Analyzing IT Value and Cost Considerations – Maximizing The Value of Your Mainframe

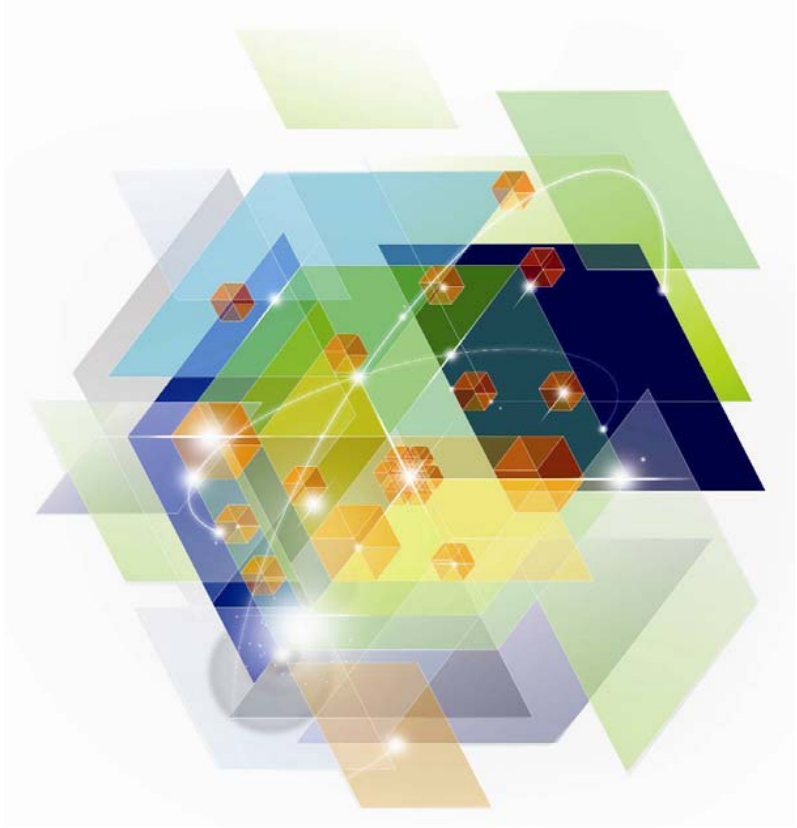

Ray Jones, Vice President,
Worldwide System z Software Sales,
IBM Software Group

March 21, 2012

© 2012 IBM Corporation

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

© 2012 IBM Corporation

A Closer Look At Fit-For-Purpose Workload Assignment

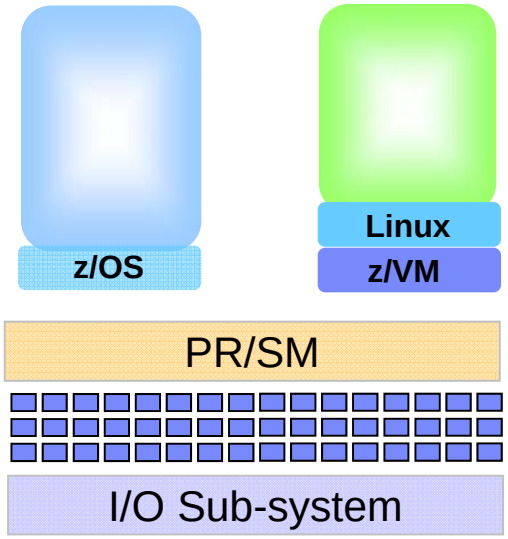

- Scale up to 80 cores in a frame (z/OS clusters with sysplex)
- Dedicated I/O sub-system
- Superior qualities of service

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

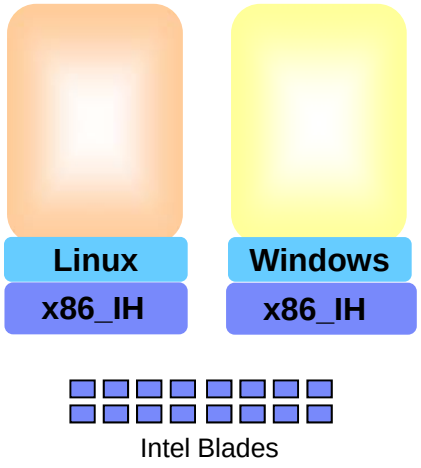

- Scales to 16 cores per blade
- 2 fast processing threads per core
- Commodity I/O
- Modest qualities of service

© 2012 IBM Corporation

Workload Characteristics Influence The Best Fit Deployment Decision

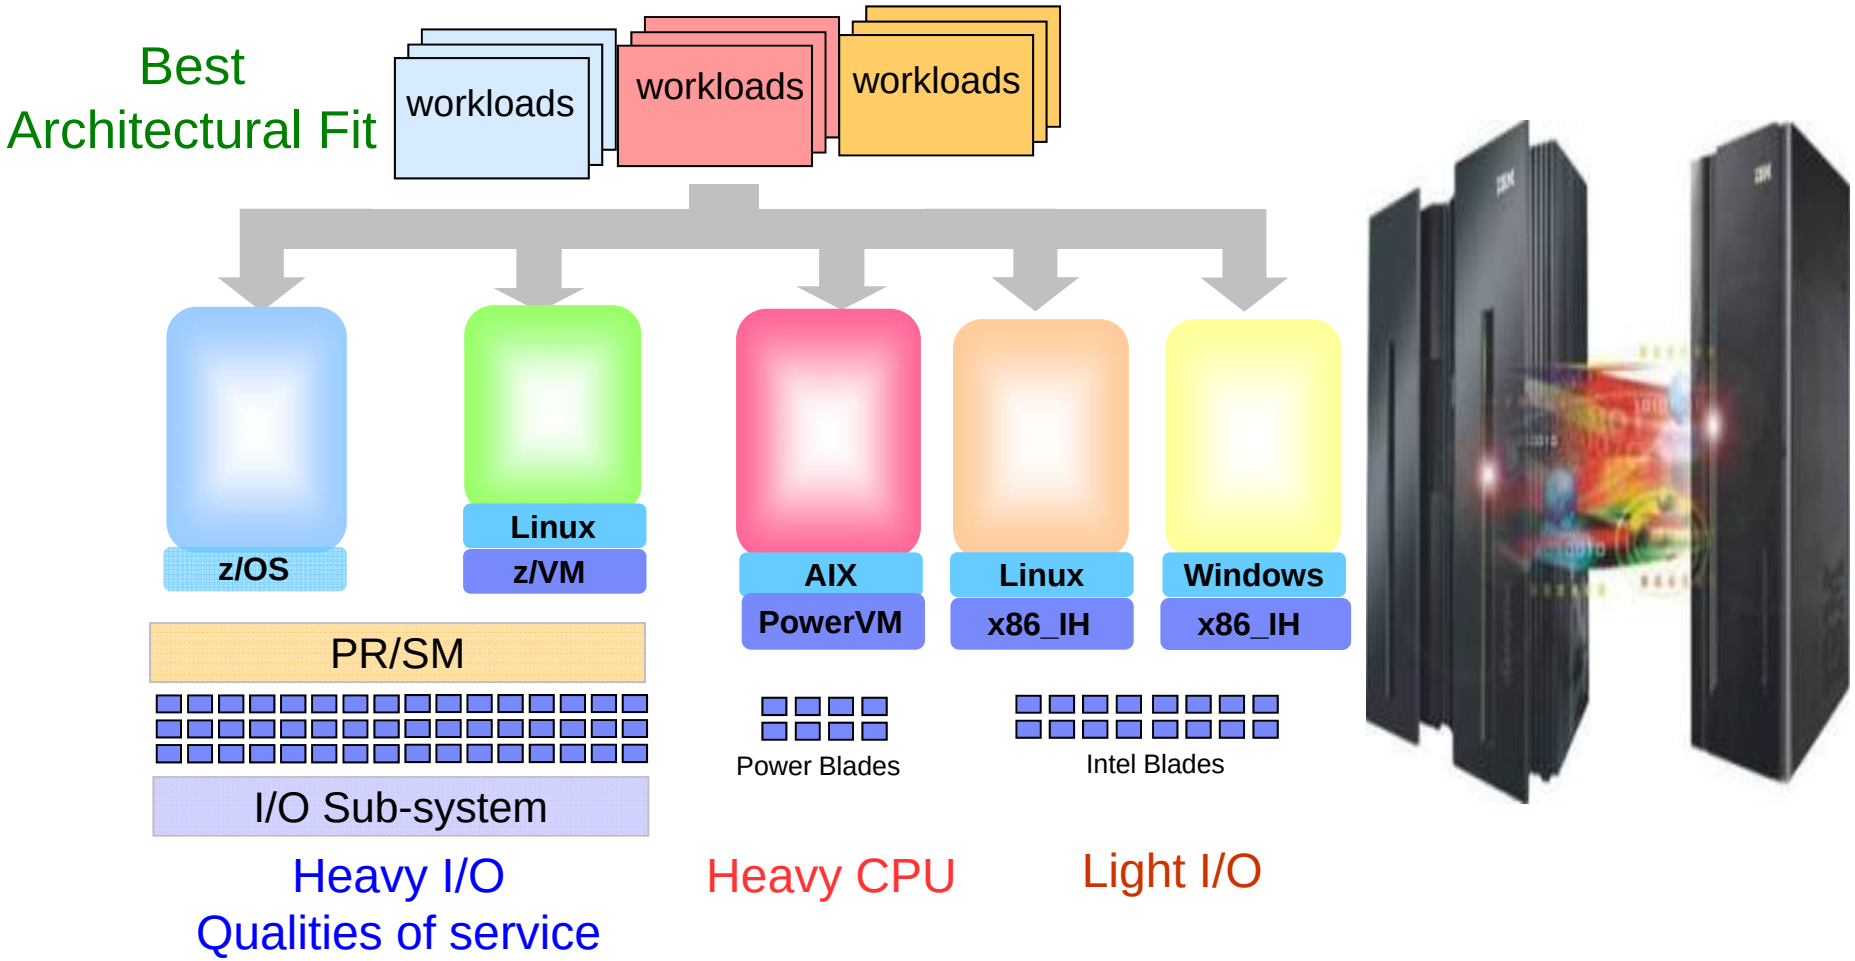

Deploy or consolidate workloads on the environment best suited for each workload to yield lowest cost

© 2012 IBM Corporation

Deploying Stand Alone Workloads With Heavy CPU Requirements

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

Heavy CPU workloads

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

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
$200,055 per workload

Best Fit

1 workload per POWER7 blade

PowerVM on PS701
8 core POWER7 Blade
$216,658 per workload

10 workloads per 32-way z/VM

z/VM on z196 CPC
32 IFLs
$328,477 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.

© 2012 IBM Corporation

Deploying Stand Alone Workloads With Light CPU Requirements

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

Light workloads

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

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
$8,165 per workload

28 workload per POWER7 blade

Fast low cost threads

PowerVM on PS701
8 core POWER7 Blade
$7,738 per workload
Best Fit

155 workloads per 32-way z/VM

z/VM on z196 CPC
32 IFLs
$21,192 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.

© 2012 IBM Corporation

Deploying Stand Alone Workloads With Heavy I/O Requirements

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

Heavy I/O 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**

1 workload per Intel blade

Virtualized on Intel
16 core HX5 Blade
$400,109 per workload

1 workload per POWER7 blade

PowerVM on PS701
8 core POWER7 Blade
$216,658 per workload

40 workloads per 32-way z/VM

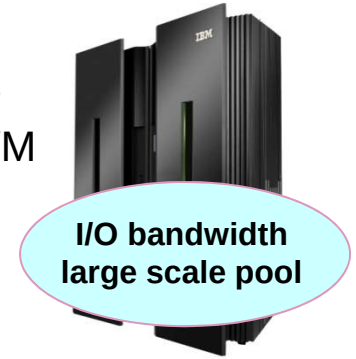

z/VM on z196 CPC
32 IFLs
$82,119 per workload
Best Fit

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.

© 2012 IBM Corporation

Benchmarks Show System z And z/OS Are Optimized For Batch Processing

	Intel x3550	Power PS701	Linux on z	z/OS
	12 processors 128 GB RAM DS8300 	8 processors 128 GB RAM DS8300 	8 processors 128 GB RAM DS8800 	8 processors 128 GB RAM DS8800 
	Sorting Average CPU 89%	Sorting Average CPU 92%	Sorting Average CPU 90%	Sorting Average CPU 72%
SORT *Job: Sort a 3 GB transaction file – Repetitions: 300*				
Total Time (secs)	7,680	6,900	2,590	644
Concurrency	12	20	18	45
Rate (MB/sec)	240	280	746.2	3,000
MERGE *Job: Merge 30 sorted files into a 90 GB master file – Repetitions: 10*				
Total Time (secs)	11,709	7,920	2,799	558
Concurrency	10	10	10	10
Rate (MB/sec)	157	244	690.5	3,460

Results:

1. **Running same software, x86 batch window is 3.6x greater than System z**
2. **On System z, Linux batch window is 4.5x greater than z/OS**
3. **Off-loading batch from z/OS to x86 leads to as much as 16x increase in batch window**

Source: Internal IBM study

© 2012 IBM Corporation

Core Proliferation for a Mid-sized Offload Project

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

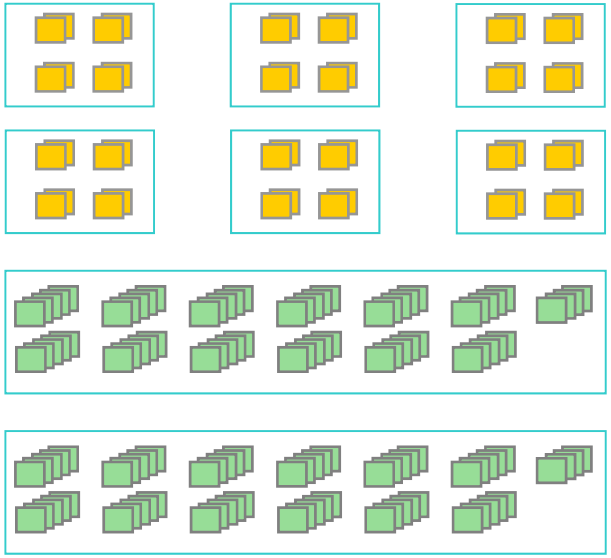

$25.4M TCO (5yr)

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

6 processors
(1,660 MIPS)

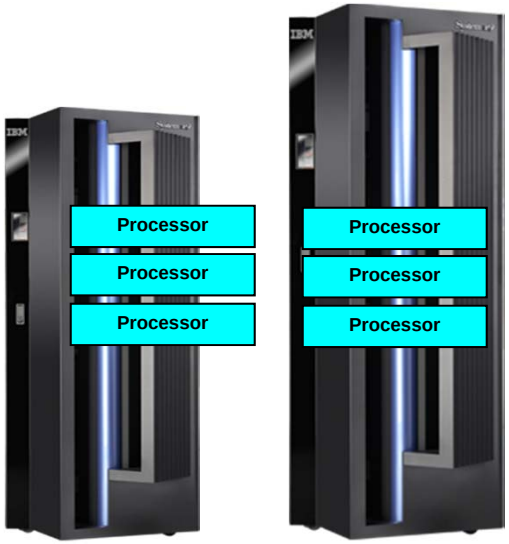

176 distributed processors
(800,072 Performance units)

$17.9M TCO (5yr)

482 Performance Units per MIPS

© 2012 IBM Corporation

Clusters Grow Database Processing Power Beyond Single Server Solutions

DB2 for z/OS

Centralized Coupling Facility Design

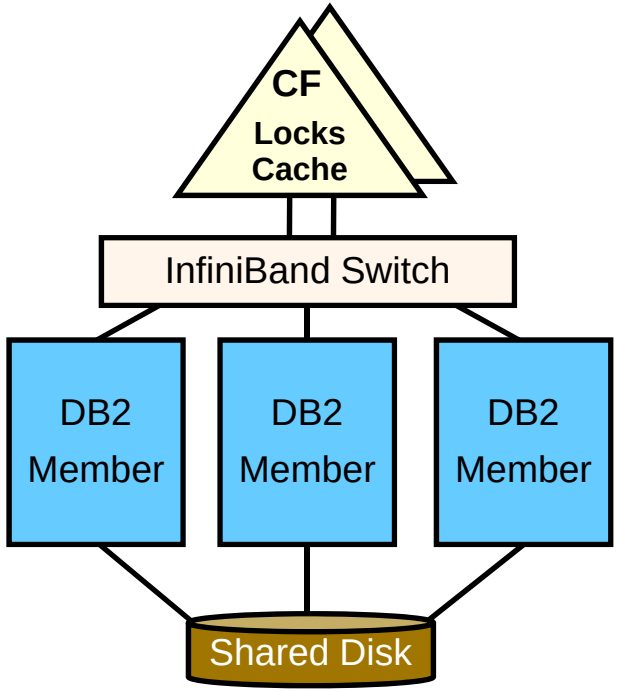

Efficient lock and buffer management achieve near linear scalability

Oracle RAC

Distributed Design

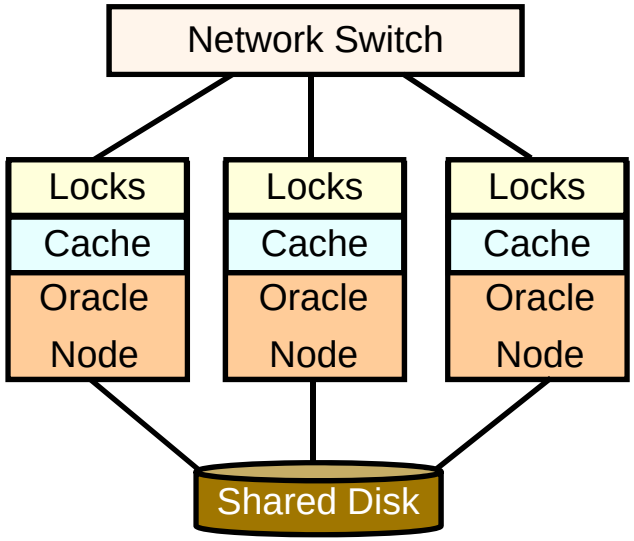

Inefficient distributed locking and buffer management limits scaling

© 2012 IBM Corporation

ISAS 9700 + IDAA Delivers

5X Performance At 25% The Unit Cost

Competitor

Quarter Rack

Unit Cost (3yr TCA)
$97/RpH

RpH (Reports/Hour)	29,572
Exadata V2 (HW+SW+Storage)	$2.9M

ISAS 9700

DB2 (ISAS 9700)

12 GP+12 zIIP

Unit Cost (3yr TCA)
$62/RpH

RpH (Reports/Hour)	57,904
ISAS 9700 (HW+SW+Storage)	$3.6M

ISAS 9700 + IDAA

DB2 (ISAS 9700)

z/OS 5 GP+5 zIIP

Netezza TwinFin 12

Unit Cost (3yr TCA)
$24/RpH

RpH (Reports/Hour)	154,893
ISAS 9700 10-cores (HW+SW+Storage)	$1.5M
NZ TF12 (HW+SW+Storage)	$2.1M

Source: IBM Competitive Project Office
Customer Study running 161,166 concurrent operational reports.
Results will vary based on customer workload profiles/characteristics.

© 2012 IBM Corporation

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

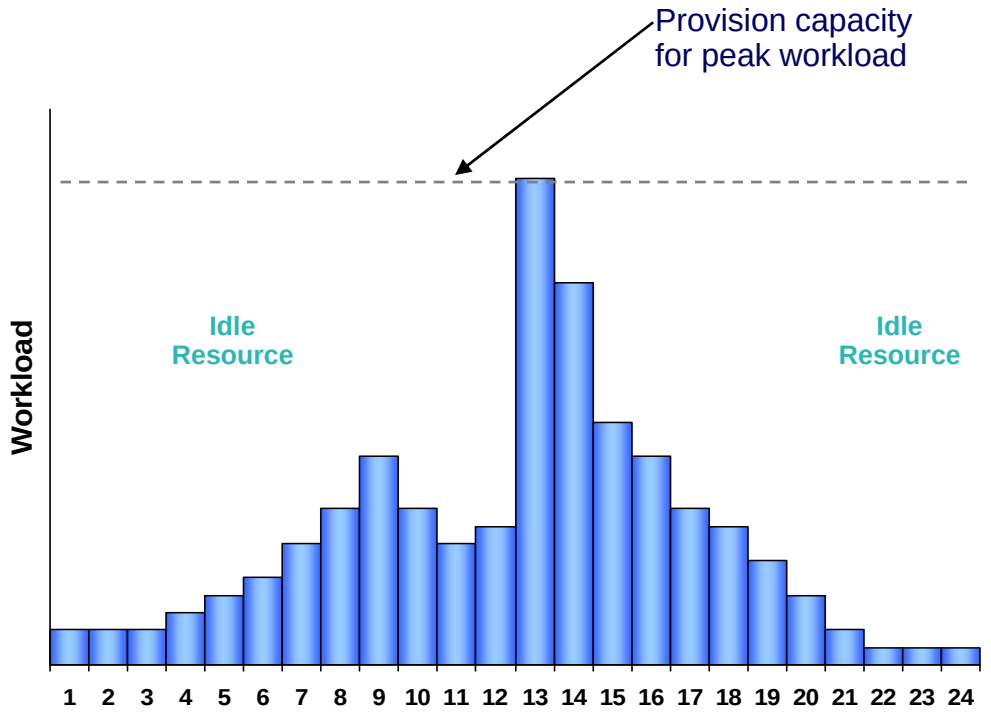

- **Storage Allocation**
  - Application-specific resulting in over-allocations
  - Fine grained storage allocation mechanisms characteristic of mainframe storage are uncommon in distributed environments.
- **Storage Utilization**
  - Single digit utilization for distributed environments is not uncommon
  - Storage utilization of 80% + is typical for mainframe
- **Storage Management**
  - Data disaster recovery, synchronization, and transfer requirements add complexity and cost

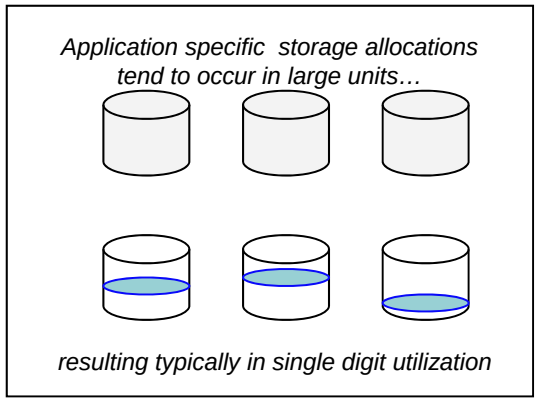

© 2012 IBM Corporation

What Is A Typical Value Of Sigma?

IBM Survey Of Workload Variability In 3200 Servers

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

IBM System x™ Servers and VMware Virtual Machine Sizing Guide

Legacy workloads on XEON 2.5-2.8GHz Servers

Normal probability distribution

© 2012 IBM Corporation

New Workload Scenarios – Beware Benchmarks

- **Stress test benchmarks have no variability!**
  - They drive the system under test to100% 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 Sigma=2.5*Mean requires processing capacity equal to **6 times the Mean** workload demand

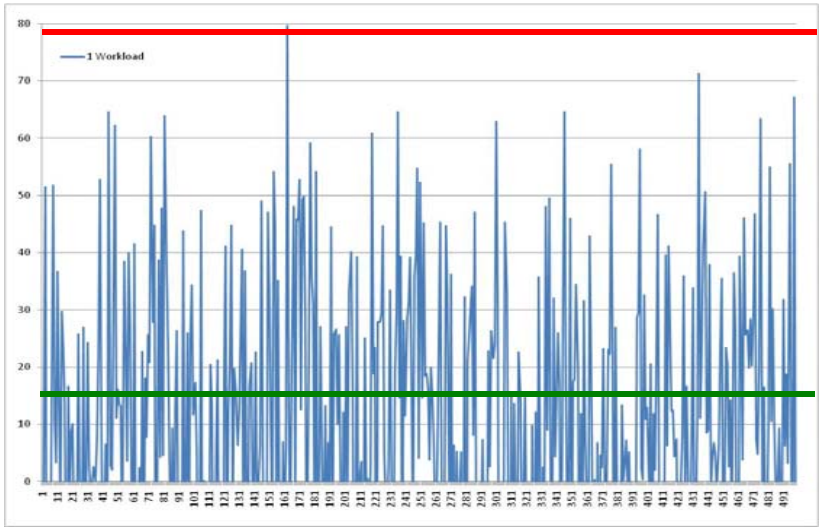

Adding a new workload to a pool of 256 existing workloads will require incremental processing capacity equal* to the **Mean** workload demand

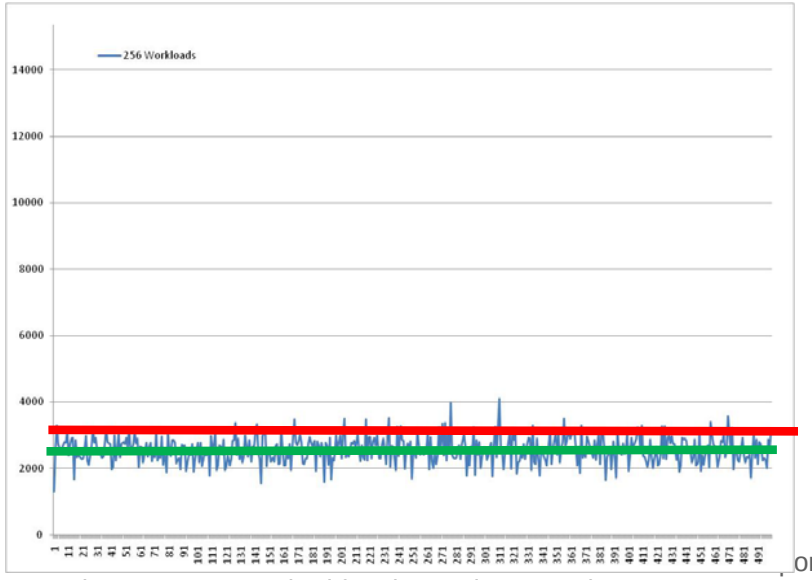

oration

* If we add one more workload to a pool of 256 consolidated workloads the computing resource required for the pool goes up by 1.00047 * Mean

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

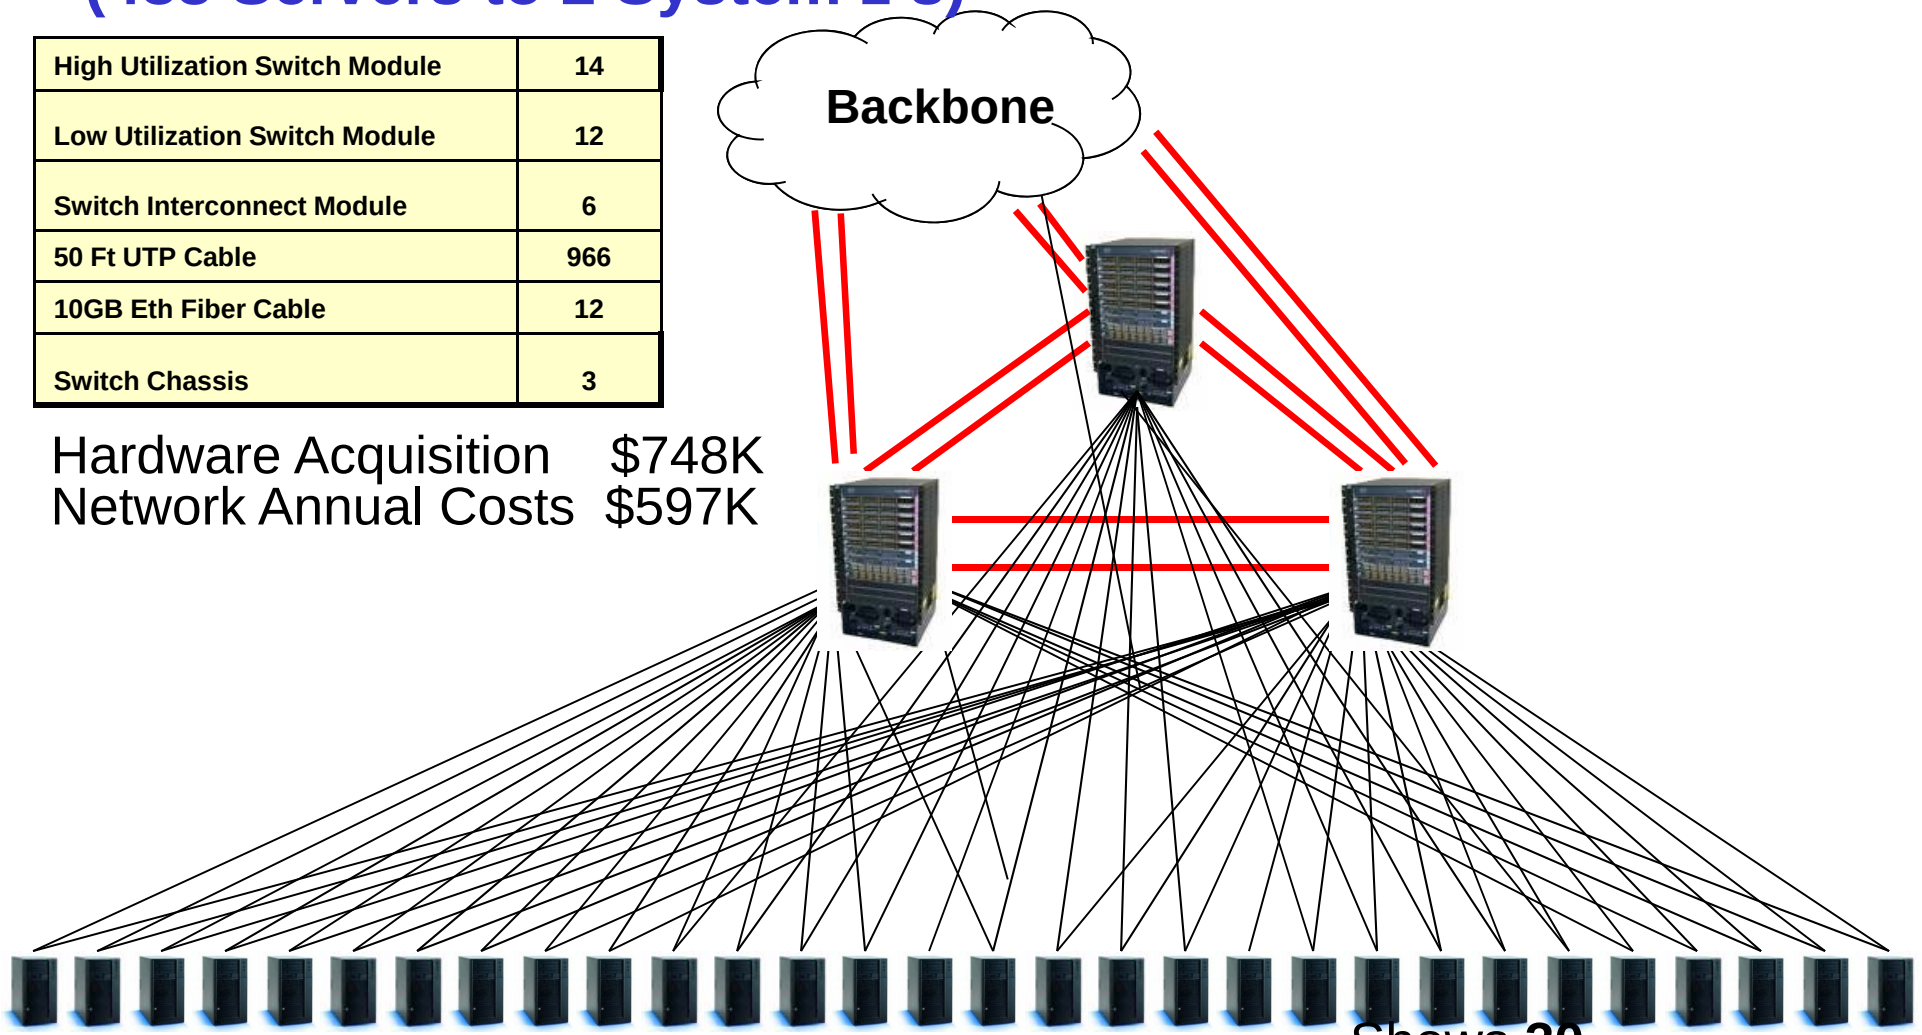

Shows **30** of the **483** Servers

© 2012 IBM Corporation

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

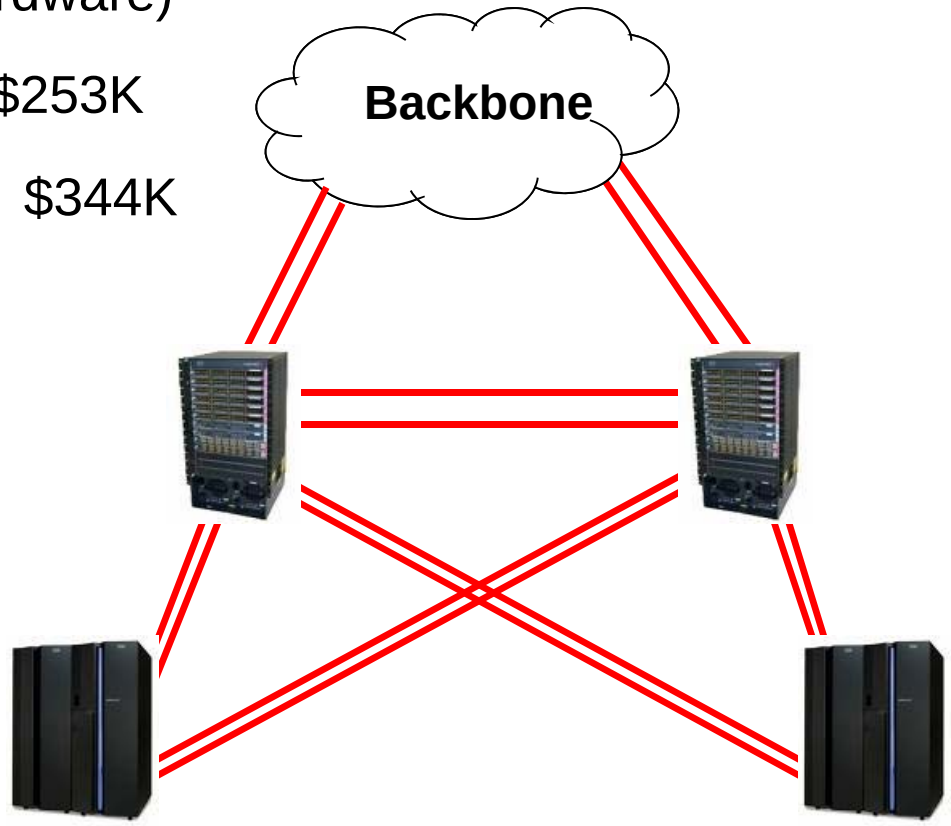

© 2012 IBM Corporation

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)

© 2012 IBM Corporation

System z Labor Cost Trends Favor A Centralized Approach To Management

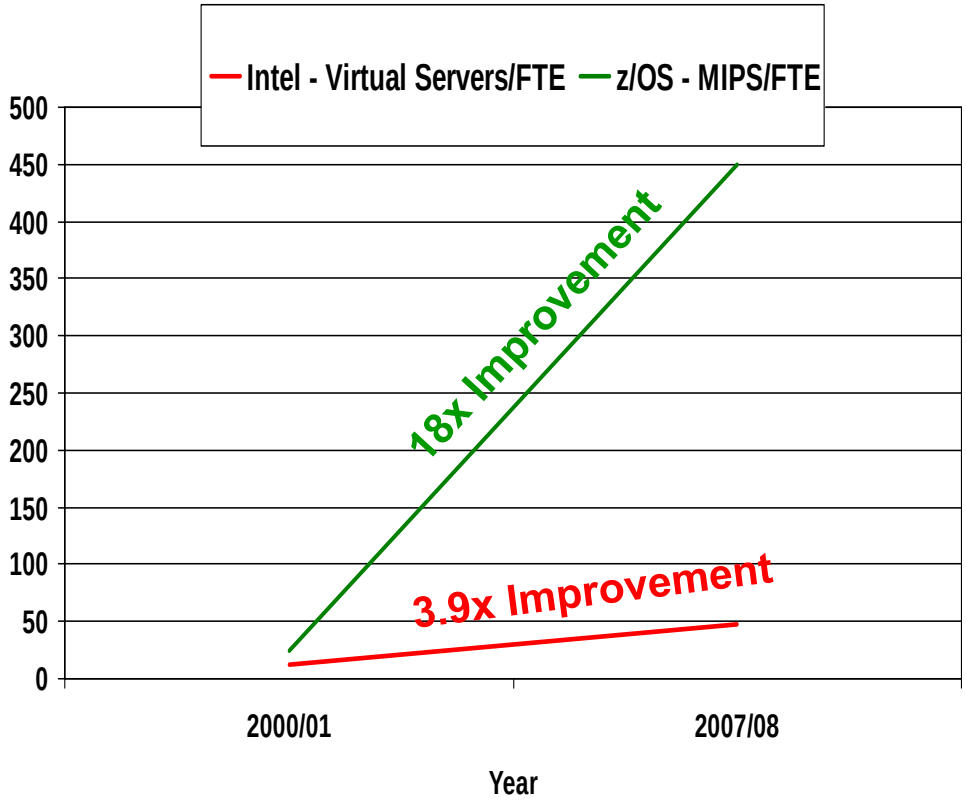

Large scale consolidation and structured management practices drive increases in labor productivity

Small scale consolidation achieves lesser gains

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

Source: IBM Scorpion Studies

© 2012 IBM Corporation

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

Labor model based on customer data from IBM studies

© 2012 IBM Corporation

Five Key IT Processes For Infrastructure Administration

Time spent on each activity

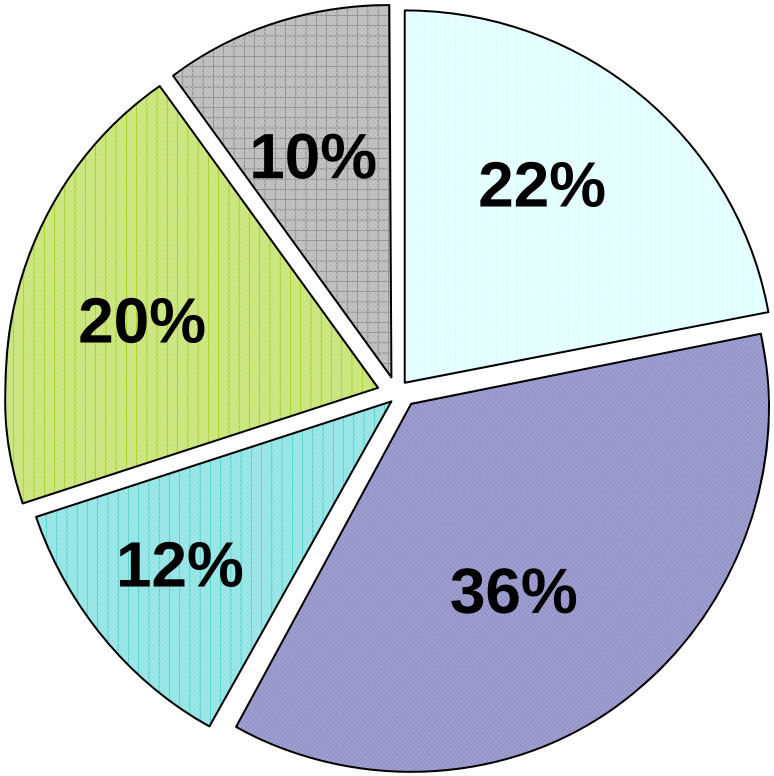

- **Deployment Management**
  - *Hardware set-up and software deployment*
- **Incident/Capacity Management**
  - *Monitor and respond automatically*
- **Asset Management**
  - *Hardware and software asset tracking*
- **Security Management**
  - *Access control*
- **Change Management**
  - *Hardware and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

Allocation based on customer data from IBM study

© 2012 IBM Corporation

zManager Labor Cost Reduction Benefits Case Study

5032 total hours per year reduced by 38% to 3111 hours per year

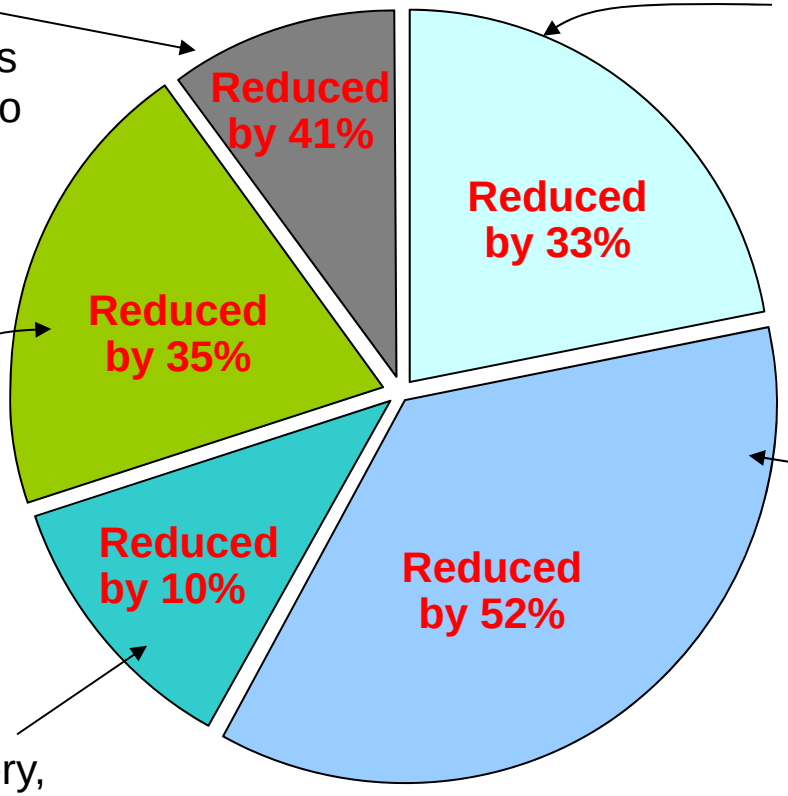

Change Management
Standardization of images and firmware, visibility into relationships among resources

Deployment Management
Automatic setup and configuration of the hypervisor and out-of-the-box networks

Security Management
Centralized fine-grain administrator access control

Incident/Capacity Management
Automation to isolate and fix issues

Asset Management
Automated discovery, entitlement management

© 2012 IBM Corporation

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

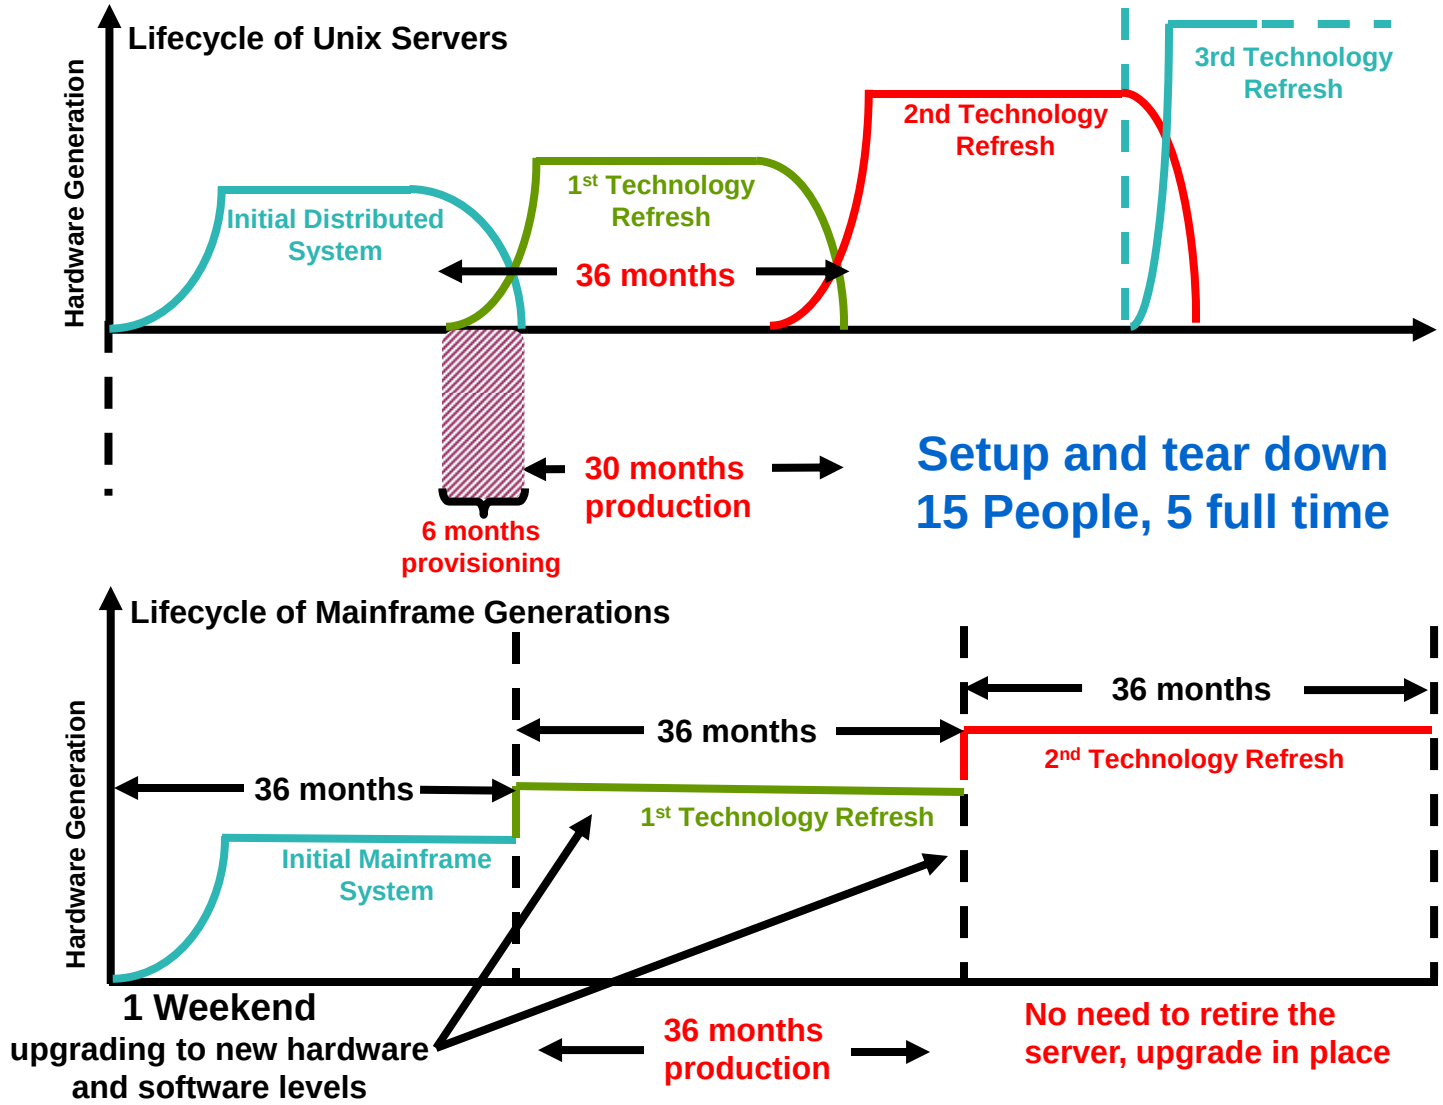

Observed at a large financial service customer

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

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

© 2012 IBM Corporation

Cost Ratios in all TCO Studies

Average Cost Ratios (z vs Distributed)

		z	Distributed	z vs distributed (%)
Offload	**5-Year TCO**	**$16,351,122**	**$31,916,262**	**51.23%**
	Annual Operating Cost	$2,998,951	$4,405,510	68.07%
	Software	$10,932,610	$16,694,413	65.49%
	Hardware	$3,124,013	$3,732,322	83.70%
	System Support Labor	$3,257,810	$4,429,166	73.55%
	Electricity	$45,435	$206,930	21.96%
	Space	$59,199	$154,065	38.42%
	Migration	$438,082	$10,690,382	4.10%
	DR	$854,266	$2,683,652	31.83%
	Average MIPS	3,954		
	Total MIPS	217,452		
Consolidation	**5-Year TCO**	**$5,896,809**	**$10,371,020**	**56.86%**
	Annual Operating Cost	$716,184	$1,646,252	43.50%
	Software	$2,240,067	$6,689,261	33.49%
	Hardware	$2,150,371	$1,052,925	204.23%
	System Support Labor	$1,766,403	$2,395,693	73.73%
	Electricity	$129,249	$365,793	35.33%
	Space	$84,033	$205,860	40.82%
	Migration	$678,449	$0	
	DR	$354,735	$411,408	86.22%
	Average MIPS	10,821		
	Total MIPS	292,165		

© 2012 IBM Corporation

Case Study – Consolidate 880 Standalone Workloads And Integrate 44 Hybrid Workloads On zEnterprise

- **Standalone distributed workload profile is a mix of**
  - **784 light**
  - **56 heavy CPU**
  - **40 heavy I/O**
- **Hybrid workload profile is a mix of**
  - **24 Web front-end workloads to CICS on z/OS**
  - **20 SAP application workloads with DB2 on z/OS**
- **What is the most cost effective way to consolidate/deploy all these workloads?**

Sun Fire X4170

zEnterprise

© 2012 IBM Corporation

What Is Best Fit For 880 Standalone Workloads On zEnterprise?

784 light workloads

28 workloads per POWER7 blade

28

IBM POWER7 blades

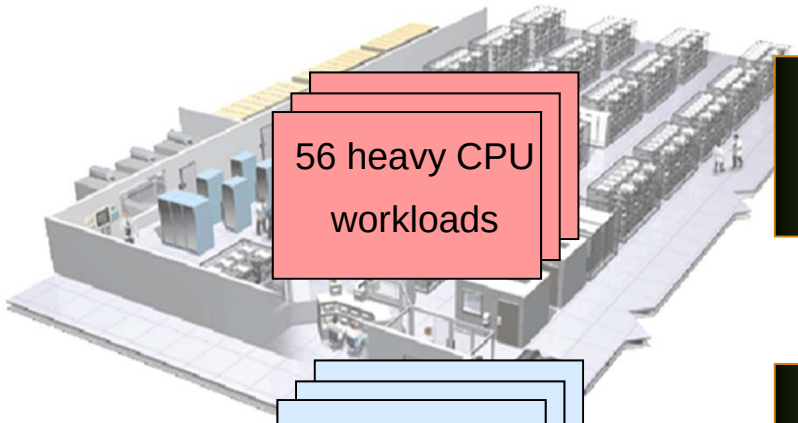

56 heavy CPU workloads

2 workloads per HX5 blade

28

IBM HX5 Blades

40 heavy I/O workloads

40 workloads per 32 IFLs

1

z196 CPC

32 IFLs

Best fit assignments

Server configurations are based on consolidation ratios derived from IBM internal studies. Projected Sun Fire X4470 2.0GHz 2ch/16co from x3550 2.66GHz 2ch/12co measurements. Prices are in US currency, prices will vary by country

© 2012 IBM Corporation

What Is Best Fit For 44 Hybrid Workloads On zEnterprise?

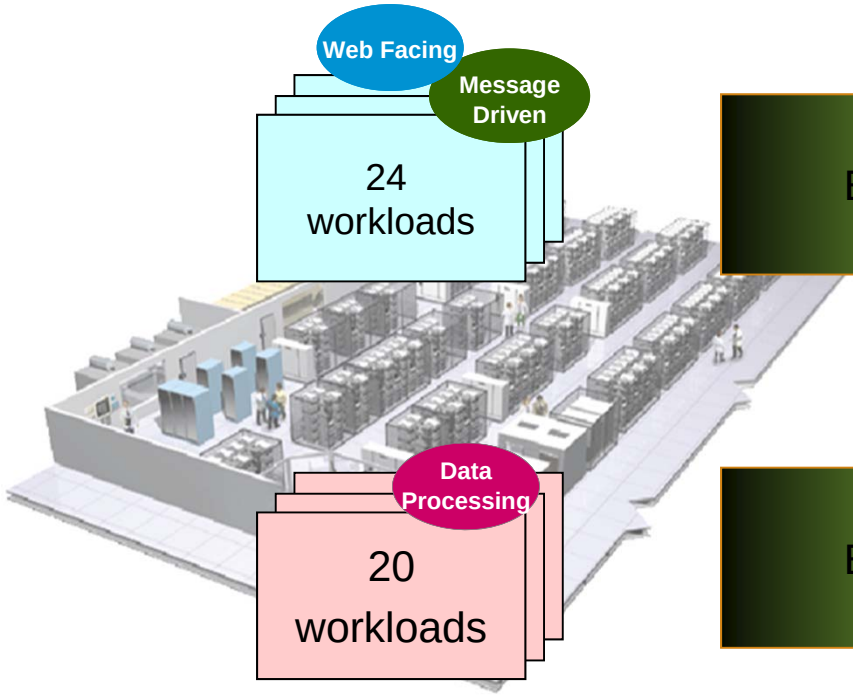

Best Fit

Best Fit

24 IBM POWER7 blades

2 DataPower XI50z blades

23 IBM POWER7 blades

CICS and DB2 components are Best Fit on z/OS

© 2012 IBM Corporation

Compare Server Hardware And Software Cost Of Acquisition

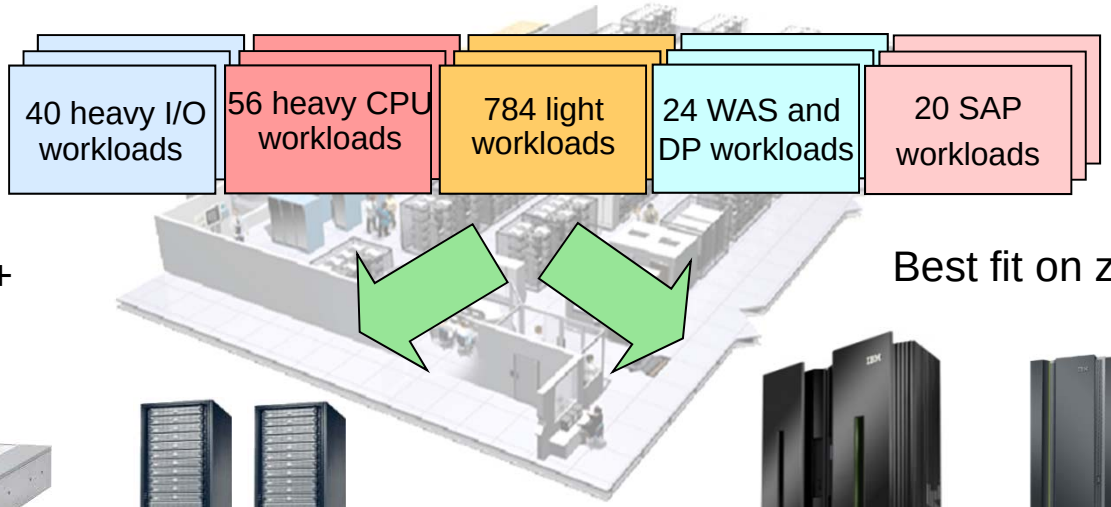

Deployed on Sun + HP servers

123 Sun Fire X4170

1476 cores

183 servers

2,060 cores

$46.0M Total

3yr TCA HW+SW

24 Sun Fire X4170

34 Sun T4-1

560 cores

2 DL380

24 cores

Best fit on zEnterprise

z196

32 IFLs

105 Blades

1,048 cores

106 servers
1,080 cores

43% less

$26.1M Total

3yr TCA HW+SW

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

© 2012 IBM Corporation

Compare Network Cost Of Acquisition

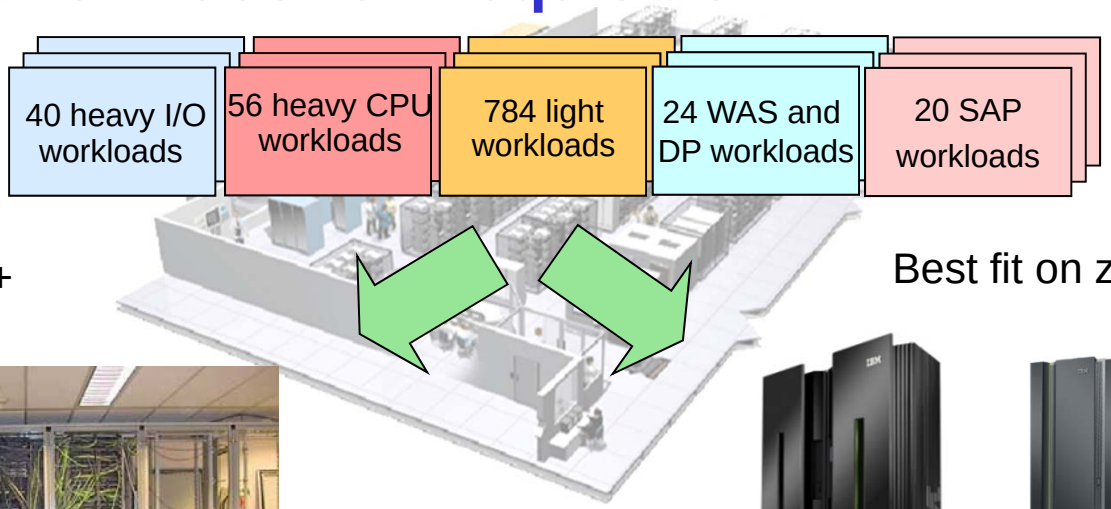

Deployed on Sun + HP servers

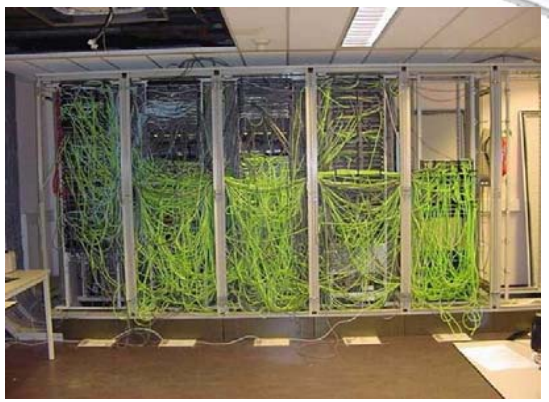

Additional network parts

37 switches

814 cables

740 adapters

1,591 total network parts

$0.45M **Total**

Best fit on zEnterprise

Additional network parts

1 switch

10 cables

10 adapters

21 total network parts

$0.03M **Total**

94% less

Network configuration is based on IBM internal studies.

Prices are in US currency, prices will vary by country

© 2012 IBM Corporation

Compare Power Consumption

40 heavy I/O workloads

56 heavy CPU workloads

784 light workloads

24 WAS and DP workloads

20 SAP workloads

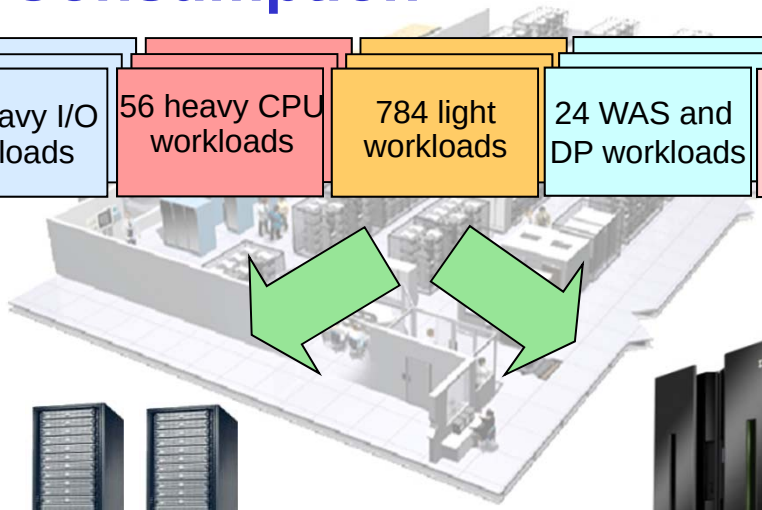

Deployed on Sun + HP servers

183 servers

124.1 kW

$0.33M **Total**

3 years
@ $0.10 per kWh

Best fit on zEnterprise

106 servers

53.4 kW

$0.14M **Total**

3 years
@ $0.10 per kWh

57% less

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

© 2012 IBM Corporation

Compare Server Infrastructure Labor Costs

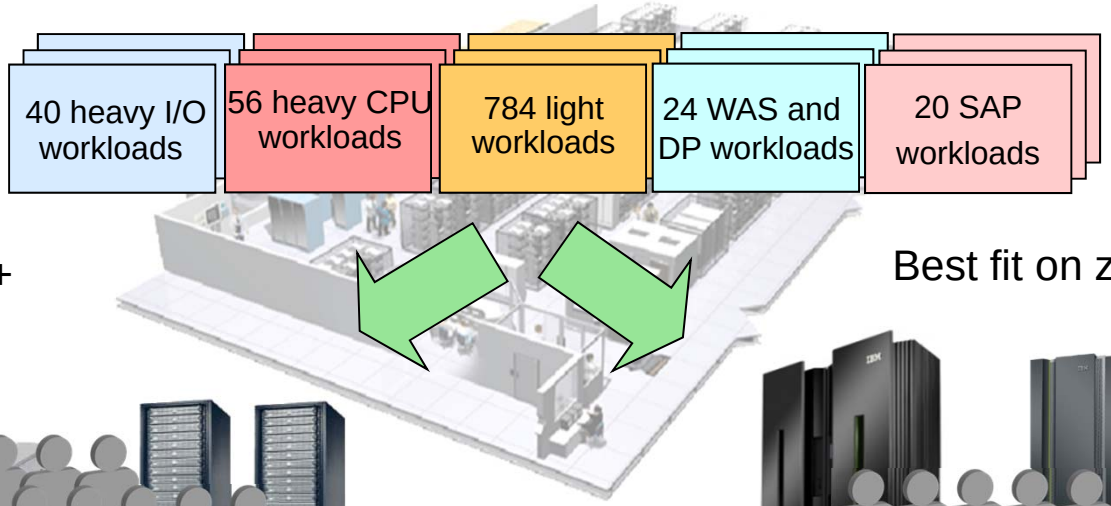

Deployed on Sun + HP servers

39,184 labor hrs/yr

18.83 administrators

$9.02M **Total**

3 years
@ $159,600/yr

Best fit on zEnterprise

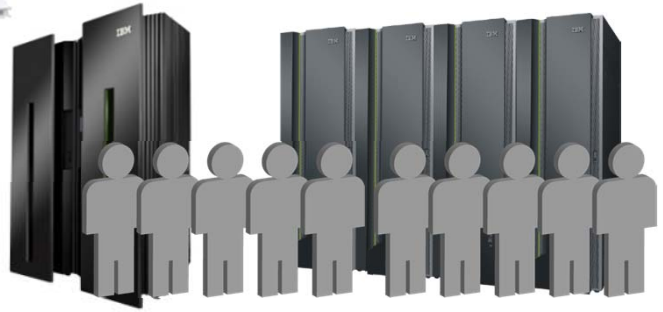

26,441 labor hrs/yr

12.71 administrators

$6.09M **Total**

3 years
@ $159,600/yr

32% less

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

© 2012 IBM Corporation

Compare Storage Costs

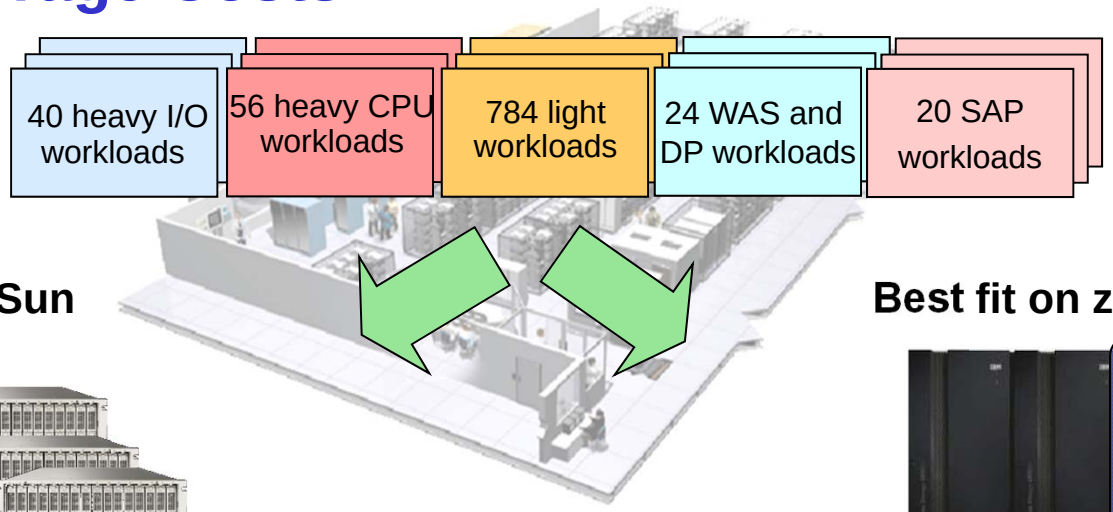

Deployed on Sun

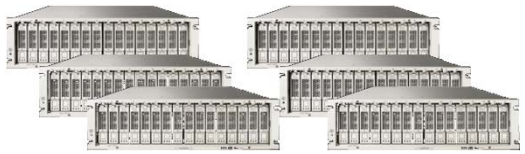

Sun Storage 6180 Array Sun F5100 Storage Flash Array

232.8TB embedded storage
36.57% utilization
70 points of admin

$8.58M TCO(3 years)

Best fit on zEnterprise

Incremental add on DS8800

143.04TB provisioned storage
59.52% utilization
1 points of admin

$4.6M TCO (3 years)

45% less

75GB/240GB active storage required per workload

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

© 2012 IBM Corporation

Compare Total Cost Of Ownership

40 heavy I/O workloads | 56 heavy CPU workloads | 784 light workloads | 24 WAS and DP workloads | 20 SAP workloads

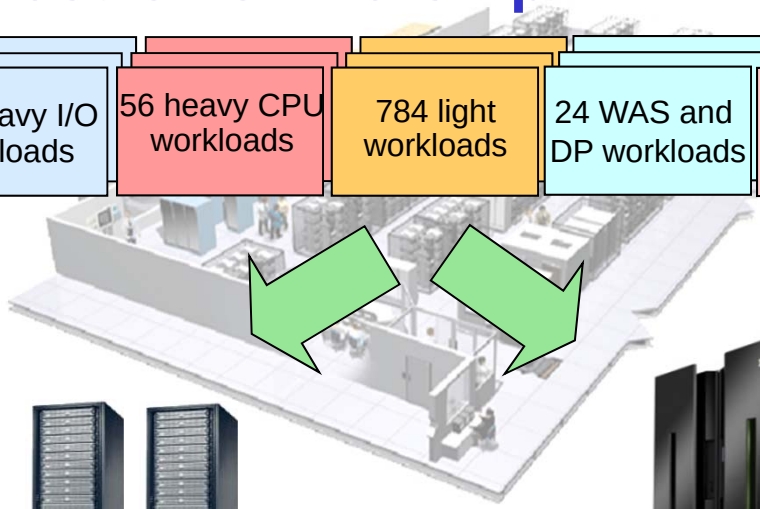

Deployed on Sun + HP servers

183 servers

2,060 cores

$64.38M Total

or $70K per workload

3yr TCO

Best fit on zEnterprise

106 servers

1,080 cores

$36.96M Total

or $40K per workload

3yr TCO

43% less

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

© 2012 IBM Corporation

Fewer Parts to Assemble and Manage

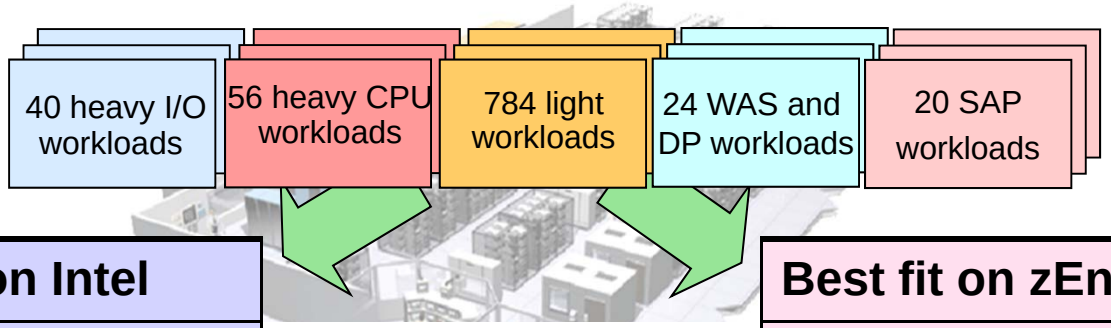

Deployed on Intel		Best fit on zEnterprise
183	Servers	1 z196 + 1 zBX (with 105 blades total)
1592	Network (parts)	21
124	Power (KW)	53
19	Administrators	13
70	Storage points	1

The Savings Are Cumulative

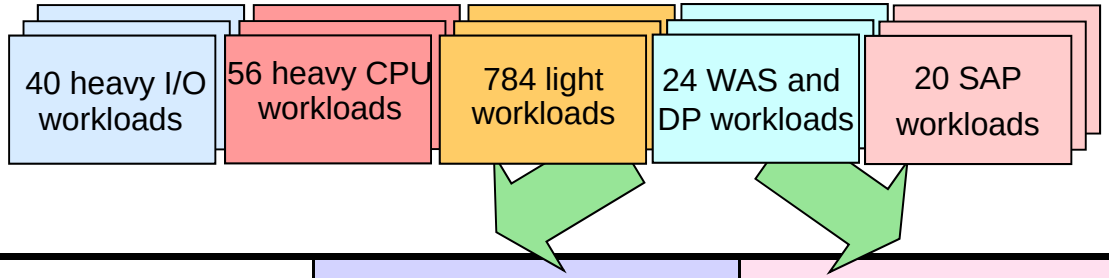

Three Year Cost Of	Deployed on Intel	Best fit on zEnterprise
Servers	$46.0M	$26.1M
Network	$0.45M	$0.03M
Power	$0.33M	$0.14M
Labor	$9.02M	$6.09M
Storage	$8.58M	$4.6M
Total	**$64.38M**	**$36.96M**
Total cost per workload	**$70K**	**$40K**

43% less

Results may vary based on customer workload profiles/characteristics. Prices are in US currency. Prices will vary by country

© 2012 IBM Corporation

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**

© 2012 IBM Corporation

Thank you

© 2012 IBM Corporation

Surveys Confirm Mainframes Are Lowest Cost For Core Business Workloads

Industry	Measure	Average IT Cost of Goods	Mainframe Biased	Server Biased	% Improvement
Bank	Per Teller Transaction	$0.31	$0.12	$0.35	-66%
Mortgage	Per Approved Loan	$263.67	$98.38	$290.80	-66%
Credit Card	Per Transaction	$0.16	$0.10	$0.18	-44%
Railroads	Per Ton Mile	$0.0014	$0.0012	$0.0018	-33%
Armed Service	Per Person	$8,036	$6,871	$9,839	-30%
Automotive	Per Vehicle	$333	$275	$370	-26%
Retail	Per Store (Door)	$494,818	$421,346	$560,300	-25%
Utilities	Per MegaWatt Hour	$2.63	$2.21	$2.94	-25%
Hospitals	Per Bed per Day	$64.30	$54.4	$71.7	-24%
Oil & Gas	Per Barrel of Oil	$2.10	$1.78	$2.32	-23%
Consulting	Per Consultant	$53,060	$48,900	$62,344	-22%
Trucking	Per Road Mile	$0.177	$0.155	$0.194	-20%
Airlines	Per Passenger Mile	$0.007	$0.0061	$0.0076	-20%
Chemicals	Per Patent	$57,717	$55,800	$59,552	-6%
Web Sites	Per Search	$0.042	$0.046	$0.041	12%

Most businesses running core workloads on mainframes had 6% to 66% lower IT costs per good than those using distributed servers

From Rubin Worldwide analysis of customer data and Gartner Research IT costs

© 2012 IBM Corporation

Mainframe Scales 2.5 to 15X Superdome
More Performance / Watt

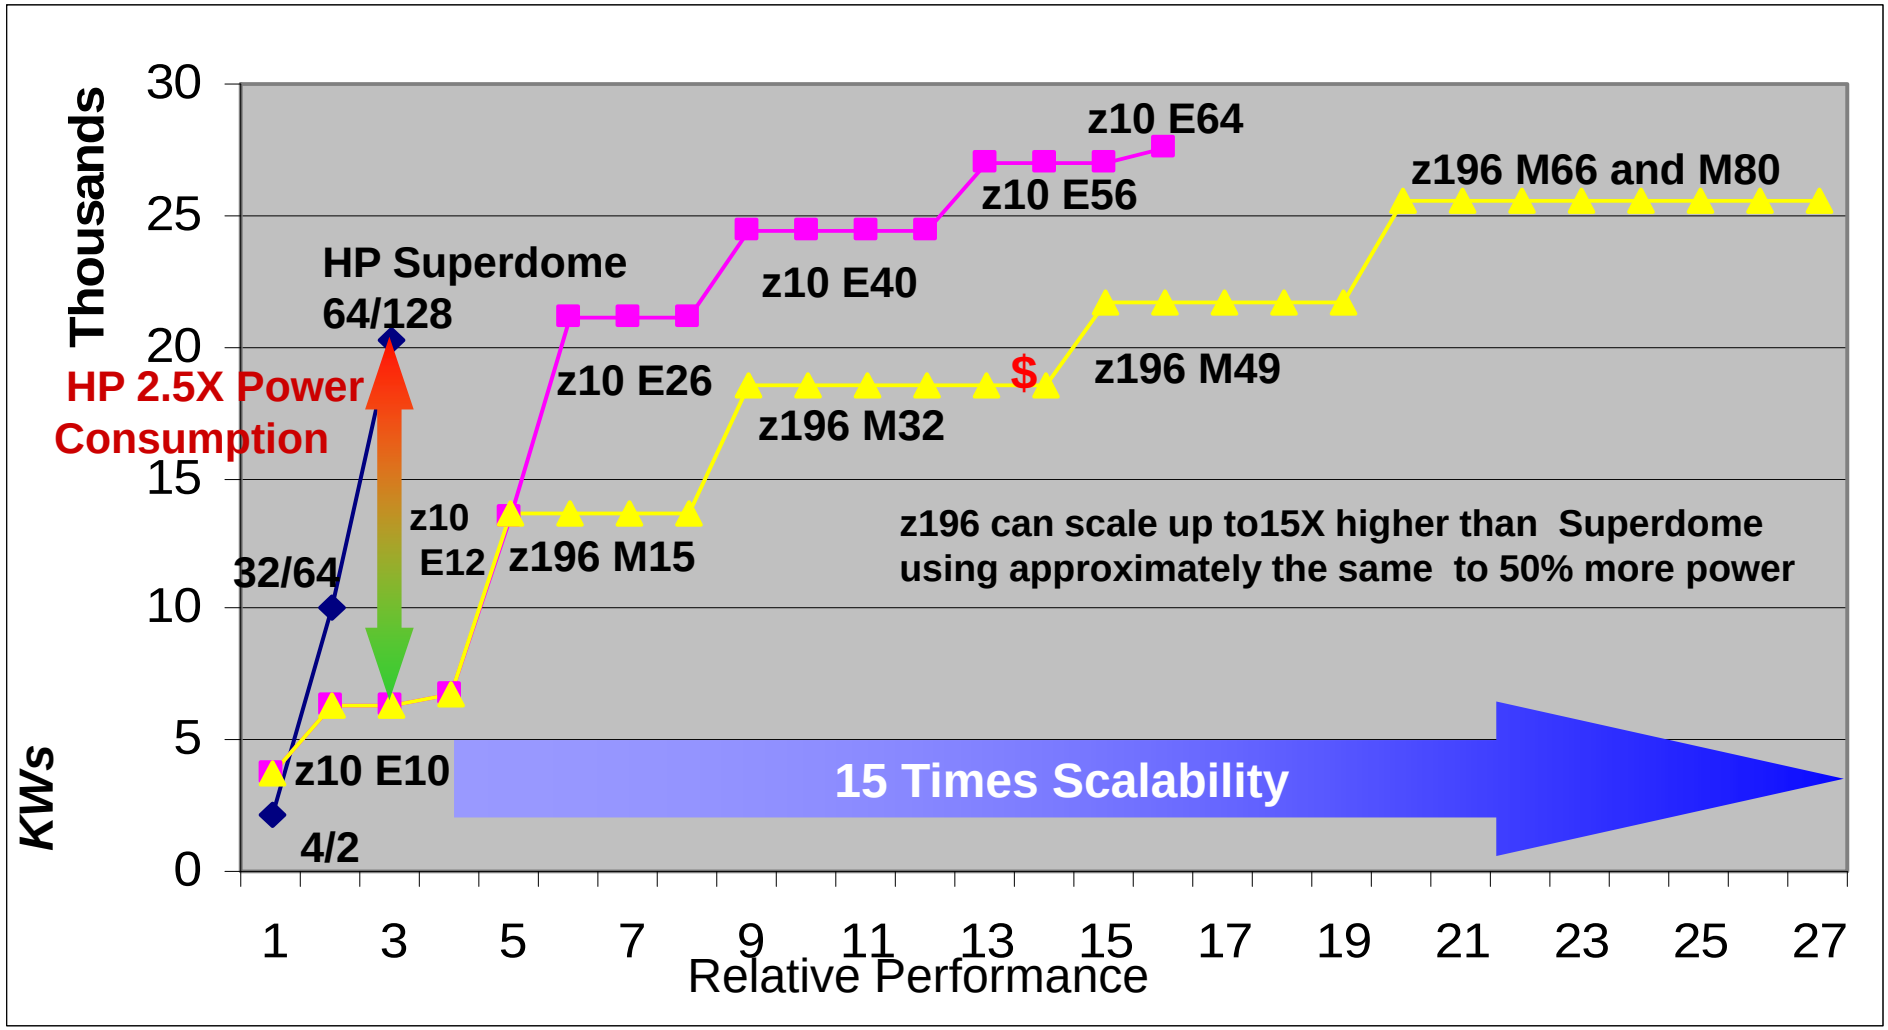

Notes: Performance as per Eagle TCO studies. Multiply by 2 for MIPS. HP performance based on 122 perf units / MIPS. z10 and z196 power is max value. It is very rare that any mainframe is even 80% of max. Typical mainframe power is less - approximately 60% of maximum as per field data. Mainframe Power scales by model or book package.

© 2012 IBM Corporation

z196 Capacity per Watt improvements

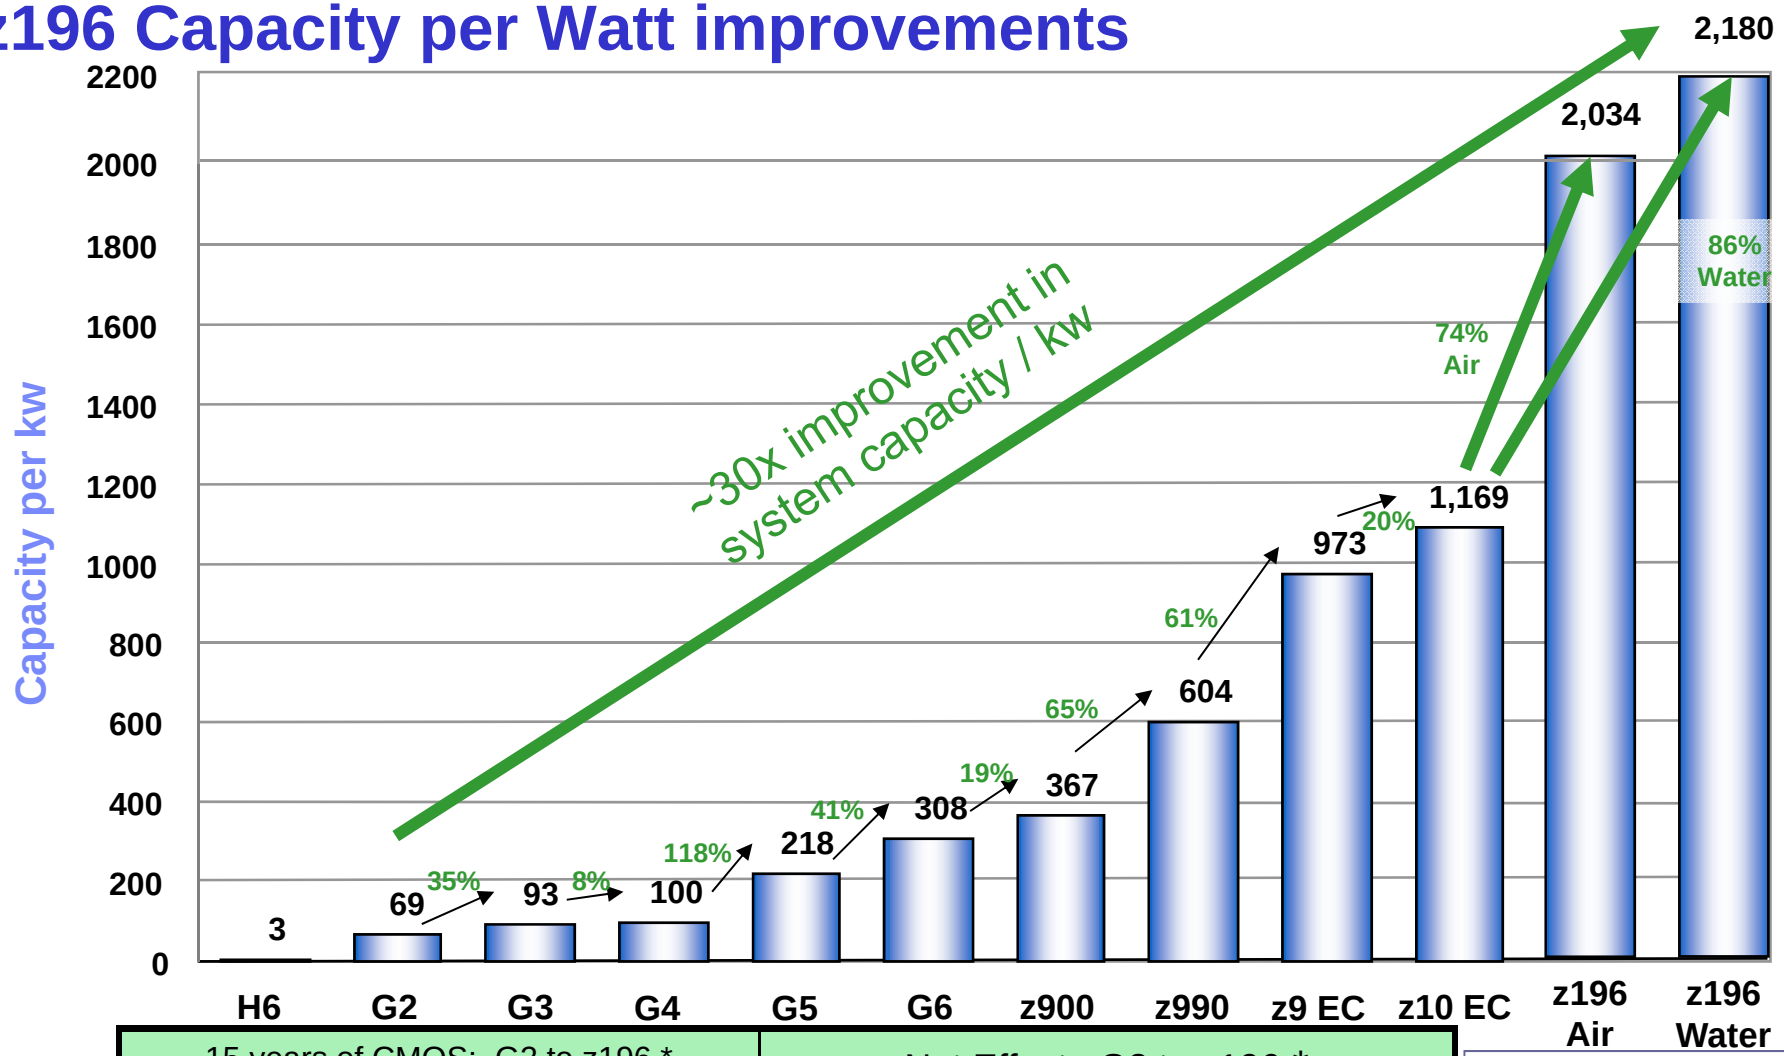

15 years of CMOS: G2 to z196 *		Net Effect: G2 to z196 *	
Power Increase:	17% **per year**	**Performance increased by:**	300x
Performance increase:	46% **per year**	Performance / kWatt increased by:	30x
Power density	13% **per year**	**Performance / sq ft increased by:**	190x

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.

z/OS Sysplex - Optimized For Efficient Clustering

- **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**

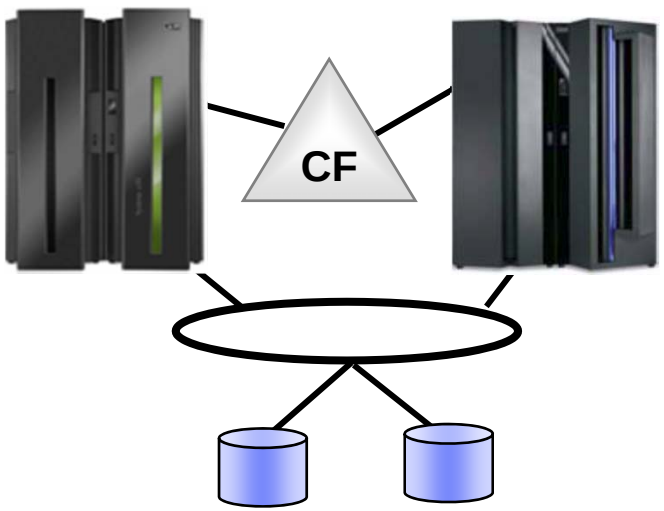

Data Sharing

A single 80-way zEnterprise delivers 52,286 transaction processing MIPs. Up to 32 of these can be clustered in a parallel sysplex, delivering ultimate scalability and availability.

© 2012 IBM Corporation