



# Analyzing IT Value and Cost Considerations - Maximizing the value of your mainframe

Ray Jones, Vice President,  
System z Software Sales,  
IBM Software Group



# Smarter Computing

Strategies to achieve breakthrough reductions in IT cost

Ascertain true elements of cost:


Hardware/Software/Maintenance  
Networking  
Energy  
Labor  
Storage

New metric  
for the age  
of Smarter  
Computing

**COST PER  
WORKLOAD**

# Smarter Computing With zEnterprise Delivers Breakthrough Economics

**Platforms Optimized For Different Workloads**




*Best fit for workload*

**Consistent Structured Management**



*Consistent structured practices*




***Lowest Cost Of Acquisition Per Workload***



***Lowest Cost Of Operation Per Workload***

**Lowest Cost Per Workload**


# A Closer Look At Fit-For-Purpose Workload Assignment



- Scale up to 80 cores in a frame (z/OS clusters with sysplex)
- Dedicated I/O sub-system
- Superior qualities of service

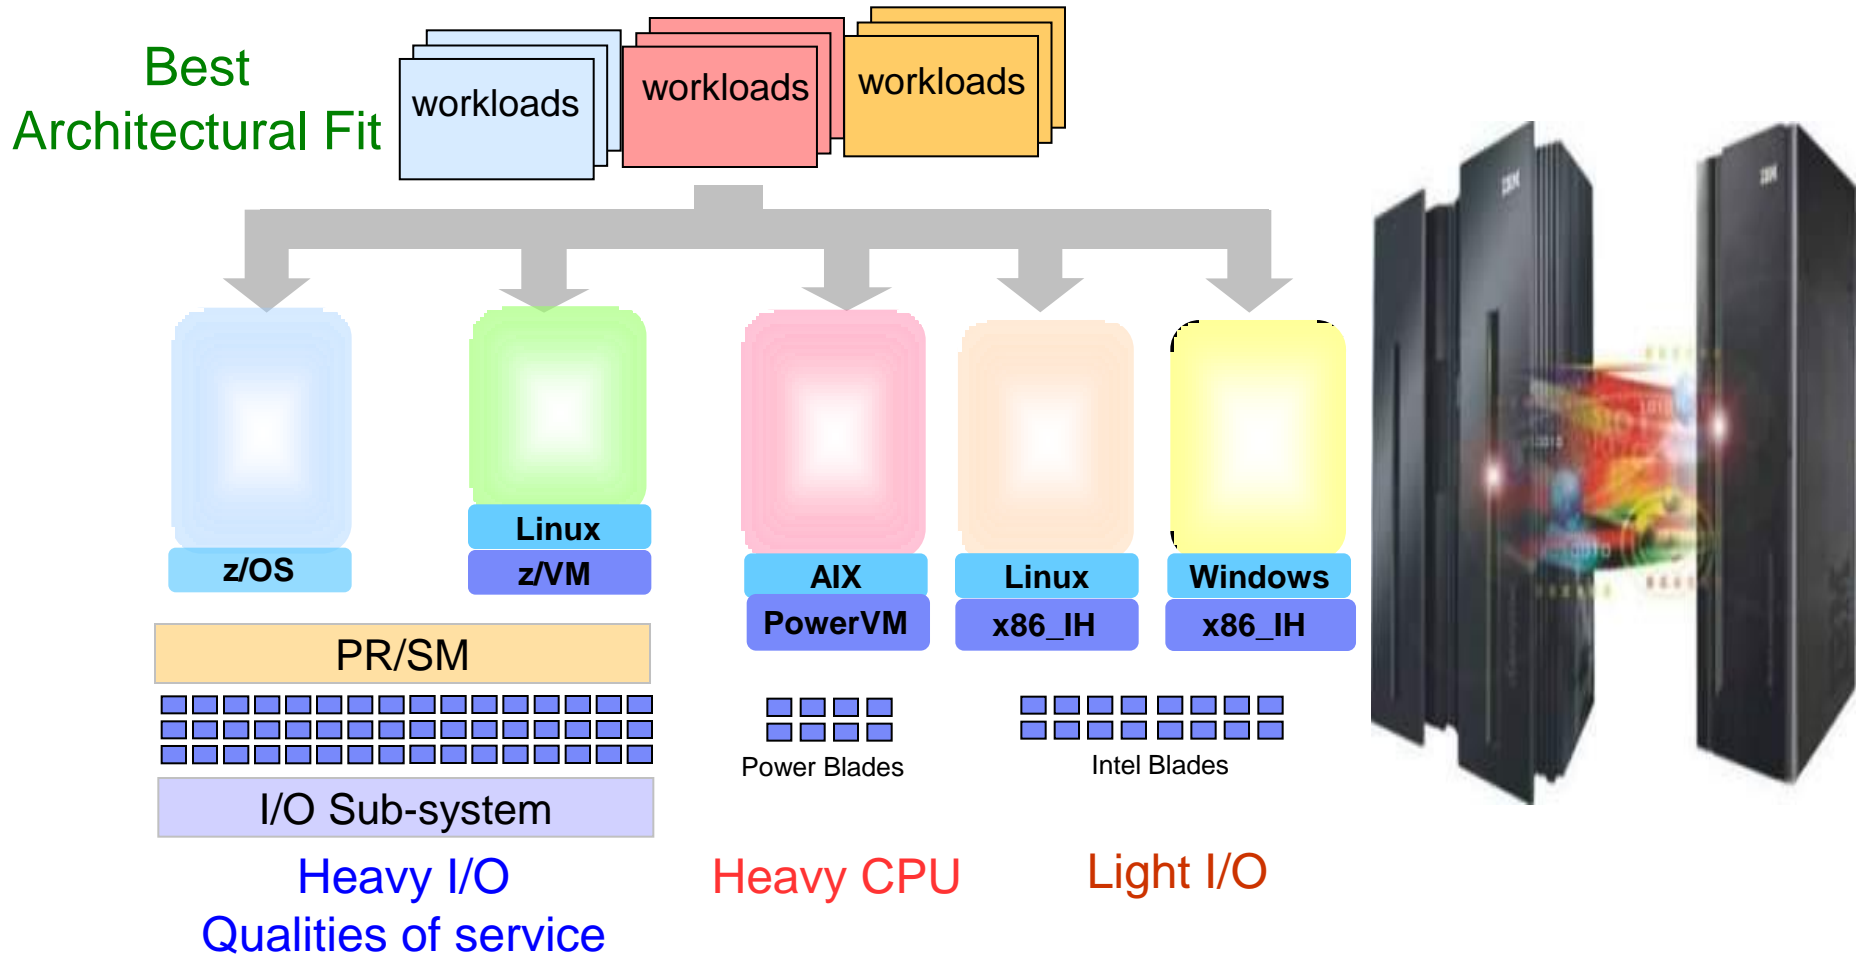


- Scales to 8 cores per blade
- 4 fast processing threads per core
- Floating point accelerators



- Scales to 16 cores per blade
- 2 fast processing threads per core
- Commodity I/O
- Modest qualities of service


# Workload Characteristics Influence The Best Fit Deployment Decision



Deploy or consolidate workloads on the environment best suited for each workload to yield lowest cost

# Deploying Stand Alone Workloads With Light I/O Requirements

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



- IBM WebSphere ND
- Monitoring software

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

2 workloads per Intel blade



Scale to 16 cores

Virtualized on Intel  
16 core HX5 Blade  
**\$190,023** per workload  
**Best Fit**

1 workload per POWER7 blade



PowerVM on PS701  
8 core POWER7 Blade  
**\$204,036** per workload

10 workloads per 32-way z/VM




z/VM on z196 CPC  
32 IFLs  
**\$339,939** per workload

Consolidation ratios derived from IBM internal studies. HX5 2.13GHz 2ch/16co performance projected from x3550 2.66GHz 2ch/12co 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.

# Deploying Stand Alone Workloads With Heavy CPU Requirements

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



- IBM WebSphere ND
- Monitoring software

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

47 workloads per Intel blade



Virtualized on Intel  
16 core HX5 Blade  
**\$8,086** per workload

28 workload per POWER7 blade



Fast low cost threads

PowerVM on PS701  
8 core POWER7 Blade  
**\$7,287** per workload  
**Best Fit**

155 workloads per 32-way z/VM




z/VM on z196 CPC  
32 IFLs  
**\$21,932** per workload

Consolidation ratios derived from IBM internal studies. HX5 2.13GHz 2ch/16co performance projected from x3550 2.66GHz 2ch/12co 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.

# Deploying Stand Alone Workloads With Heavy I/O Requirements

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



- IBM WebSphere ND
- Monitoring software

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

1 workload per Intel blade




Virtualized on Intel  
16 core HX5 Blade  
**\$380,046** per workload

1 workload per POWER7 blade



PowerVM on PS701  
8 core POWER7 Blade  
**\$204,036** per workload

40 workloads per 32-way z/VM



**I/O bandwidth large scale pool**


z/VM on z196 CPC  
32 IFLs  
**\$84,985** per workload  
**Best Fit**

Consolidation ratios derived from IBM internal studies. HX5 2.13GHz 2ch/16co performance projected from x3550 2.66GHz 2ch/12co 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.




# Core Proliferation for a Mid-sized Offload Project

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



**\$25.4M TCO (5yr)**

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



**6 processors**  
 (1,660 MIPS)



**176 distributed processors**  
 (800,072 Performance units)

**\$17.9M TCO (5yr)**


**482 Performance Units  
 per MIPS**

# Optimized For High I/O Bandwidth – Reduce Batch Window By 83%


**zEnterprise + DS8300**

**Power Blade 701 + DS8300**

8 processors  
128GB memory  
16 IO channels



8 processors  
128GB memory  
2 IO channels



*300 jobs each to sort 3GB file*

Sorting Total Elapsed	1229 Seconds
Concurrency	20
Bytes Per Sec	<b>1600MB</b>

Sorting Total Elapsed	6900 Seconds
Concurrency	20
Bytes Per Sec	<b>280MB</b>

*10 jobs each to merge 30 sorted files into 90GB master file*

Merging Total Elapsed	1422 Seconds
Concurrency	10
Bytes Per Sec	<b>1350MB</b>

Merging Total Elapsed	7920 Seconds
Concurrency	10
Bytes Per Sec	<b>244MB</b>

# Utilization of Distributed Servers & Storage


Typical utilization of:

Windows Servers	5-10%
UNIX Servers	10-20%
System z Servers	85-100%



Server dedicated to one application

The cost of storage is typically three times more in distributed environments



▪ **Storage Allocation**


- Application-specific resulting in over-allocations
- Fine grained storage allocation mechanisms characteristic of mainframe storage are uncommon in distributed environments.

▪ **Storage Utilization**

- Single digit utilization for distributed environments is not uncommon
- Storage utilization of 80% + is typical for mainframe

▪ **Storage Management**

- Data disaster recovery, synchronization, and transfer requirements add complexity and cost



# What Is A Typical Value Of Sigma?

## IBM Survey Of Workload Variability In 3200 Servers

Type Of Workload	Average Utilization	Peak Utilization	Sigma
<b>Infrastructure</b>	<b>6%</b>	<b>35%</b>	<b>2.5 * Mean</b>
<b>Web Server</b>	<b>4%</b>	<b>24%</b>	<b>2.5 * Mean</b>
<b>Application</b>	<b>4%</b>	<b>34%</b>	<b>3.75 * Mean</b>
<b>Database</b>	<b>5%</b>	<b>37%</b>	<b>3.25 * Mean</b>
<b>Terminal</b>	<b>6%</b>	<b>45%</b>	<b>3.25 * Mean</b>
<b>E-Mail</b>	<b>4%</b>	<b>34%</b>	<b>3.75 * Mean</b>

IBM System x™ Servers and VMware Virtual Machine Sizing Guide

Legacy workloads on XEON 2.5-2.8GHz Servers

Normal probability distribution



# New Workload Scenarios – Beware Benchmarks

## ▪ Stress test benchmarks have no variability!

- They drive the system under test to 100% utilization with no variation
- Comparing mean throughputs at 100% utilization doesn't give a realistic view of the resources required for deployment

Running a new workload with variability  $\text{Sigma}=2.5*\text{Mean}$  requires processing capacity equal to **6 times the Mean** workload demand


Adding a new workload to a pool of 256 existing workloads will require incremental processing capacity equal\* to the **Mean** workload demand



# Case Study: Network Costs –Before Consolidation (483 Servers to 2 System z’s)

High Utilization Switch Module	14
Low Utilization Switch Module	12
Switch Interconnect Module	6
50 Ft UTP Cable	966
10GB Eth Fiber Cable	12
Switch Chassis	3

Hardware Acquisition \$748K  
 Network Annual Costs \$597K




Shows 30 of the 483 Servers

## Case Study: Network Costs – After Consolidation (483 Servers to 2 System z's)

New Hardware Acquisition \$0  
(reuse some of old network hardware)

“After” Network Annual Cost \$253K

Network Annual Cost Savings \$344K



# Why Does Core Proliferation Happen?

## ▪ De-consolidation of applications to dedicated servers

- Dedicated servers for functional roles - application, database, security, batch, systems management
- Separate servers for production, development, quality assurance test
- Low utilization due to provisioning for the peak on each server and pre-provisioning for growth

## ▪ Disaster Recovery

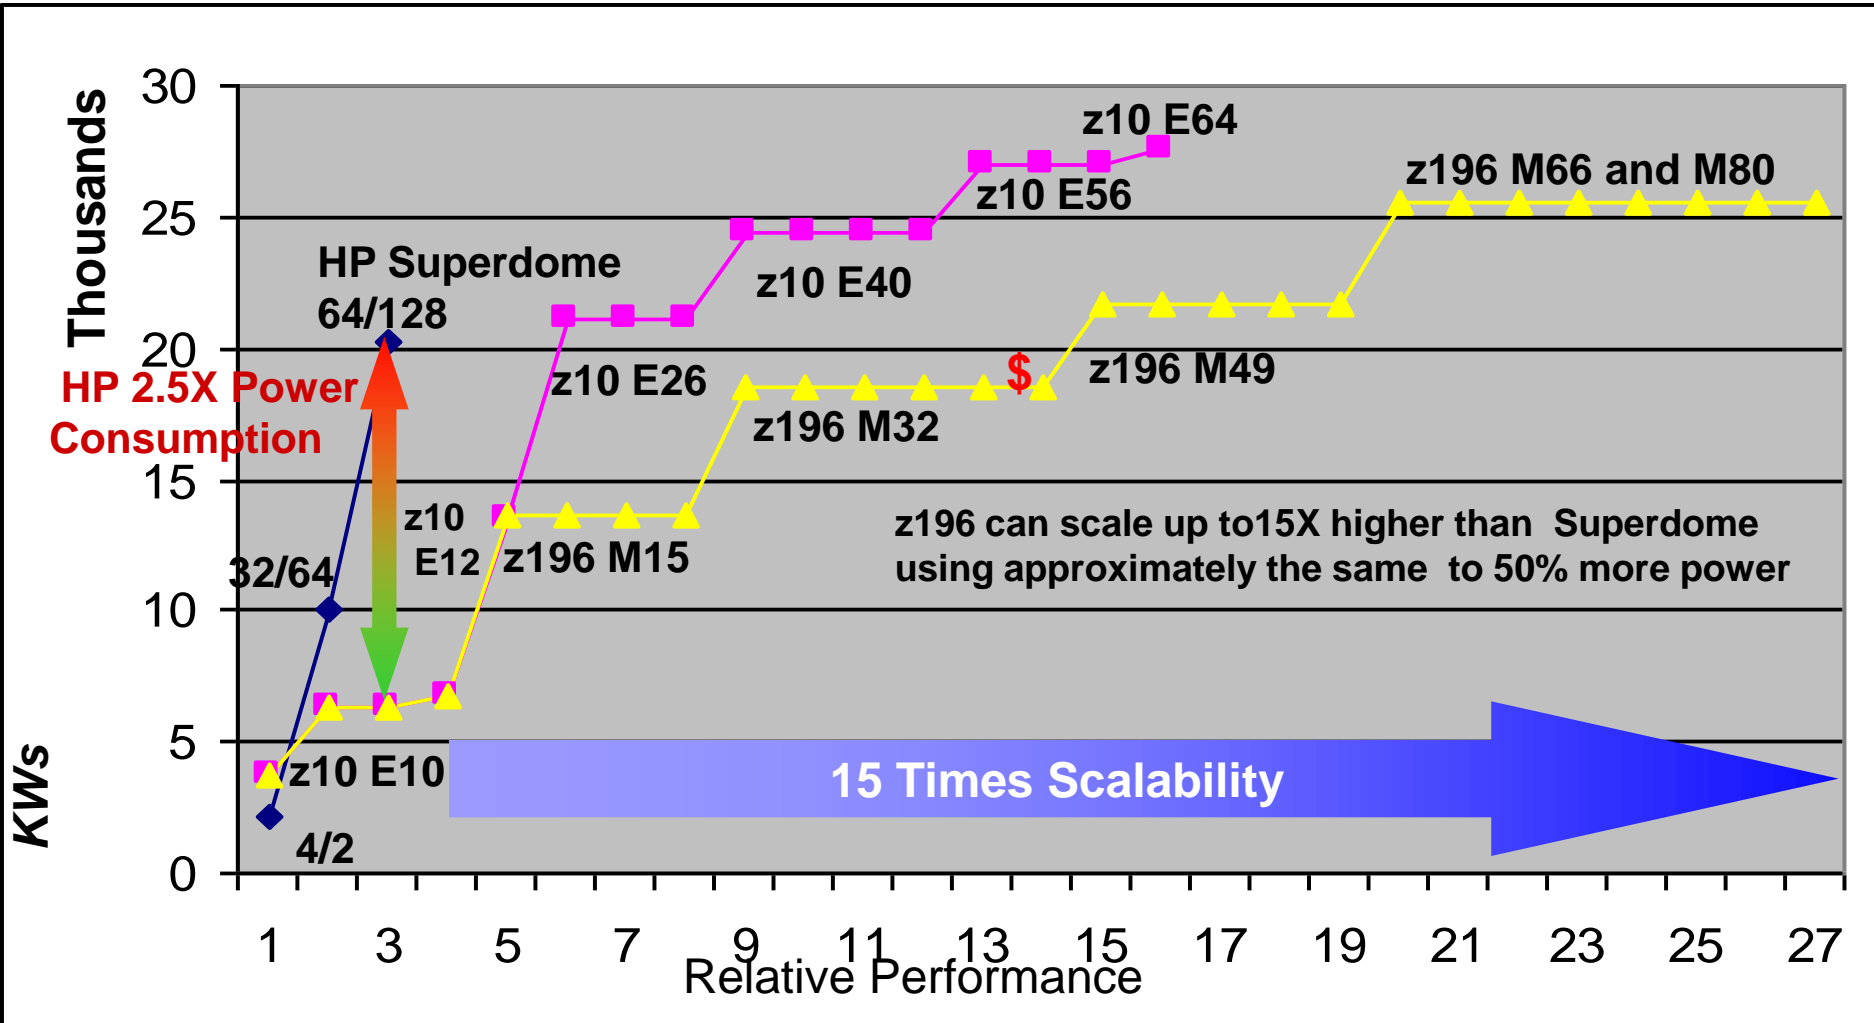
- 100% coverage doubles the number of cores required
- As a result, full DR is rarely implemented

## ▪ Processing comparisons

- Language expansion (CICS/COBOL path lengths are highly optimized)
- Networking drives up cycles spent on protocols
- Mainframe has dedicated processors for I/O operations, distributed does not
- Converting classic file systems to relational results in up to 3x expansion
- Zero network traffic on mainframe reduces computation (and latency)




# Mainframe Scales 2.5 to 15X Superdome More Performance / Watt



Notes: Performance as per Eagle TCO studies. Multiply by 2 for MIPS. HP performance based on 122 perf units / MIPS.

z10 and z196 power is max value. It is very rare that any mainframe is even 80% of max. Typical mainframe power is less - approximately 60% of maximum as per field data. Mainframe Power scales by model or book package.


# z196 Capacity per Watt improvements



15 years of CMOS: G2 to z196 *		Net Effect: G2 to z196 *	
Power Increase:	17% per year	Performance increased by:	300x
Performance increase:	46% per year	Performance / kWatt increased by:	30x
Power density	13% per year	Performance / sq ft increased by:	190x

Note: Capacity/kWatt assumes hot room, max plugged I/O power, max memory power and all engines turned on. Real world max capacity system is about 3/4 of this.

# System z Labor Cost Trends Favor A Centralized Approach To Management



Large scale consolidation and structured management practices drive increases in labor productivity

Small scale consolidation achieves lesser gains

**The more workloads you consolidate and manage with structured practices...  
the lower the management labor cost**

# Accumulated Field Data For Labor Costs

- **Average of quoted infrastructure labor costs**


- **30.7** servers per FTE (dedicated Intel servers)
  - **67.8** hours per year per server for hardware and software tasks
- **52.5** Virtual Machines per FTE (virtualized Intel servers)
  - **39.6** hours per year per Virtual Machine for software tasks and amortized hardware tasks
  - Typical 8 Virtual Machines per physical server

- **Best fit data indicates**

- Hardware tasks are **32** hours per physical server per year
  - Assume this applies to Intel or Power servers
  - Internal IBM studies estimate **320** hours per IFL for zLinux scenarios
- Software tasks are **36** hours per software image per year
  - Assume this applies to all distributed and zLinux software images

# Five Key IT Processes For Infrastructure Administration


Time spent on each activity




- Deployment Management**  
– Hardware set-up and software deployment
- Incident/Capacity Management**  
– Monitor and respond automatically
- Asset Management**  
– Hardware and software asset tracking
- Security Management**  
– Access control
- Change Management**  
– Hardware and software changes

# zManager Labor Cost Reduction Benefits Case Study

5031 total hours per year **reduced by 38%** to 3111 hours per year



# New York Financial Services Company – Useful Lifetime Of 36 Month Lease

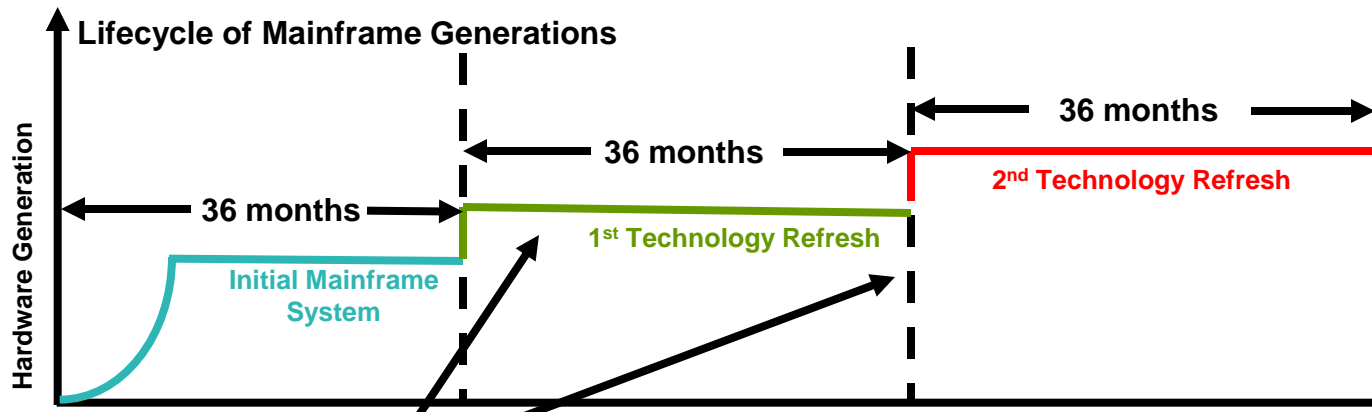


Observed at a large financial service customer

In each 36 month lease there are only 30 months production use

6 months provisioning  
 30 months production  
**Setup and tear down 15 People, 5 full time**

Setup and tear-down time costs 25% more. Plus . . . 41 hours of FTE setup and tear down labor per server = \$3,075



Weekend upgrades performed by IBM

Capacity on demand pricing

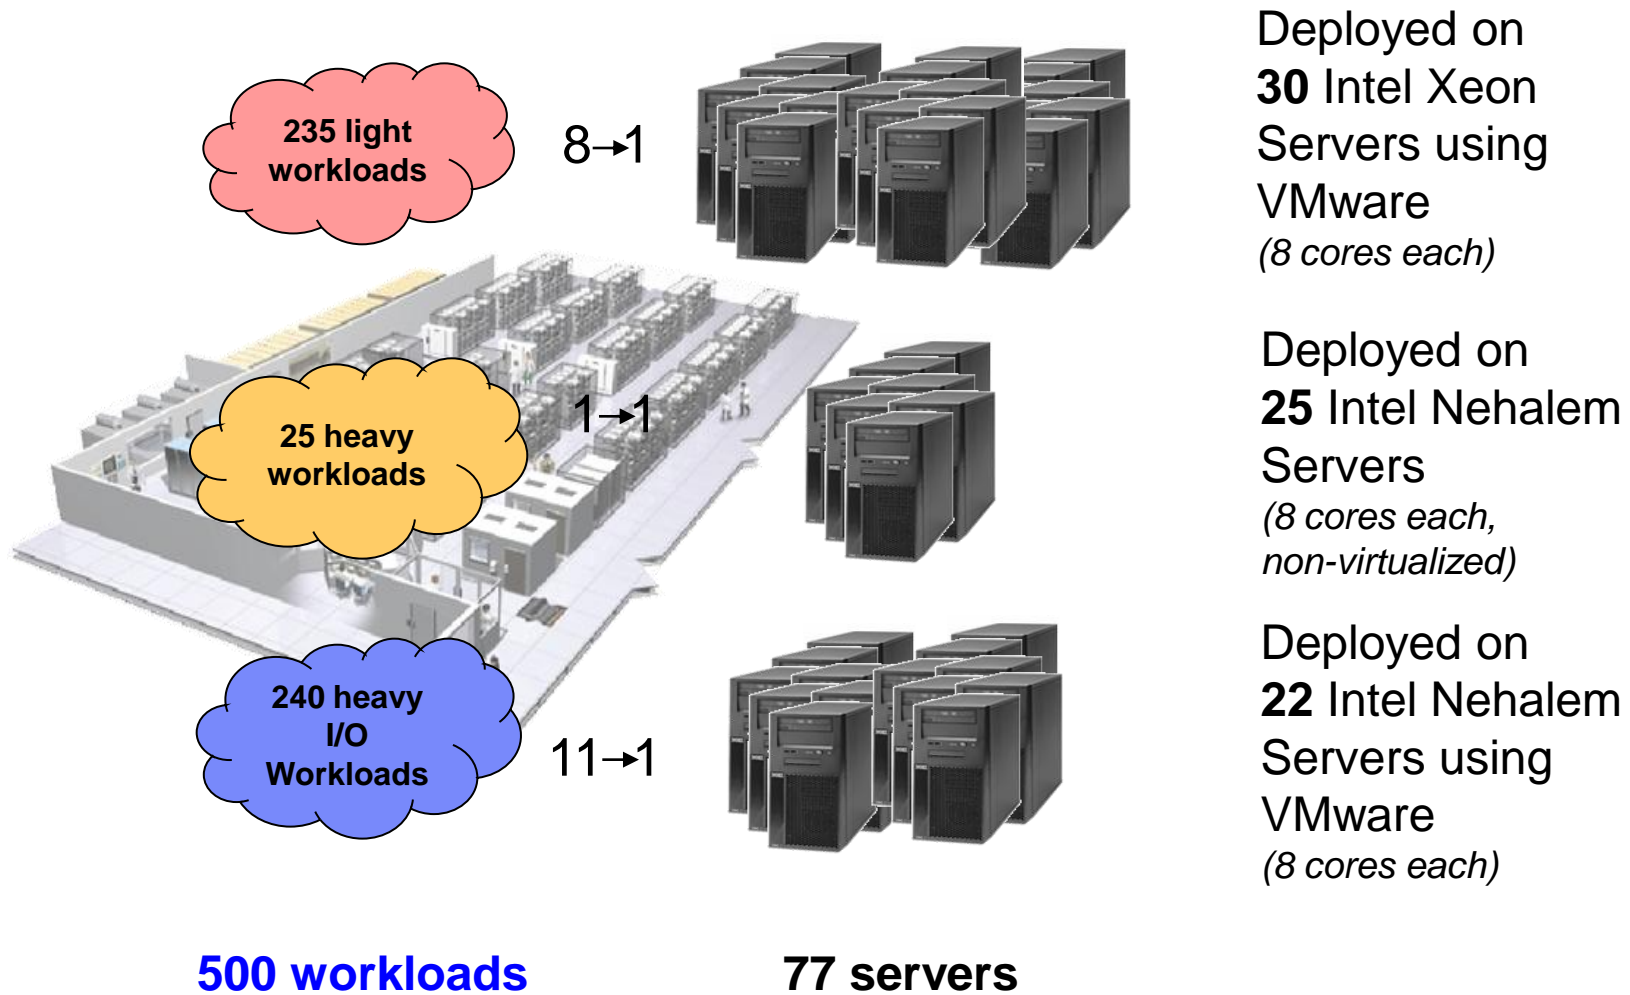
1 Weekend upgrading to new hardware and software levels  
 36 months production  
 No need to retire the server, upgrade in place

## Average Cost Ratios (z vs Distributed)

		z	Distributed	z vs distributed (%)
<b>Offload</b>	<b>5-Year TCO</b>	<b>\$14,617,537</b>	<b>\$25,016,633</b>	<b>58.43%</b>
	Annual Operating Cost	\$2,930,180	\$3,342,404	87.67%
	Software	\$9,349,434	\$10,045,104	93.07%
	Hardware	\$3,045,738	\$4,007,849	75.99%
	System Support Labor	\$3,207,949	\$5,109,879	62.78%
	Electricity	\$36,144	\$191,862	18.84%
	Space	\$56,027	\$148,727	37.67%
	Migration	\$586,808	\$8,716,612	6.73%
	DR	\$715,357	\$2,707,487	26.42%
	Average MIPS	3,128		
Total MIPS	140,759			
<b>New Workload</b>	<b>5-Year TCO</b>	<b>\$2,295,560</b>	<b>\$6,821,249</b>	<b>33.65%</b>
	Annual Operating Cost	266,530	693,442	38.44%
	Software	1,073,625	2,785,542	38.54%
	Hardware	669,311	1,313,598	50.95%
	System Support Labor	1,418,025	1,247,685	113.65%
	Electricity	13,920	\$85,569	16.27%
	Space	7,993	291,656	2.74%
	Migration	0	0	
	DR	68,005	2,269,640	3.00%
	Average MIPS	5,012		
Total MIPS	15,035			
<b>Consolidation</b>	<b>5-Year TCO</b>	<b>\$8,713,071</b>	<b>\$14,347,493</b>	<b>60.73%</b>
	Annual Operating Cost	\$1,087,137	\$2,328,635	46.69%
	Software	\$3,641,376	\$9,734,725	37.41%
	Hardware	\$3,068,105	\$1,570,789	195.32%
	System Support Labor	\$2,380,009	\$4,491,882	52.98%
	Electricity	\$192,962	\$375,922	51.33%
	Space	\$130,731	\$270,787	48.28%
	Migration	\$2,294,437	\$0	
	DR	\$416,326	\$632,933	65.78%
	Average MIPS	10,635		
Total MIPS	15,035			



# Large Data Center – What Did It Cost to Deploy 500 Workloads on Virtualized Intel Servers?




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

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

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


# Large Data Center – What Does it Cost to Deploy 500 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 x blades is a statement of direction only. Results may vary based on customer workload profiles/characteristics. © 2010 IBM Corporation

# Compare Server Cost of Acquisition



Deployed on Intel

Best fit on zEnterprise



**77 Intel Servers**

616 cores

**2 Frames**

192 cores

**\$15.2M** TCA (3 years)


**\$7.5M** TCA (3 years)

**51% less**

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


# Compare Network Cost of Acquisition

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**

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

# Compare Power Consumption



77 Servers  
**289 kW**

**\$0.25M**  
 3 years @ \$0.10 per kWh

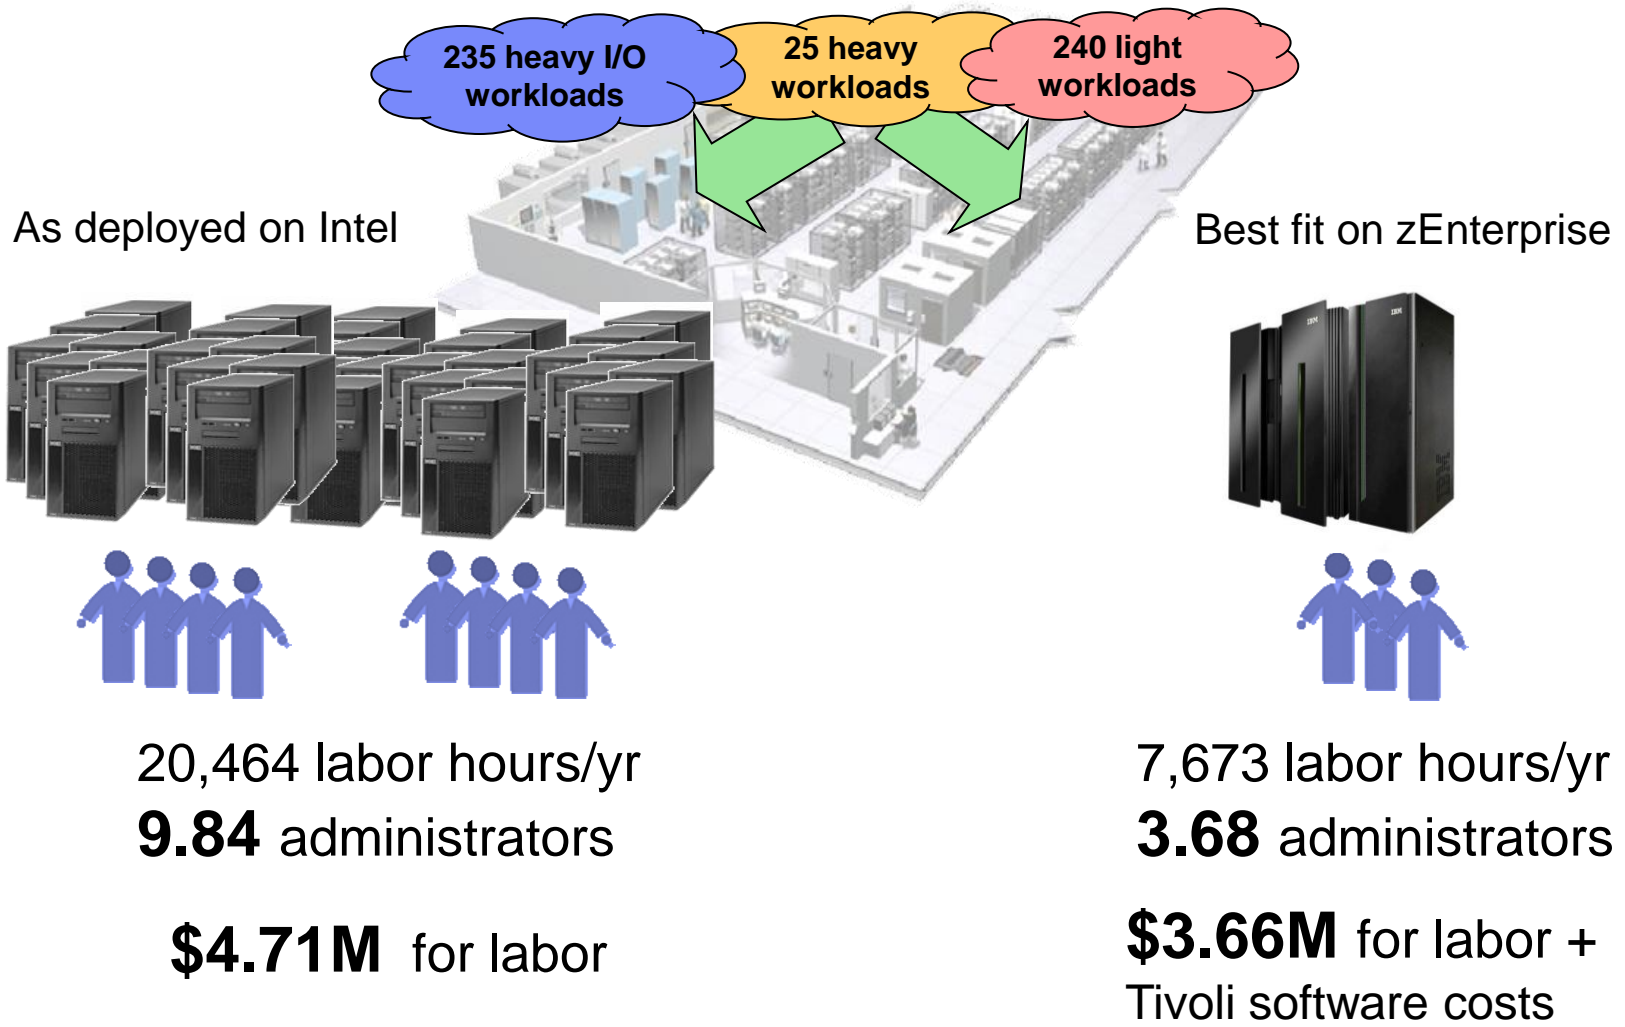
2 frames  
**67 kW**

**\$0.06M**  
 3 years @ \$0.10 per kWh

**77% less**

Server configuration based on IBM internal studies.  
 Calculations for Intel servers based on published power ratings and industry standard rates. Prices are in US currency, prices will vary by country

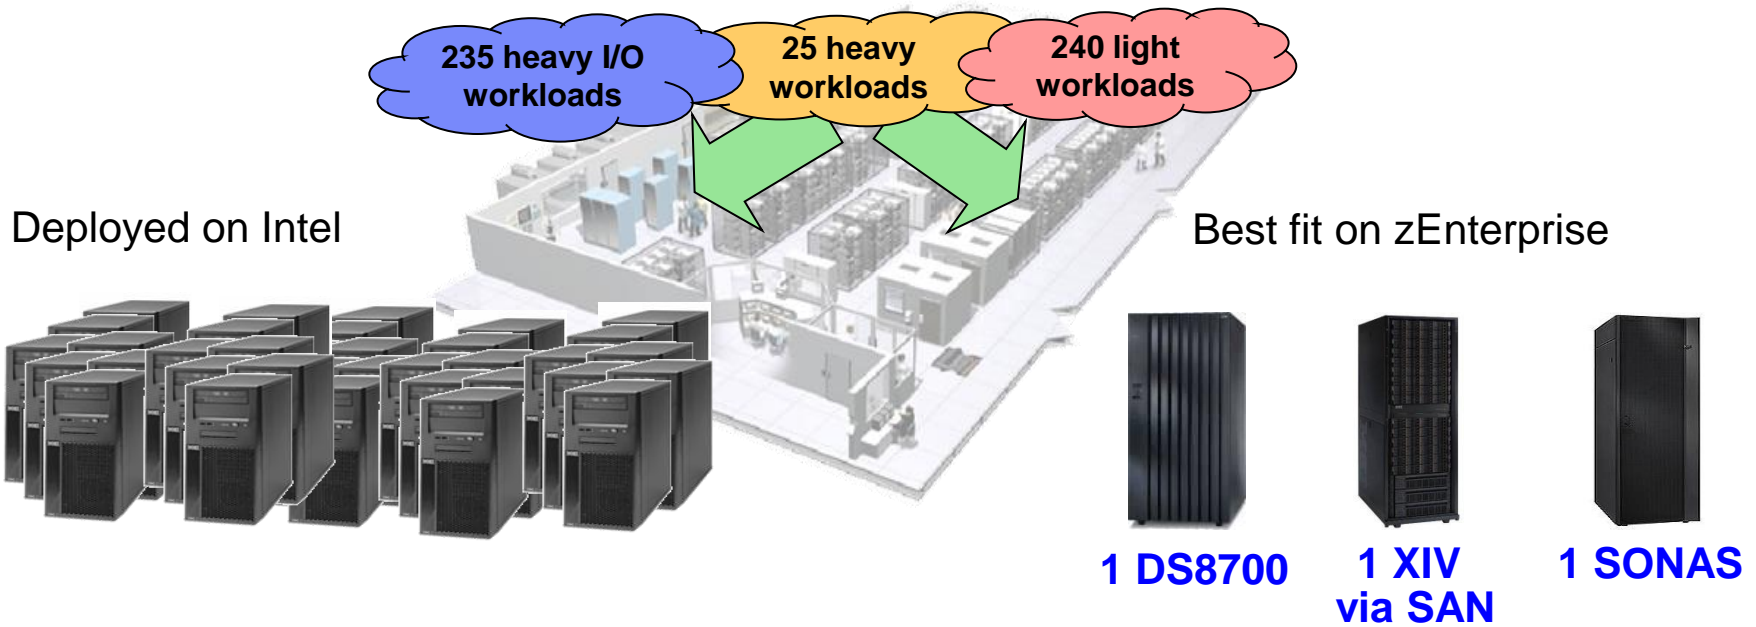
# Compare Server Infrastructure Labor Cost



**22% less**

Configuration based on IBM internal studies. Labor model based on customer provided data from IBM studies. Labor rates will vary by country

# Compare Storage Cost



**484.4 TB** embedded storage  
 24% utilization  
 580 points of admin

**\$9.1M** TCO(3 years)

**172.3 TB** provisioned storage  
 67% utilization  
 3 points of admin

**\$6M** TCO (3 years)

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

**34% less**

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

# Fewer Parts to Assemble and Manage



Deployed on Intel
77
664
289
10
580


Servers  
 Network (parts)  
 Power (KW)  
 Administrators  
 Storage admin points

Best fit on zEnterprise
2 frames
21
67
4
3





# The Savings are Cumulative




Three Year Cost Of	Deployed on Intel	Best fit on zEnterprise
Servers	\$15.2M	\$7.5M
Network	\$0.20M	\$0.03M
Power	\$0.25M	\$0.06M
Labor	\$4.71M	\$3.66M
Storage	\$9.1M	\$6.0M
Total	\$29.46M	\$17.25M
<b>Total cost per workload</b>	<b>\$59K</b>	<b>\$35K</b>

**41% less**

# Summary

- **Cost per workload is the key metric for the new IT economics**
  - Mainframe cost per work goes down as workload increases



- **Fit for purpose reduces cost of acquisition per workload**
- **zEnterprise's integrated management reduces cost per workload with extreme automation for simplicity**

# Thank you

