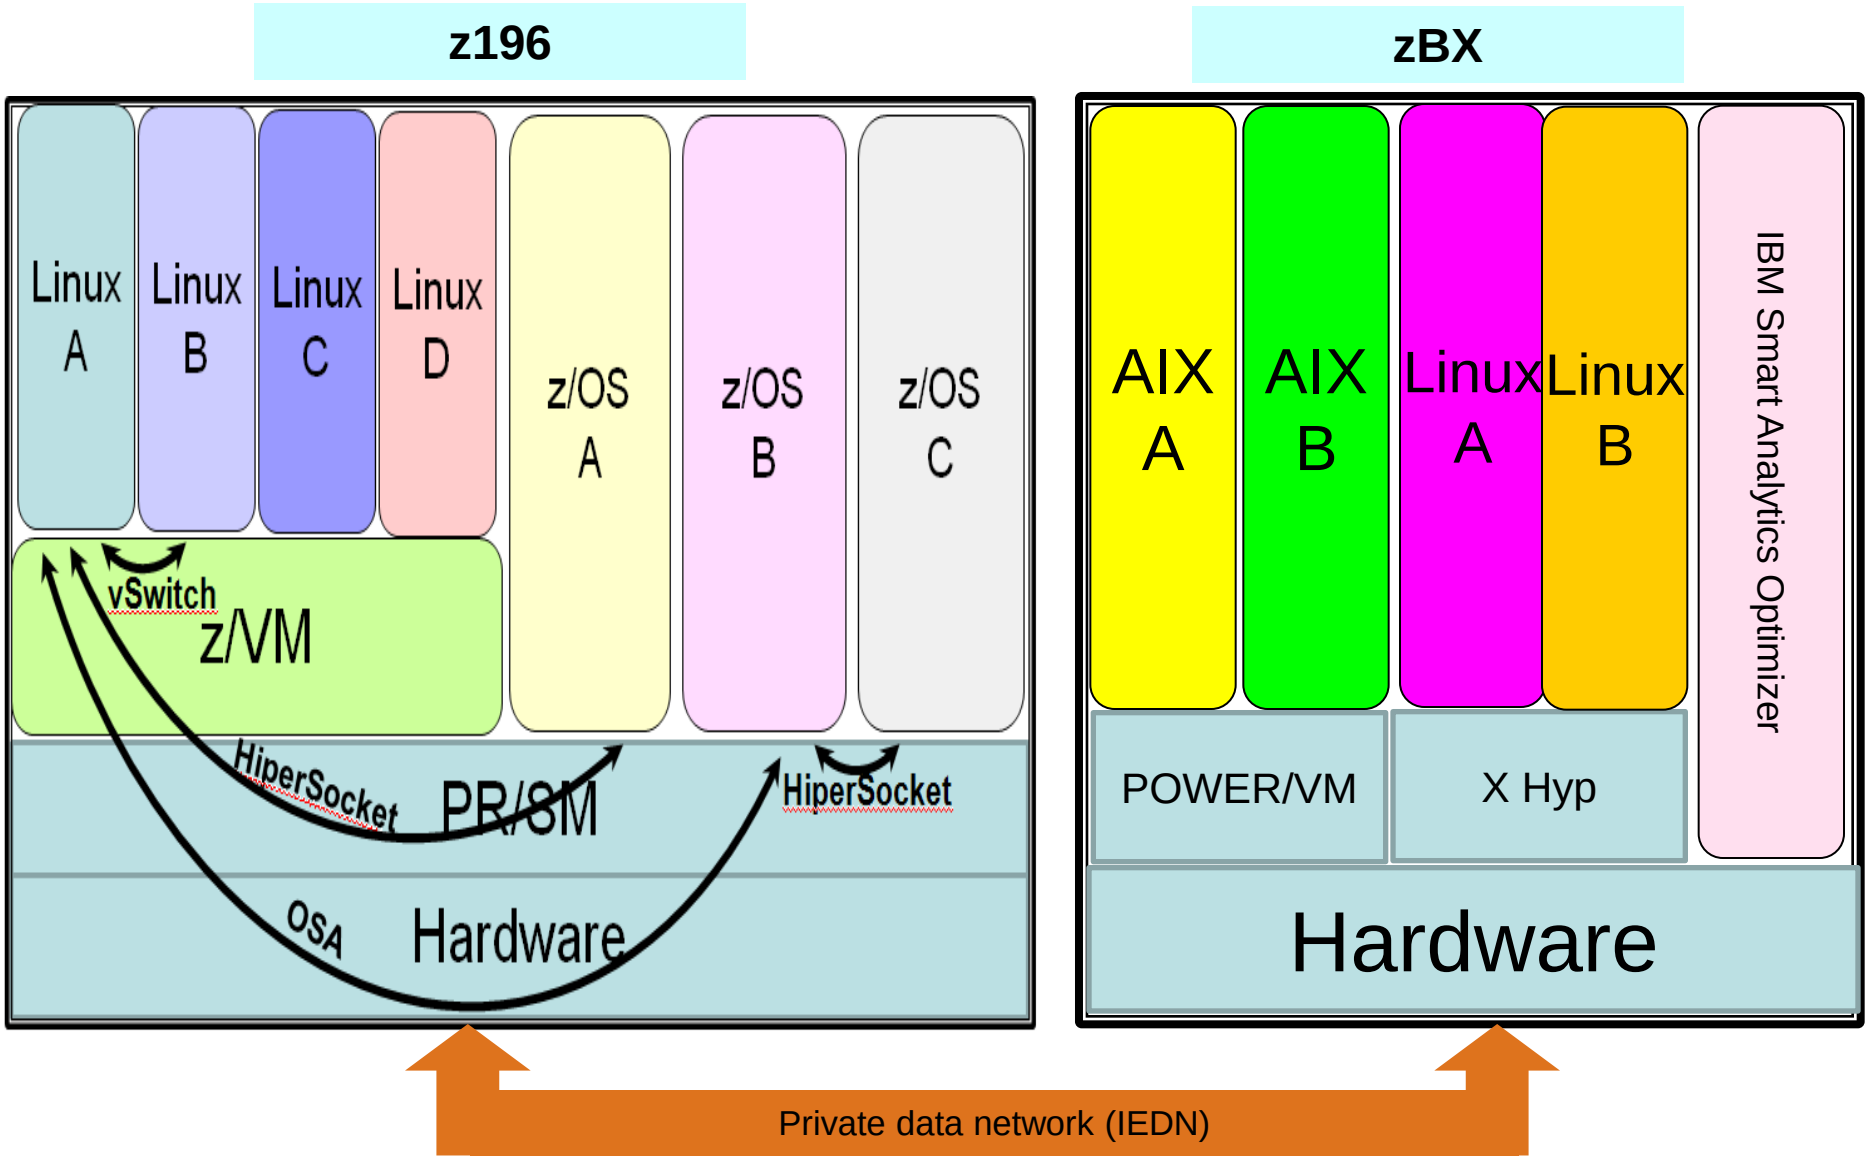

Compare Storage Cost

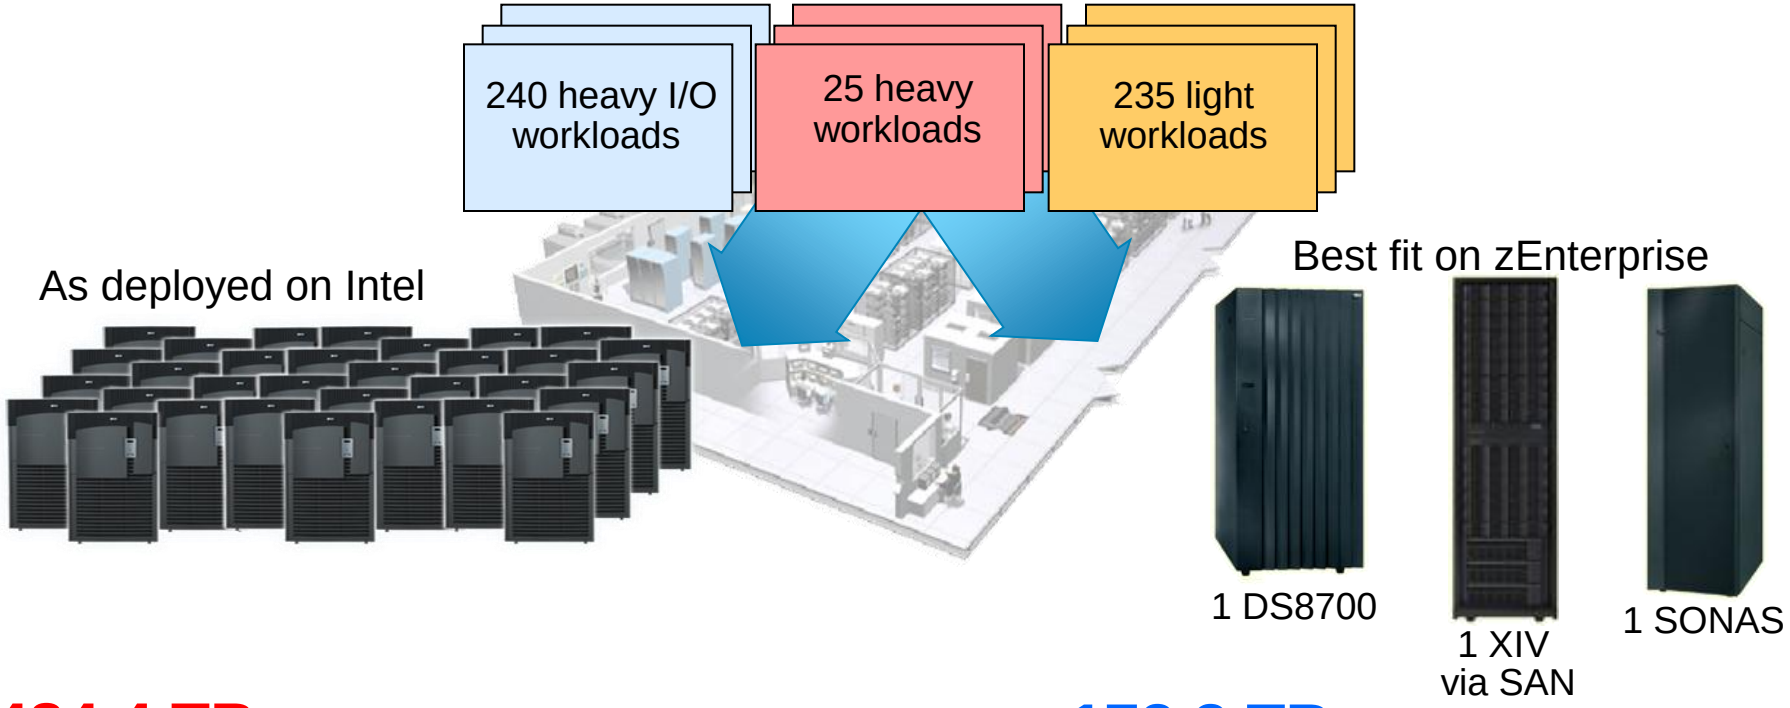

484.4 TB embedded storage
24% utilization
580 points of admin

$9.1M TCO (3 years)

240GB active storage required per workload

172.3 TB provisioned storage
67% utilization
3 points of admin

$6M TCO (3 years)

34%
less

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

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

235 workloads

32TB to XIV
26TB to SONAS

zBX rack with
7 x blades

+

XIV
1 via SAN

SONAS
1

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

25 workloads

3.3TB to XIV
2.7TB to SONAS

1 zBX rack
with 13 p blades

+

XIV
1 via SAN

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

240 workloads

58TB to DS8700

1 zEnterprise CPCs

+

DS8700
1

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

zEnterprise Is A Roadmap To The Data Center Of The Future

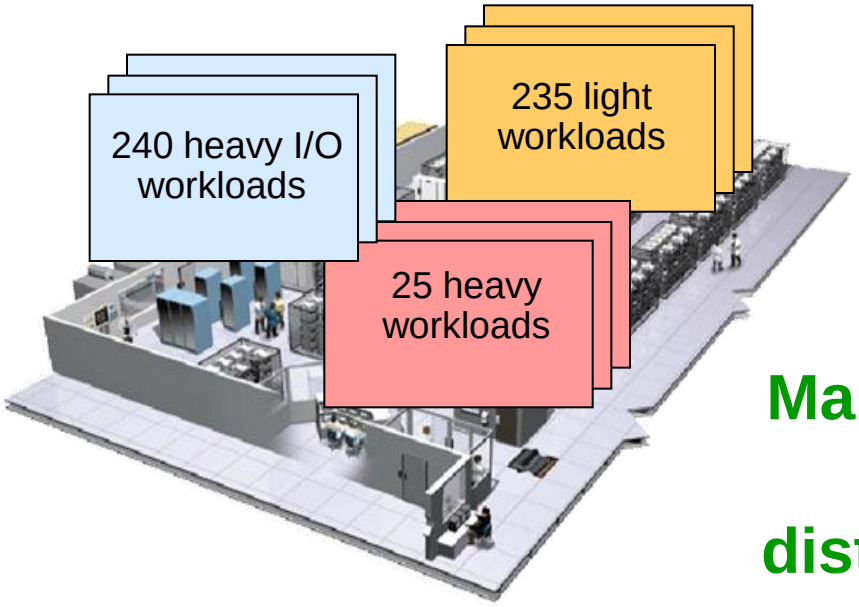

- Lower cost per unit of work for large scale workloads
- Revolutionary cost reductions for smaller scale workloads
- Data center simplification
- Improve quality of service
- No other platform can match!

Mainframe workloads
+
distributed workloads
best fit for cost