Leveraging the Mainframe Capabilities to Improve Total Cost of Ownership

The future runs on System z

Ray Jones
WW Vice President, z Software

Let's Break Down the Elements of Cost

Total Cost of Ownership =

Hardware/Maintenance

\+ IBM Software

\+ Environmentals

\+ Labor

\+ required Quality-of-Service

(Availability, Security, Disaster/Recovery…)

\+ other Elements

(Chargeback)

The total cost requires a total picture of your I/T assets and expenses

Mainframe Cost/Unit of Work Decreases as Workload Increases

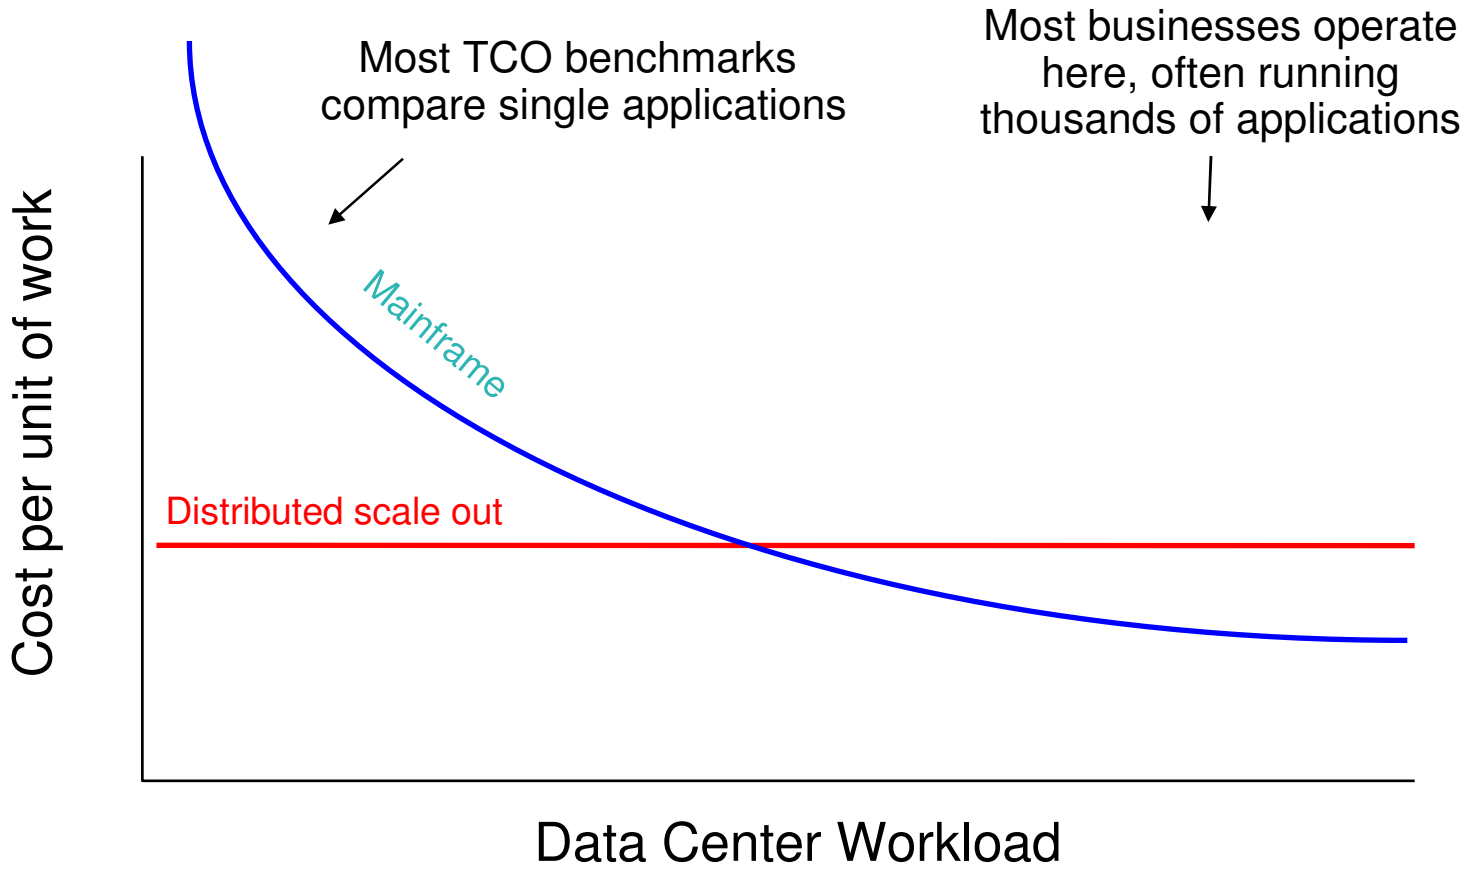

Observed Consolidation Ratios

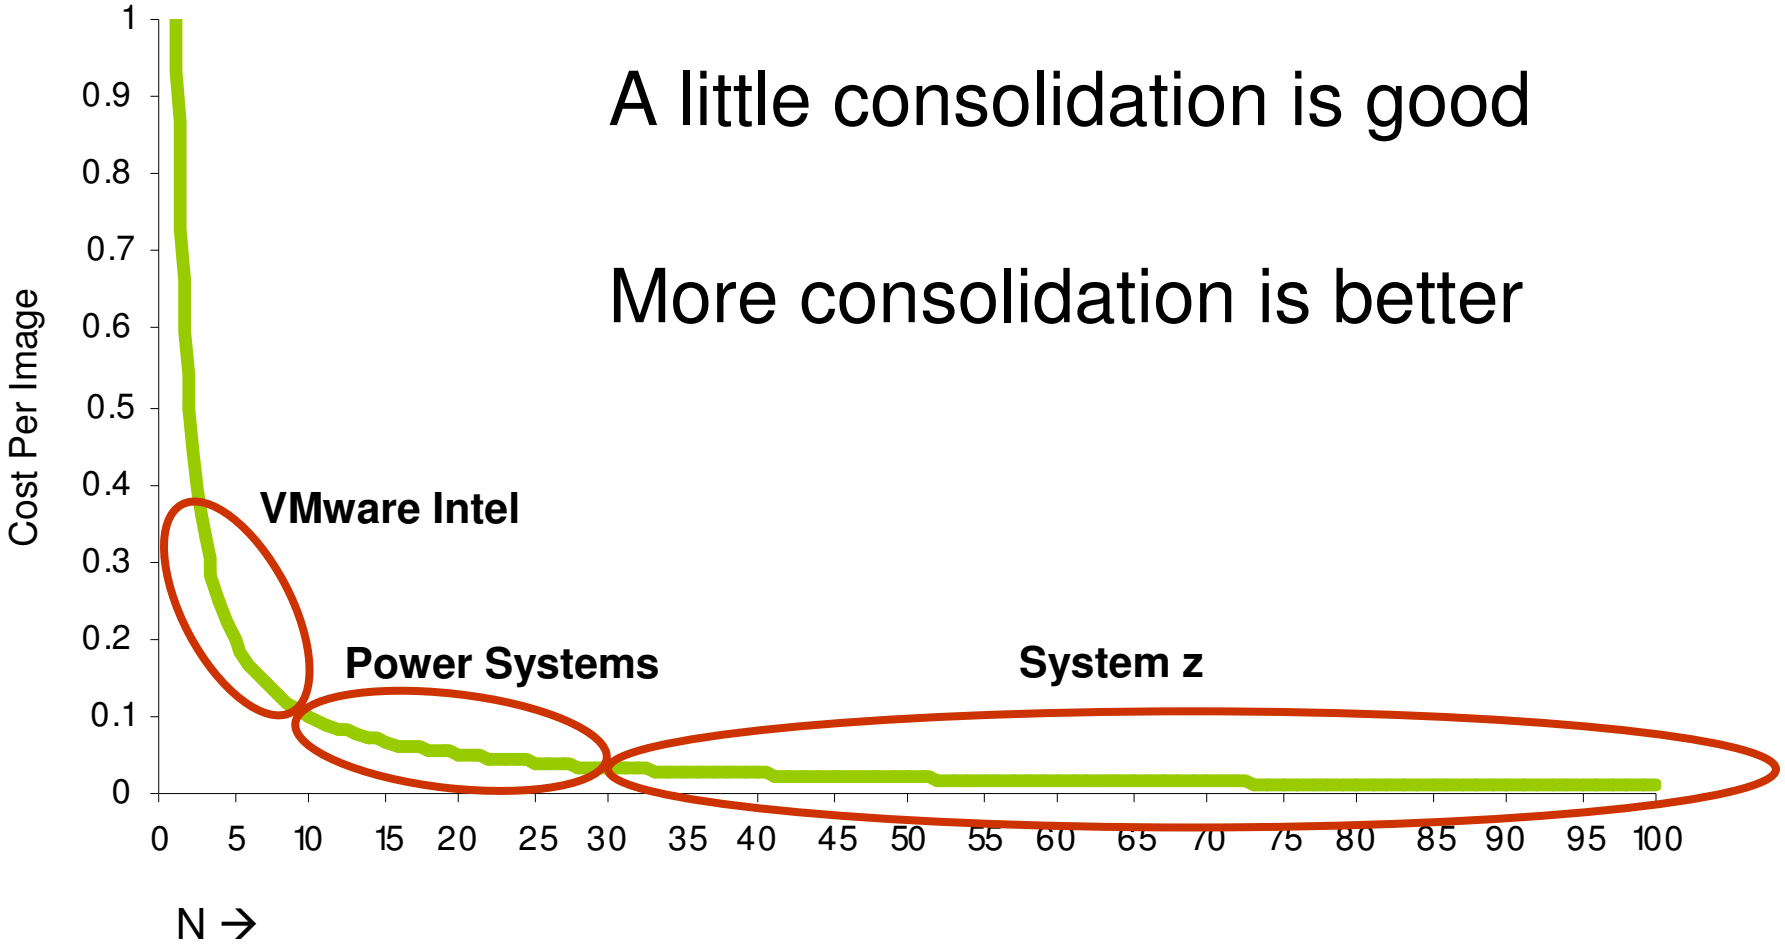

Utilization of Distributed Servers & Storage

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

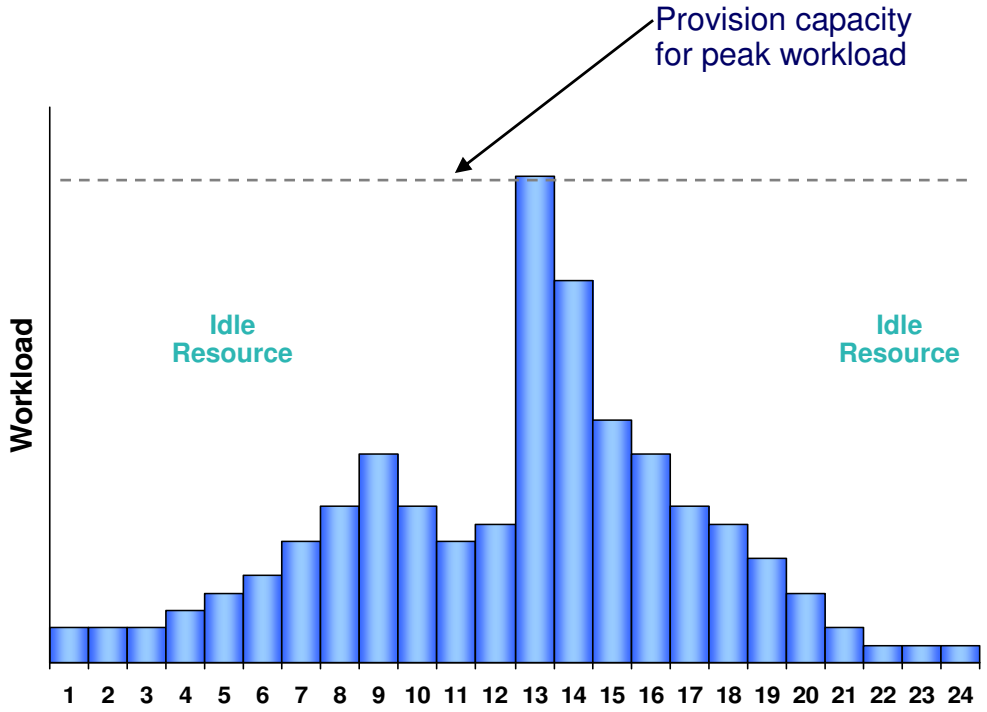

- **Storage Allocation**
  - Application-specific resulting in over-allocations
  - Fine grained storage allocation mechanisms characteristic of mainframe storage are uncommon in distributed environments.
- **Storage Utilization**
  - Single digit utilization for distributed environments is not uncommon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

Application specific storage allocations tend to occur in large units…

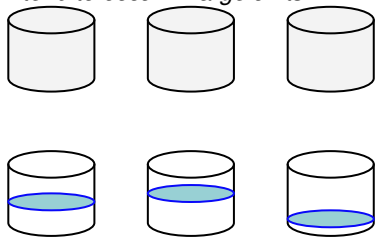

resulting typically in single digit utilization

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

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

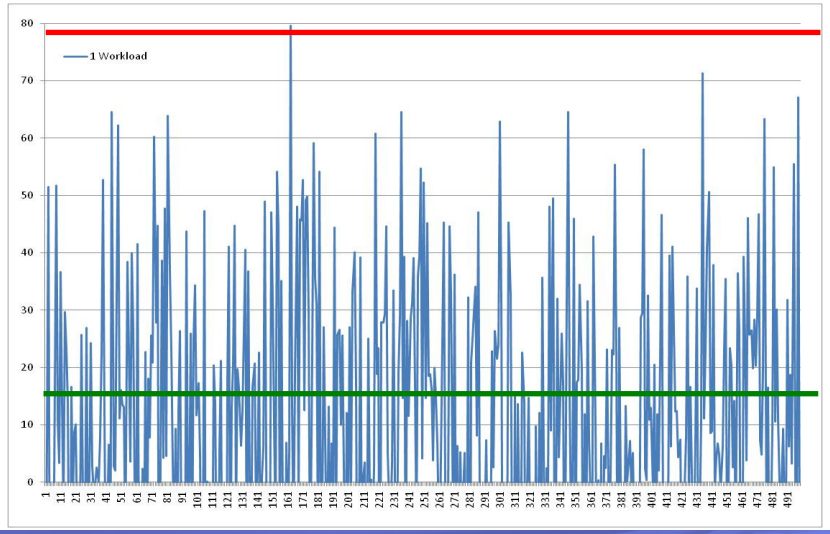

Adding a new workload to a pool of 256 existing workloads will require incremental processing capacity equal* to the **Mean** workload demand

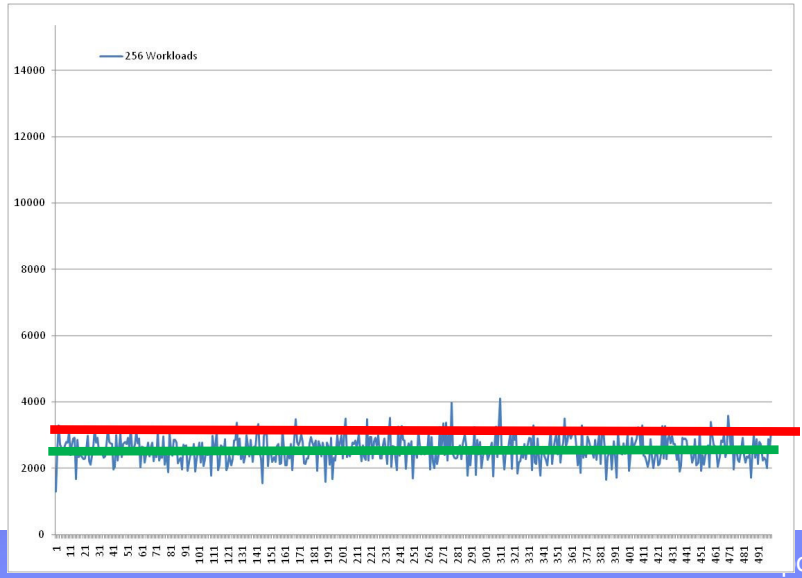

* If we add one more workload to a pool of 256 consolidated workloads the computing resource required for the pool goes up by 1.00047 * Mean

Compare The Processors Needed To Achieve 2,200 Transactions Per Second

Online Injector: 1 x HP RX7620

Temenos T24 Servers:
2 x HP RX7620
3 x HP 9000 Superdome

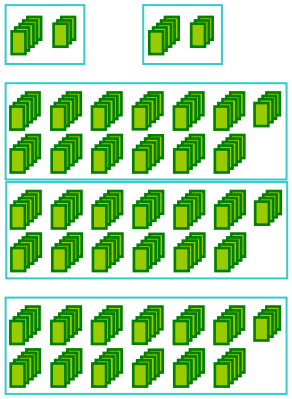

Oracle 10g: 1 x HP 9000 Superdome

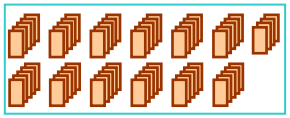

5 processors

(3,906 MIPS)

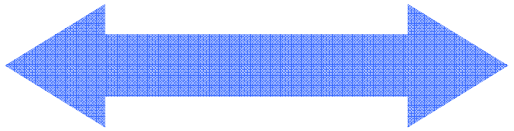

280 processors

(457,762 Performance Units)

$26.0M
TCA (3yr)

TCS BaNCS and DB2
1x z10 2097-705

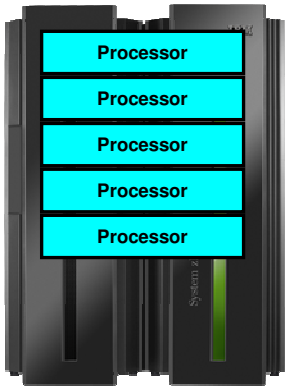

$18.9M
TCA (3yr)

117 Performance Units per MIP

HP Integrity rx7620 - (10U) 1.5GHz 6MB (8ch/8co)

HP 9000 Superdomes - 32W 1GHz 32MB (32ch/64co)

Compare The Processors Needed To Achieve 2,200 Transactions Per Second (with Dev/QA)

Online Injector: 2 x HP RX7620

Temenos T24 Servers:
4 x HP RX7620
6 x HP 9000 Superdome

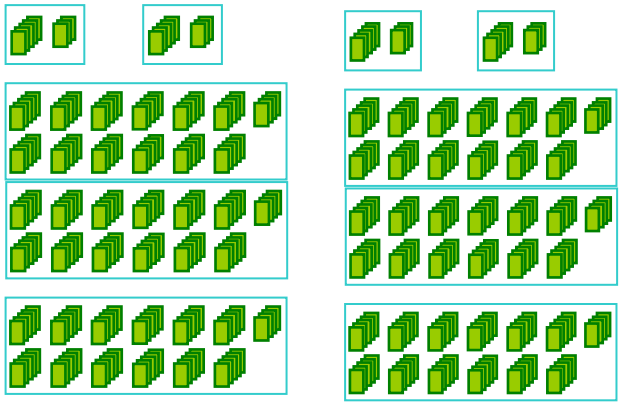

Oracle 10g: 2 x HP 9000 Superdome

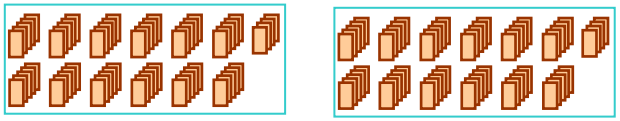

7 processors

(4,906 MIPS)

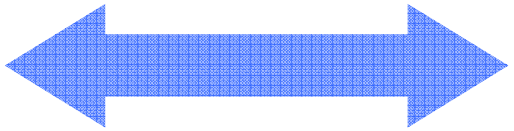

560 processors

(915,524 Performance Units)

$59.2M
TCA (3yr)

TCS BaNCS and DB2
1x z10 2097-707

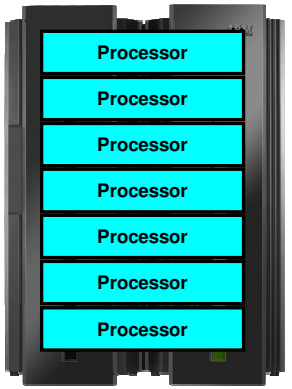

$22.7M
TCA (3yr)

187 Performance Units per MIP

NOTE: Double Distributed Servers, add 1000 MIPS to System z for Dev/QA

HP Integrity rx7620 - (10U) 1.5GHz 6MB (8ch/8co)

HP 9000 Superdomes - 32W 1GHz 32MB (32ch/64co)

Benchmark - Code Expansion When Moving From CICS/Cobol To Java On Wintel (Higher Is Worse)

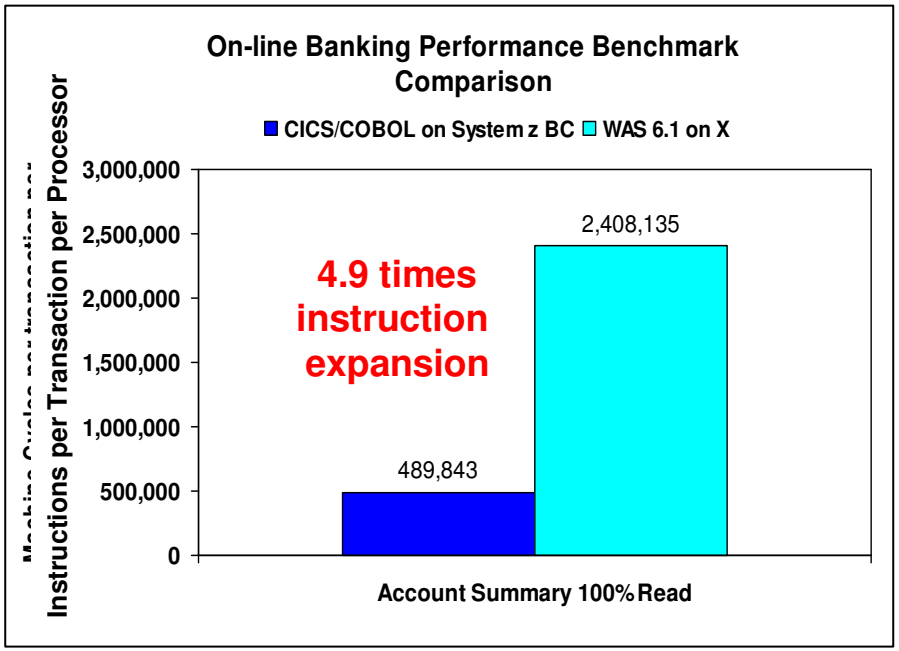

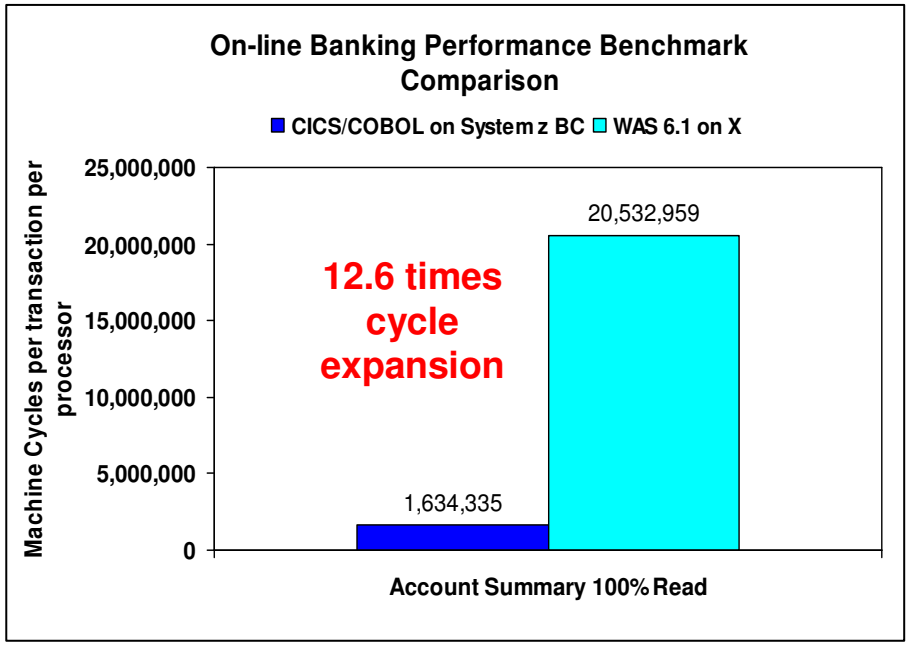

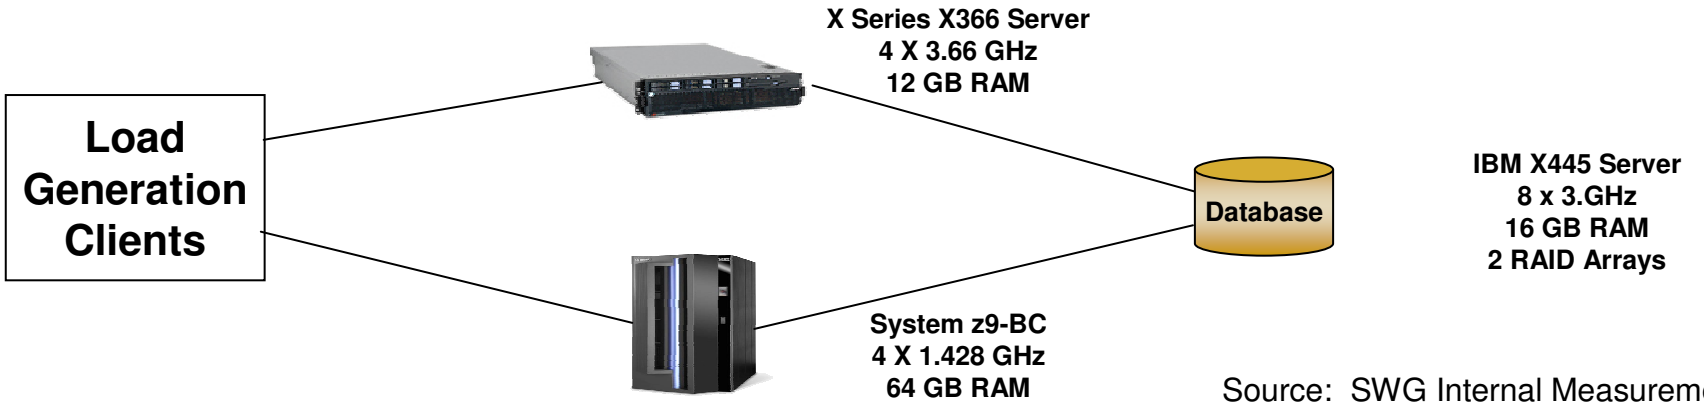

Source: SWG Internal Measurements

3. Reuse applications and data

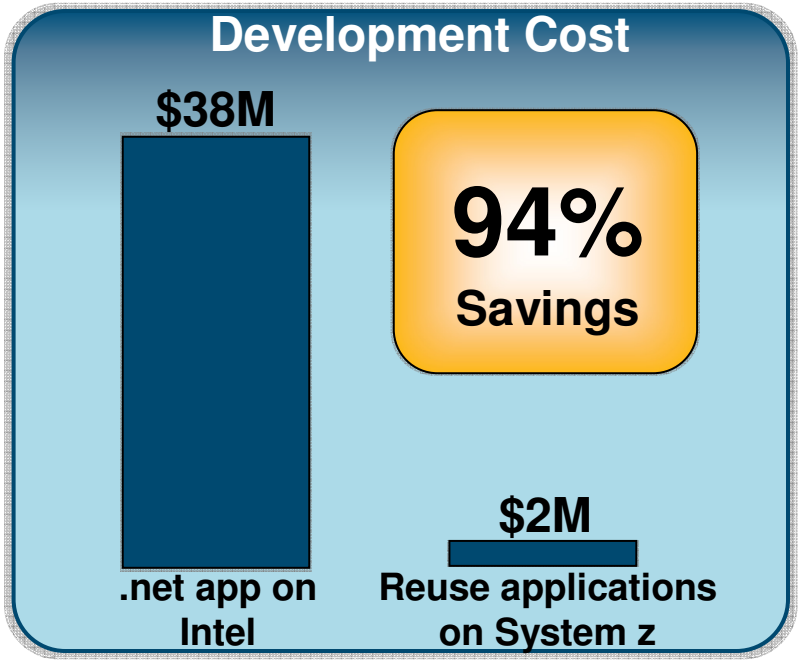

- **Complexity of recoding from scratch all the business processes into .net framework**
- **Speed of implementing System z solution was less than 29 days**
- **Additional employees to test and maintain .net application versus none for System z**

Additional benefits

Improved application functionality

Faster time to market

Quick implementation and reduced risk

System z Batch Processing Performance

- Bank of China BMT*
  - **IBM System z9**
  - **TCS BaNCS (Cobol)**
  - **380** Million Accounts
  - End of Day processing – 175M accounts finished in 55 minutes (52,970 accounts/second)

- HP/Temenos BMT**
  - **HP Itanium**
  - **Temenos T24 (Java)**
  - **13** Million Accounts
  - End of Day processing finished in 13.33 minutes (16,250 accounts/second)

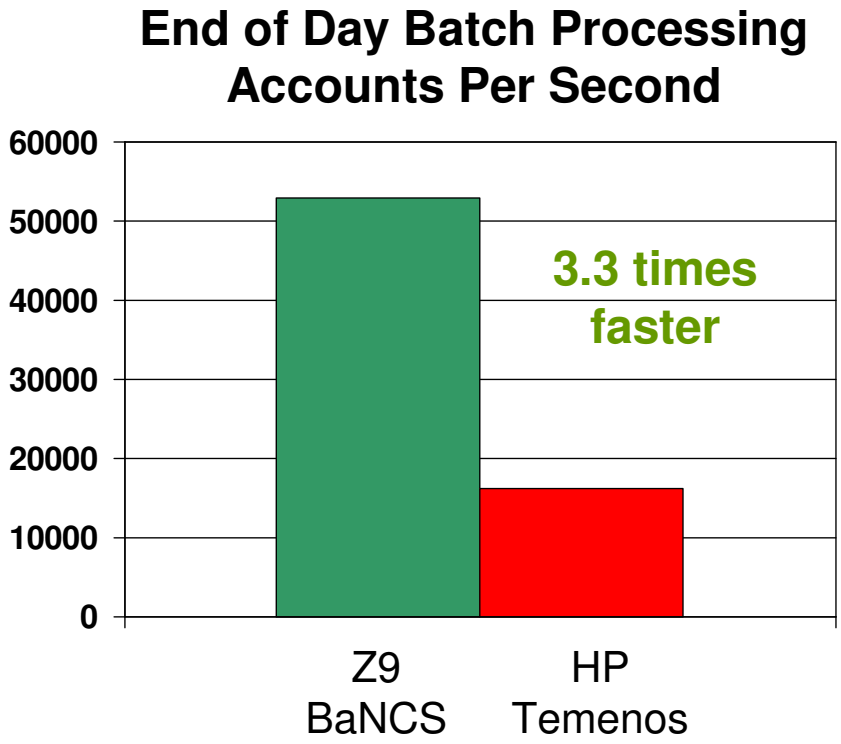

SOURCE:*http://www.enterprisenetworksandservers.com/monthly/art.php?2976 **Source**: InfoSizing FNS BANCS Scalability on IBM System z – Report Date: September 20, 2006
SOURCE:**TEMENOS BENCHMARKS; http://h71028.www7.hp.com/enterprise/downloads/TemenosBenchmark.pdf

Storage Costs: DB2 Delivers More Storage Savings Than Oracle

- **DB2 for z/OS lowers TCO by reducing storage needed**
  - TPC-H Benchmark: DB2 compression of 62% vs 27% for Oracle RAC

- **Storage savings with DB2 vs. Oracle for a 10 TB data base**

	Oracle	DB2 for z/OS*
Storage System	**HP XP24000 Storage**	**IBM System Storage DS8100**
Overall database compression ratio (using TPC-H benchmark results)	**27%**	**62%**
For 10 TB uncompressed data storage needed	**7.3 TB of HP Storage**	**3.8 TB of IBM Storage**
Cost of storage (3 year TCA)	**$888,399 + $37,560 x 3** = $1,001,079	**$192,205 + $7,992 x 2**** = $208,189
With compression, storage for DB2 costs 79% less than for Oracle		

*DB2 for z/OS achieves similar compression ratios to those of DB2 for LUW
**IBM storage maintenance fee for the first year is included in the warranty

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

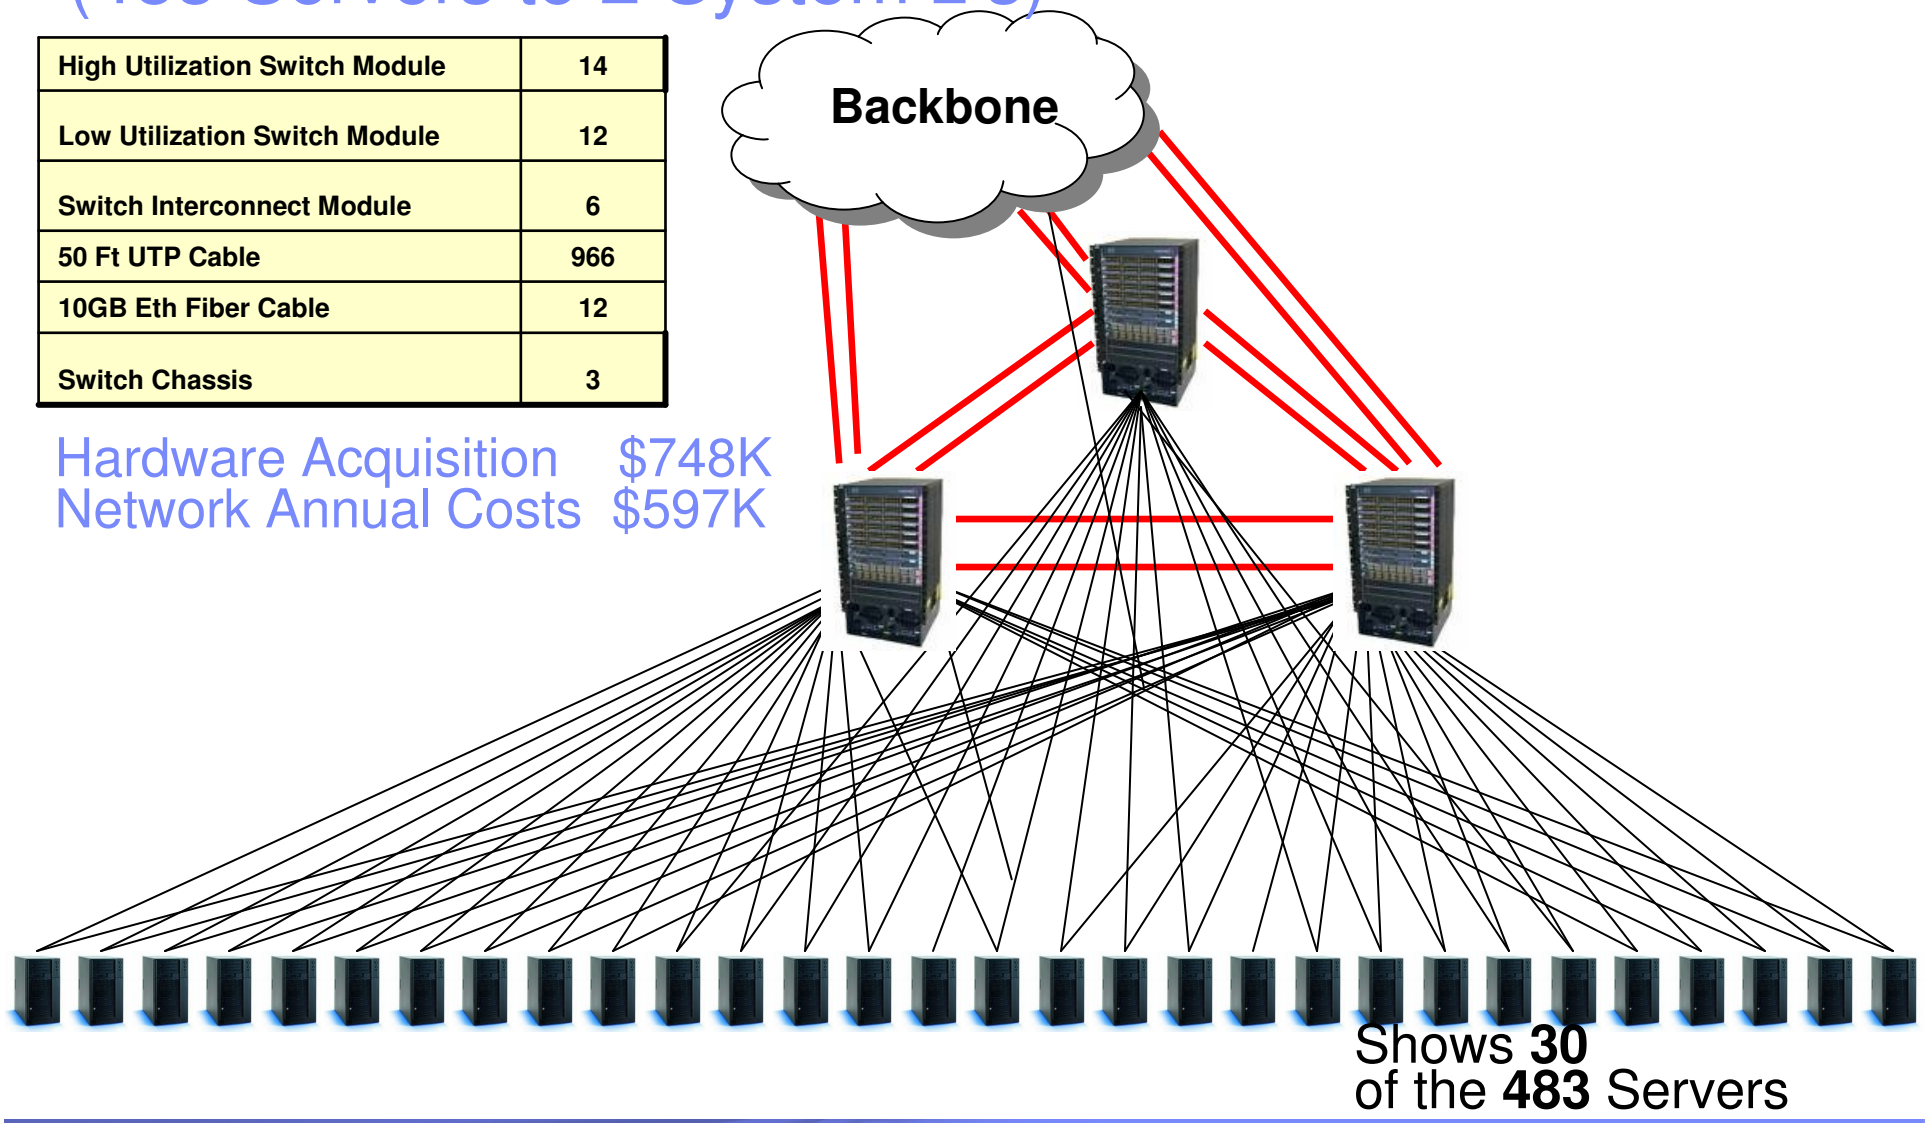

Case Study: Network Costs – After Consolidation (483 Servers to 2 System z's)

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

"After" Network Annual Cost $253K

Network Annual Cost Savings $344K

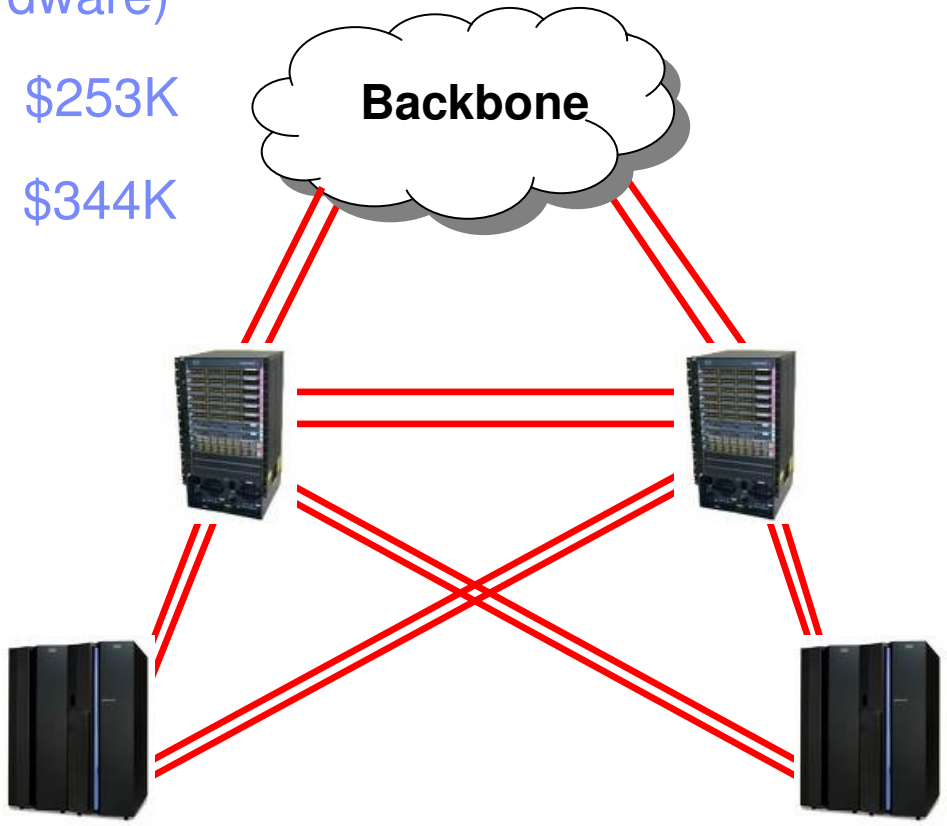

Specialty Engines Reduce Cost For New Workloads

- **Special assist processors for System z**
  - For Java workloads (zAAP)
  - For selected DB2 workloads (zIIP)
  - For Linux workloads (IFL)

- **Attractive pricing**
  - $125K for a 920 MIP processor (**90% discount**)
  - No charge for IBM software running on zAAP/zIIP
  - IBM software running on IFL costs 120 PVU's
  - Free upgrade to next generation!

- **Requirements**
  - Max number of zAAP =< number of general purpose processors
  - Max number of zIIP =< number of general purpose processors
  - No limit on the number of IFL's

International Restaurant Chain Avoids High Cost Software

- **Existing environment of 1600 MIPS included high cost ISV system management software**
- **Competitor's proposal was only a partial offload**
  - Complete offload projected to cost 2.3x more
  - $56M vs $24M over 5 years
- **System management software costs more in the offload case**
  - Mainframe systems management
    - $2.0M Stream per year (48 products, mostly third party)
  - Distributed systems management
    - $2.6M Yearly Maintenance (26 products)
    - $13.3M One Time Charge
- **Better: Replace higher cost System z ISV software with lower cost IBM Software**

Email, Calendaring, and Collaborative Application on System z is 1/3 the Cost of x86 and Saves $8M+ over 3 years

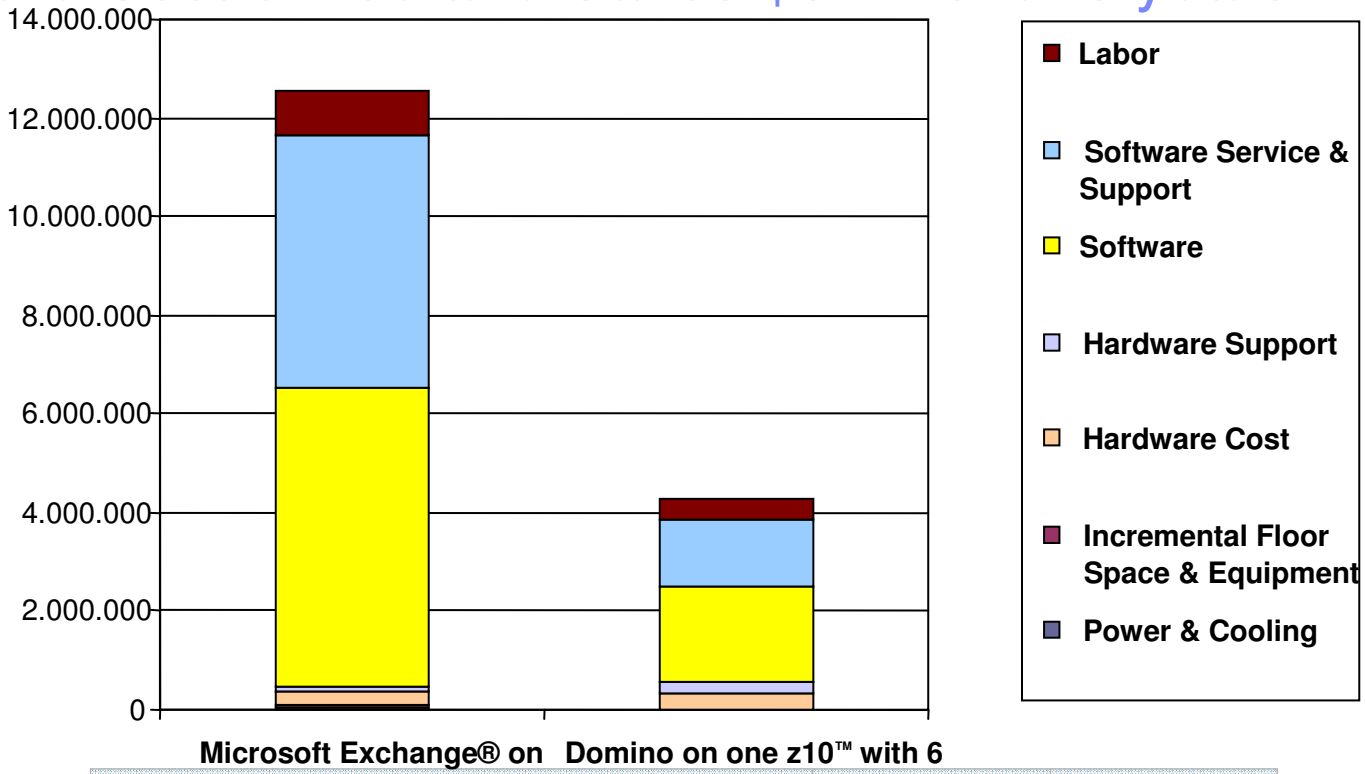

	TCO: 3 Years	Per User Cost
Microsoft Exchange on fourteen x86 Servers	$ 12,557,473	$ 1,046
Domino on one z10 BC™ with 6 IFLs	$ 4,286,997	$ 357
Savings with Domino on System z Linux	$ 8,270,476	$ 689

Assumes 12,000 users

z10 Consumes Less Power Than Superdome

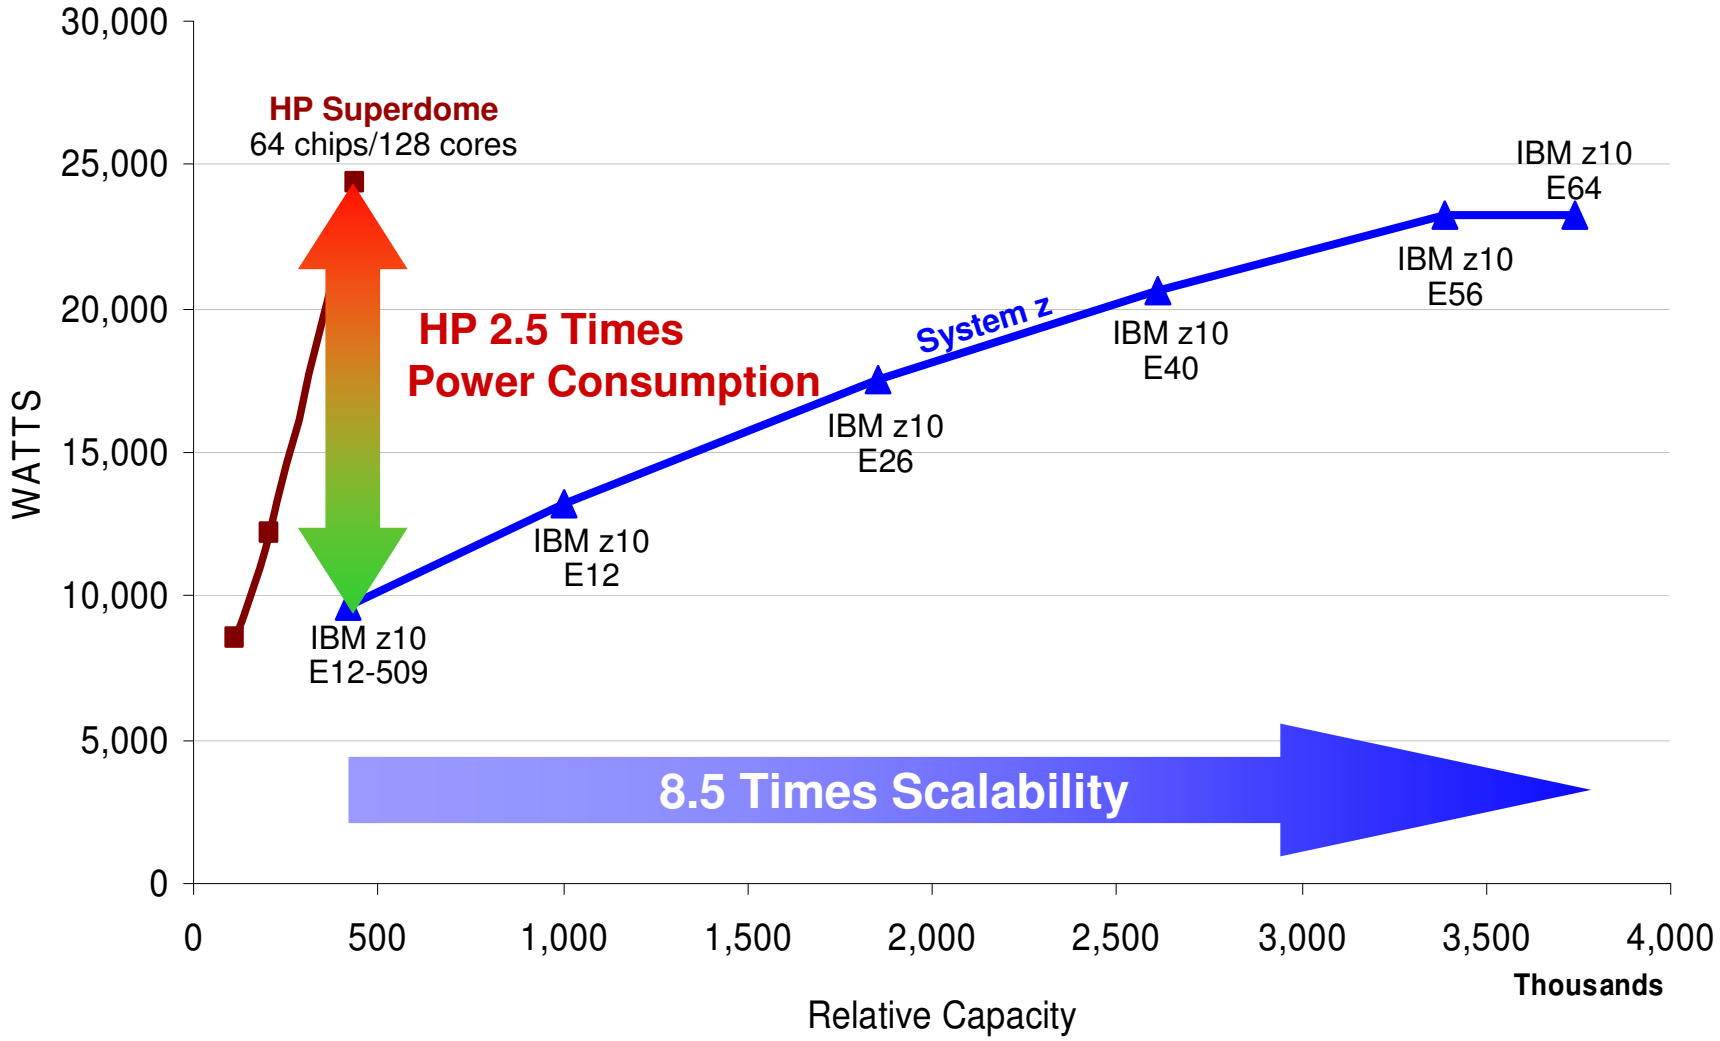

Do The Math

- **HP Itanium 2 Superdome 9050 (64ch/128co)* consumes a maximum of 24,392 watts**
  - [24,392 X $.10 X (24 X 365)]/1000 = $21,367 per year for electricity
- **Mainframe with similar computing capacity - a System z10 704 machine with 2 I/O cages using 13.26 kW (rated)***
  - $11,615 per year for electricity
- **Similar savings on cooling capacity**
  - Cost of cooling is about 60% additional
  - Superdome total $34,187 per year vs. Mainframe $18,585
  - Savings of mainframe power and cooling is $15,602 per year

* Performance equivalence determined by IBM TCO study

IBM Storage Also Saves Energy Costs

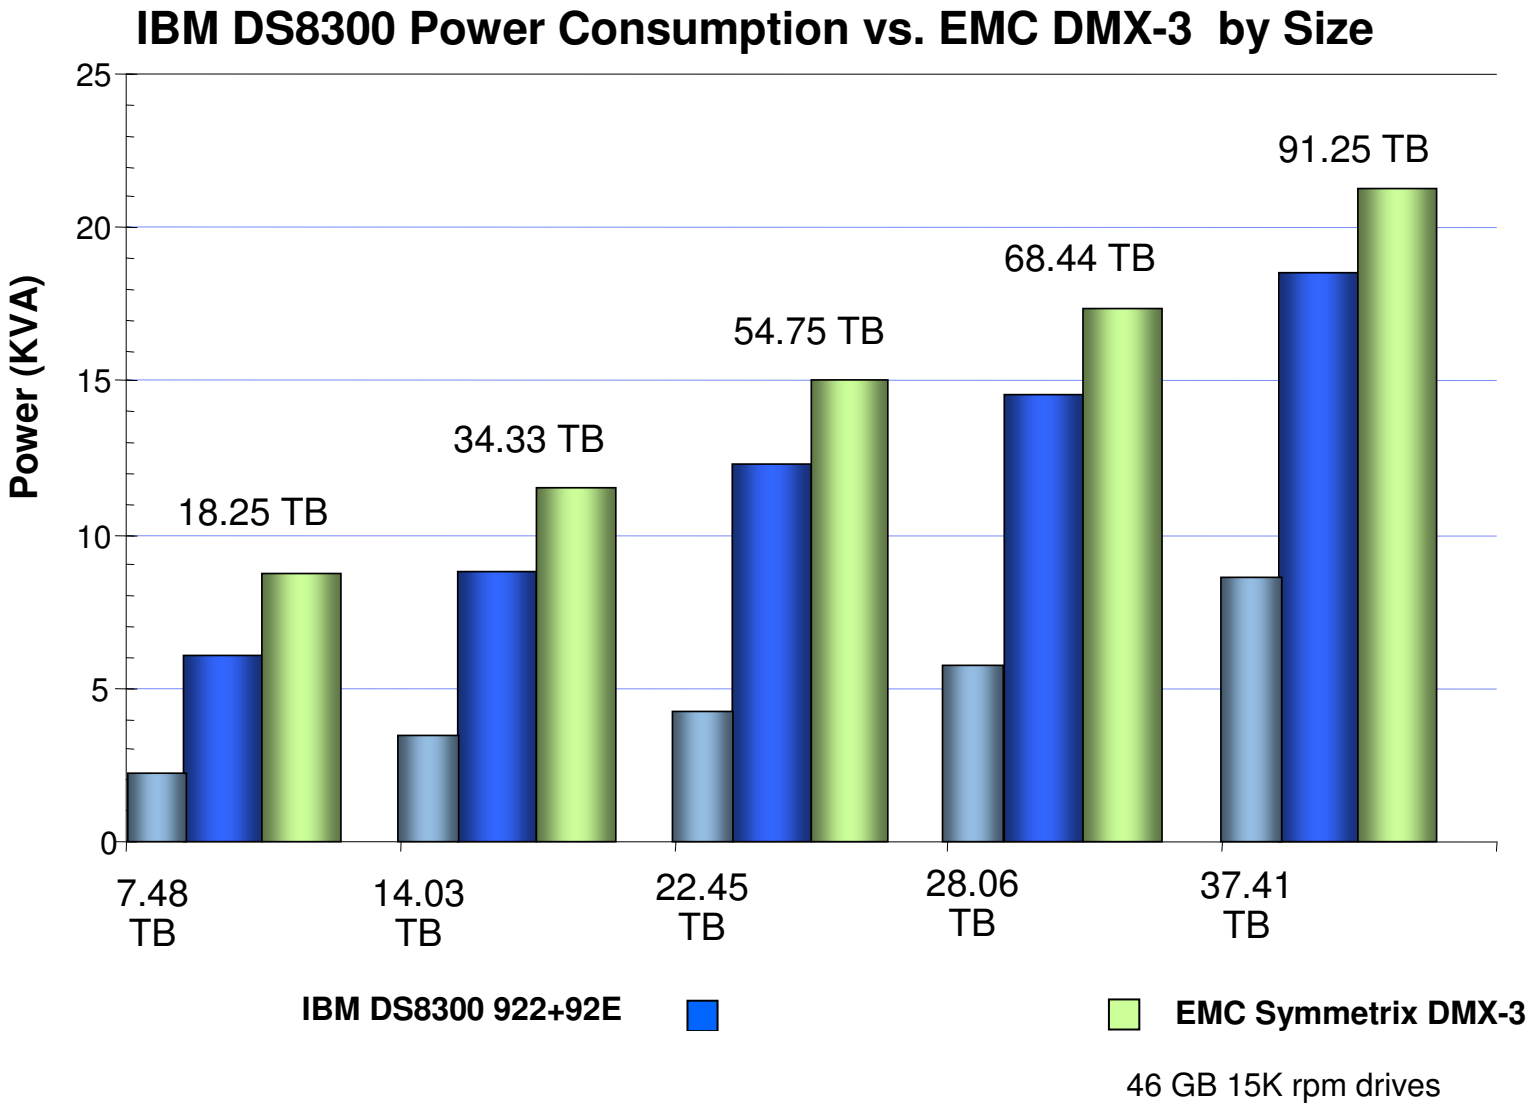

The Mainframe Also Delivers More Compute Power Per Footprint Unit

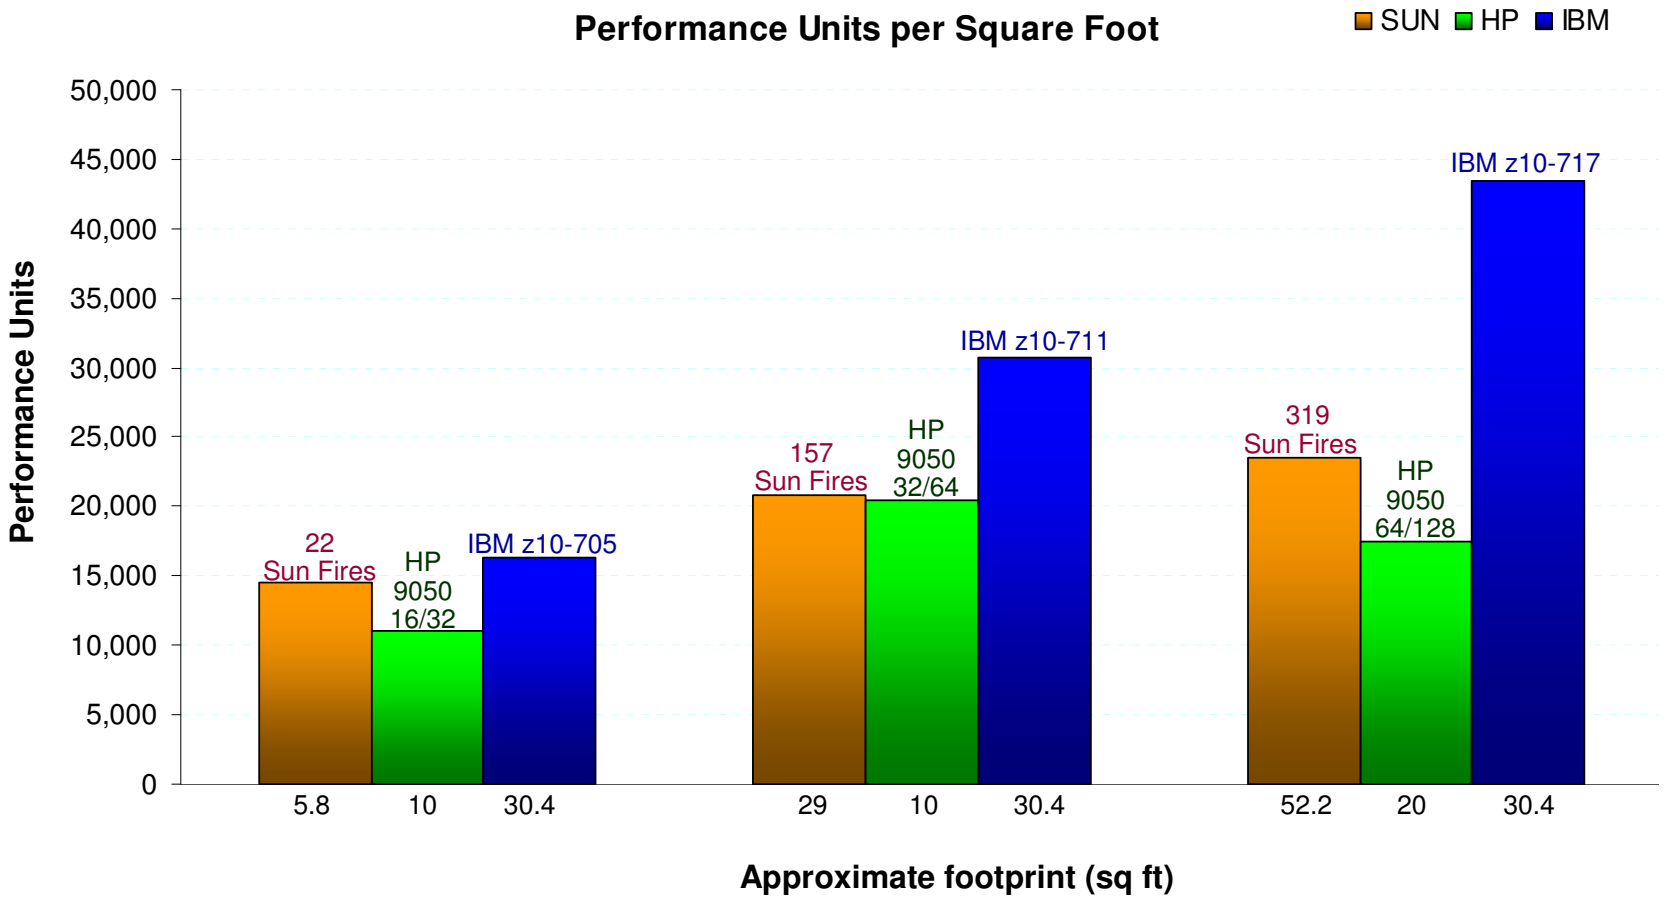

Based on 122 performance units per MIP

Mainframe footprint remains constant

Customer Survey – How Many People to Manage Servers?

# NT Servers	# People	Ratio (s/p)
1123	68	**16.5**
228	20	**14.4**
671	51	**13.1**
700	65	11.5
154	18	8.5
431	61	7.1
1460	304	4.8
293	79	**3.7**
132	54	**2.0**

# UNIX Servers	# People	Ratio (s/p)
706	99	**7.1**
273	52	5.2
69	15	4.6
187	56	3.3
170	51	3.3
85	28	**3.0**
82	32	**2.6**
349	134	**2.6**
117	50	**2.3**
52	52	**1.0**

Mainframe administration productivity surveys range 167-625 MIPS per headcount (500 is typical), so…

Source: IBM Scorpion Customer Studies NOTE: Figures for total administration cost

Manage More Workload Per Headcount

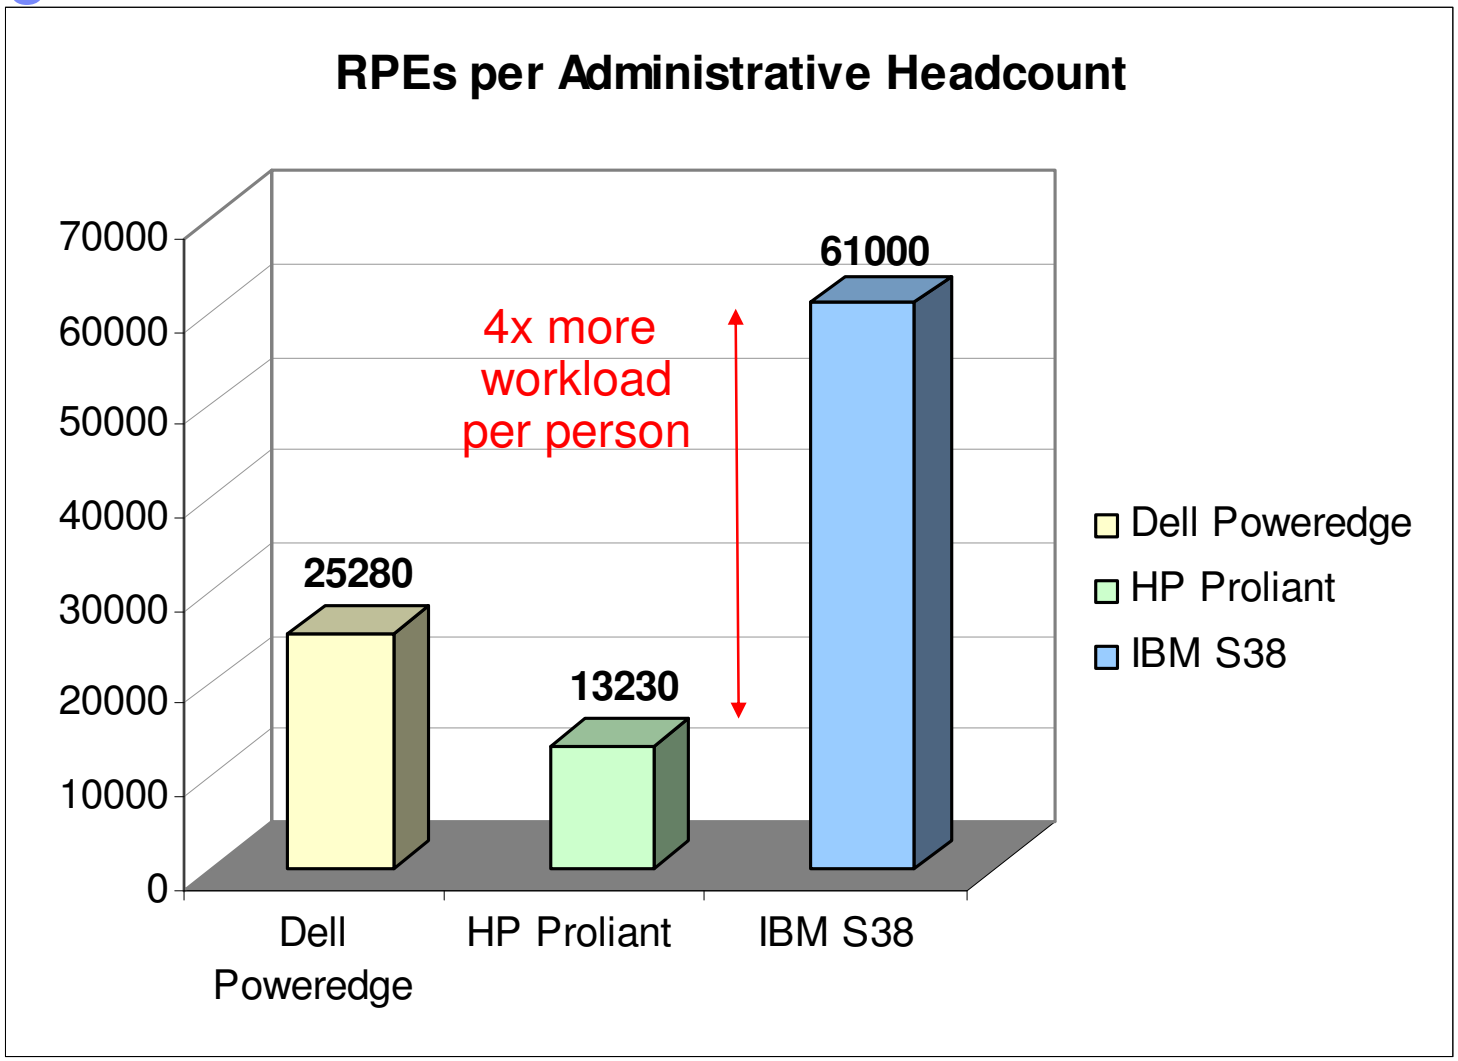

Compared at 122 RPE's = 1 MIP

Understand The Cost Components

Annual Operations Cost Per Server
(Averaged over 3917 Distributed Servers)

Power	**$731**
Floor Space	**$987**
Annual Server Maintenance	**$777**
Annual connectivity Maintenance	**$213**
Annual Disk Maintenance	**$203**
Annual Software support	**$10,153**
Annual Enterprise Network	**$1,024**
Annual Sysadmin	**$20,359**
Total Annual Costs	**$34,447**

The largest cost component was labor for administration
7.8 servers per headcount @ $159,800/yr/headcount

Source: IBM internal study

How Does Consolidation Reduce Costs?

- Costs shared by all “N” consolidated images
  - Hardware
  - Software
  - Power
  - Floor Space
  - Local Network Connectivity

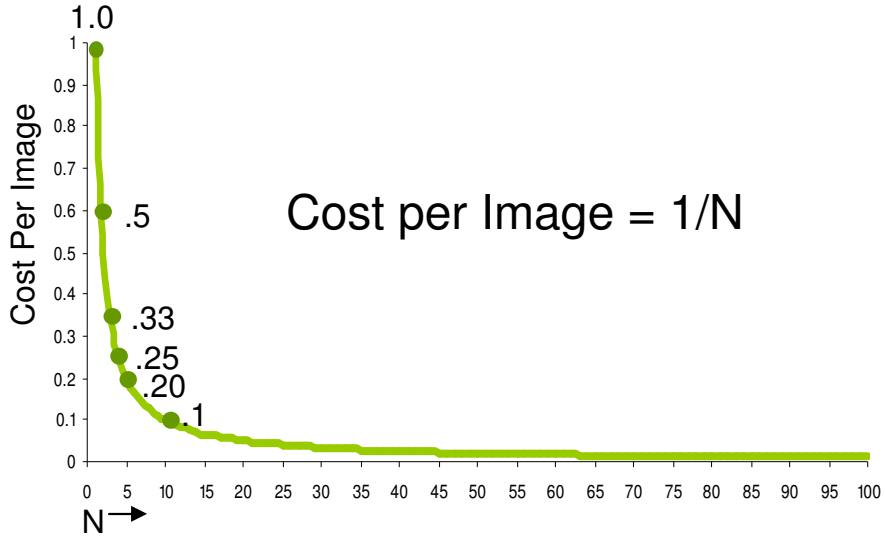

Cost per Image = 1/N

- Costs not shared by consolidated images
  - Migration cost per image — Fixed cost per image
  - Off premise network cost
  - Labor cost per image — Fixed cost per image, but typically less than unconsolidated labor cost

The more workloads you can consolidate, the lower the cost per image

26

Optimize deployment of applications and data

Deploying SAP database and application servers

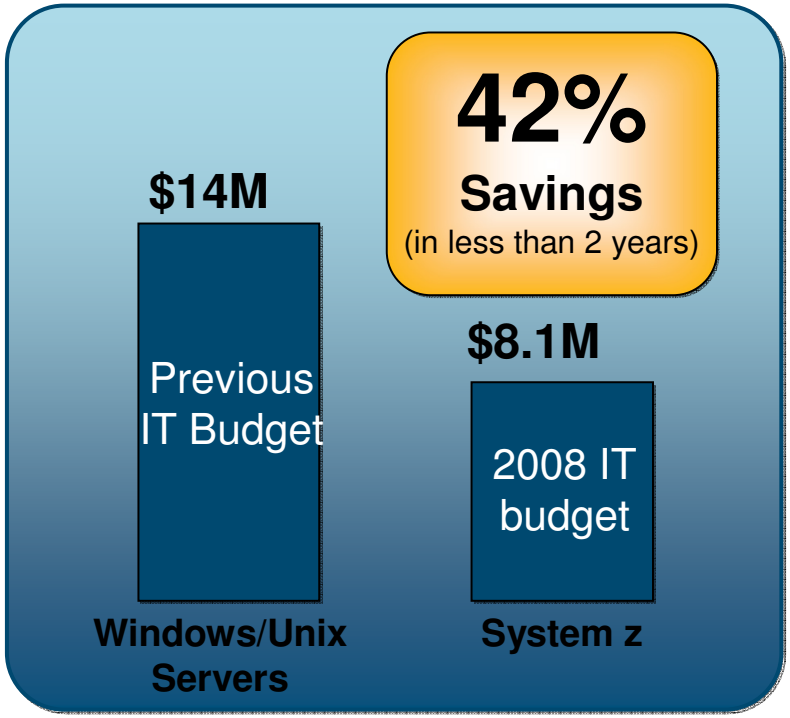

Top three reasons for savings

- Software and hardware licensing costs dramatically reduced
- Software and hardware maintenance costs are significantly down
- Networking costs plunged, while infrastructure was drastically simplified

BALDOR $1.8 billion Electric motors manufacturer

Expected Benefits Realized: Availability and Performance

The System z decision was driven by expected benefits:

- **Reduced complexity**
- **High availability**
- **Ease of maintenance**
- **Dynamic Workload**
- **Good consistent application response time (SAP)**
- **zLinux for rich toolset, ease of use**

Additional Benefits Realized: Significant Cost Savings

- **+Reduced IT budget by 42% - in less than 2 years**
- **+Reduced floor space by 70%**
- **+Reduced software and hardware maintenance by more than 50%**
- **+Reduced power consumption by more than 60%**
- **+Reduced total TCO from 2% of sales to below 1% - and realized 1 year ahead of schedule**

Key Points:

Mainframe Costs	Distributed Costs
The cost of running incremental workload on the mainframe goes down as the total workload grows	The cost of running additional workload on distributed servers goes up more linearly
– Labor costs hold steady as workload grows	– Labor is now the highest cost element in distributed environments Administrative staff costs increase in proportion to the number of servers
– IBM pricing policies designed to favor the addition of more workload	– New workload requires additional servers and licenses
– Highly Efficient Power and Cooling – Small Footprint	– Energy and Space cost is more linear
– Lower software costs per transaction as workload grows – and PRA can lower ISV tool costs	– Cost of software licenses is more linear
– High Availability and Security Translate into low cost	– Fractionally less Availability and Security can drive Significant downstream costs
Customers have learned that mainframes deliver economies of scale, especially as the workload grows	**Result – scale out strategies do not deliver equivalent economies of scale as the workload grows**

This pricing discussion uses published list prices

IBM Software Price Per Transaction is Going Down

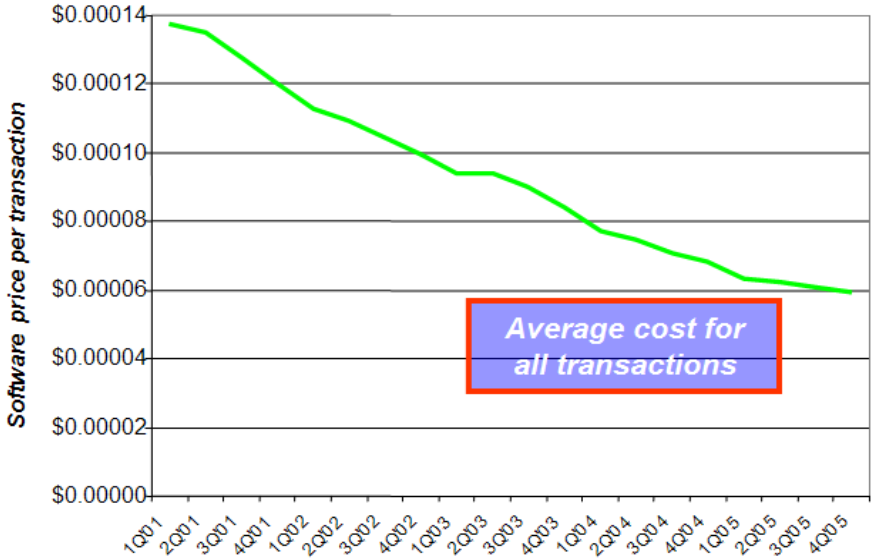

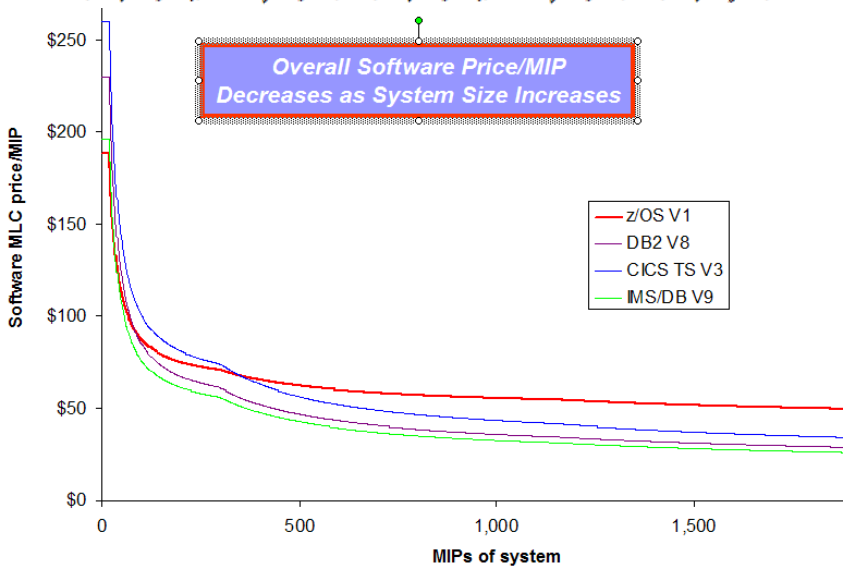

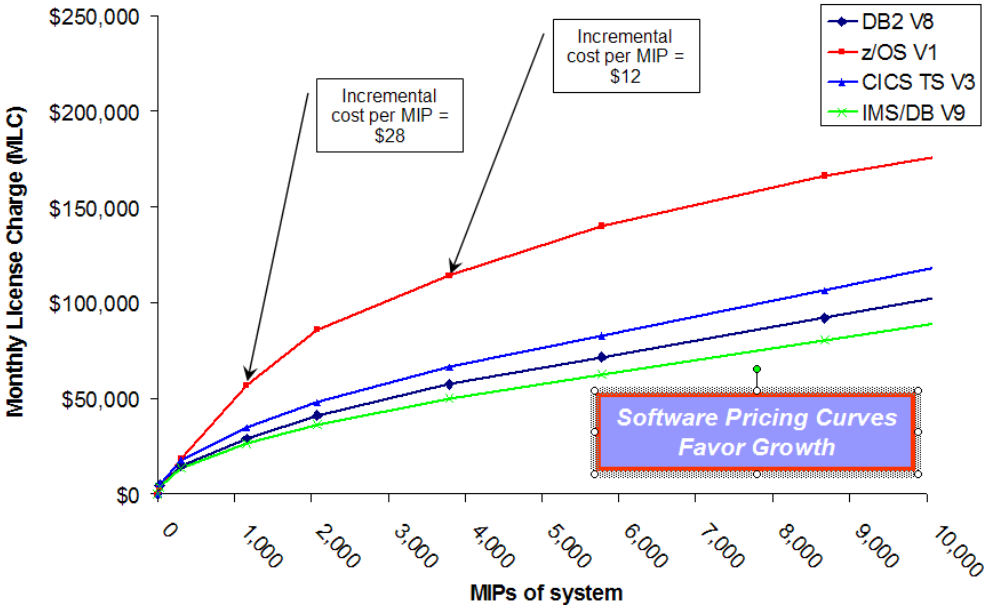

Putting This in Perspective

- **For a typical system of 1,400 MIPS, MLC software stack costs $59 per incremental MIP**
- **If a transaction is 1 million instructions, an incremental MIP can perform >2½ million additional transactions per month for Δ$59 software cost (44K transactions per dollar)**
- If these are credit card transactions **of average $100 with a commission of 2%, the business makes $5.2M per month for a software cost of $59 per month (88,000 times return)**
- If this is a bank account **averaging 3 transactions a day, the business can do 40 years of account management for a software cost of $1**

Incremental Cost Of New Workload Can Be Isolated Using Sub-Capacity Pricing…

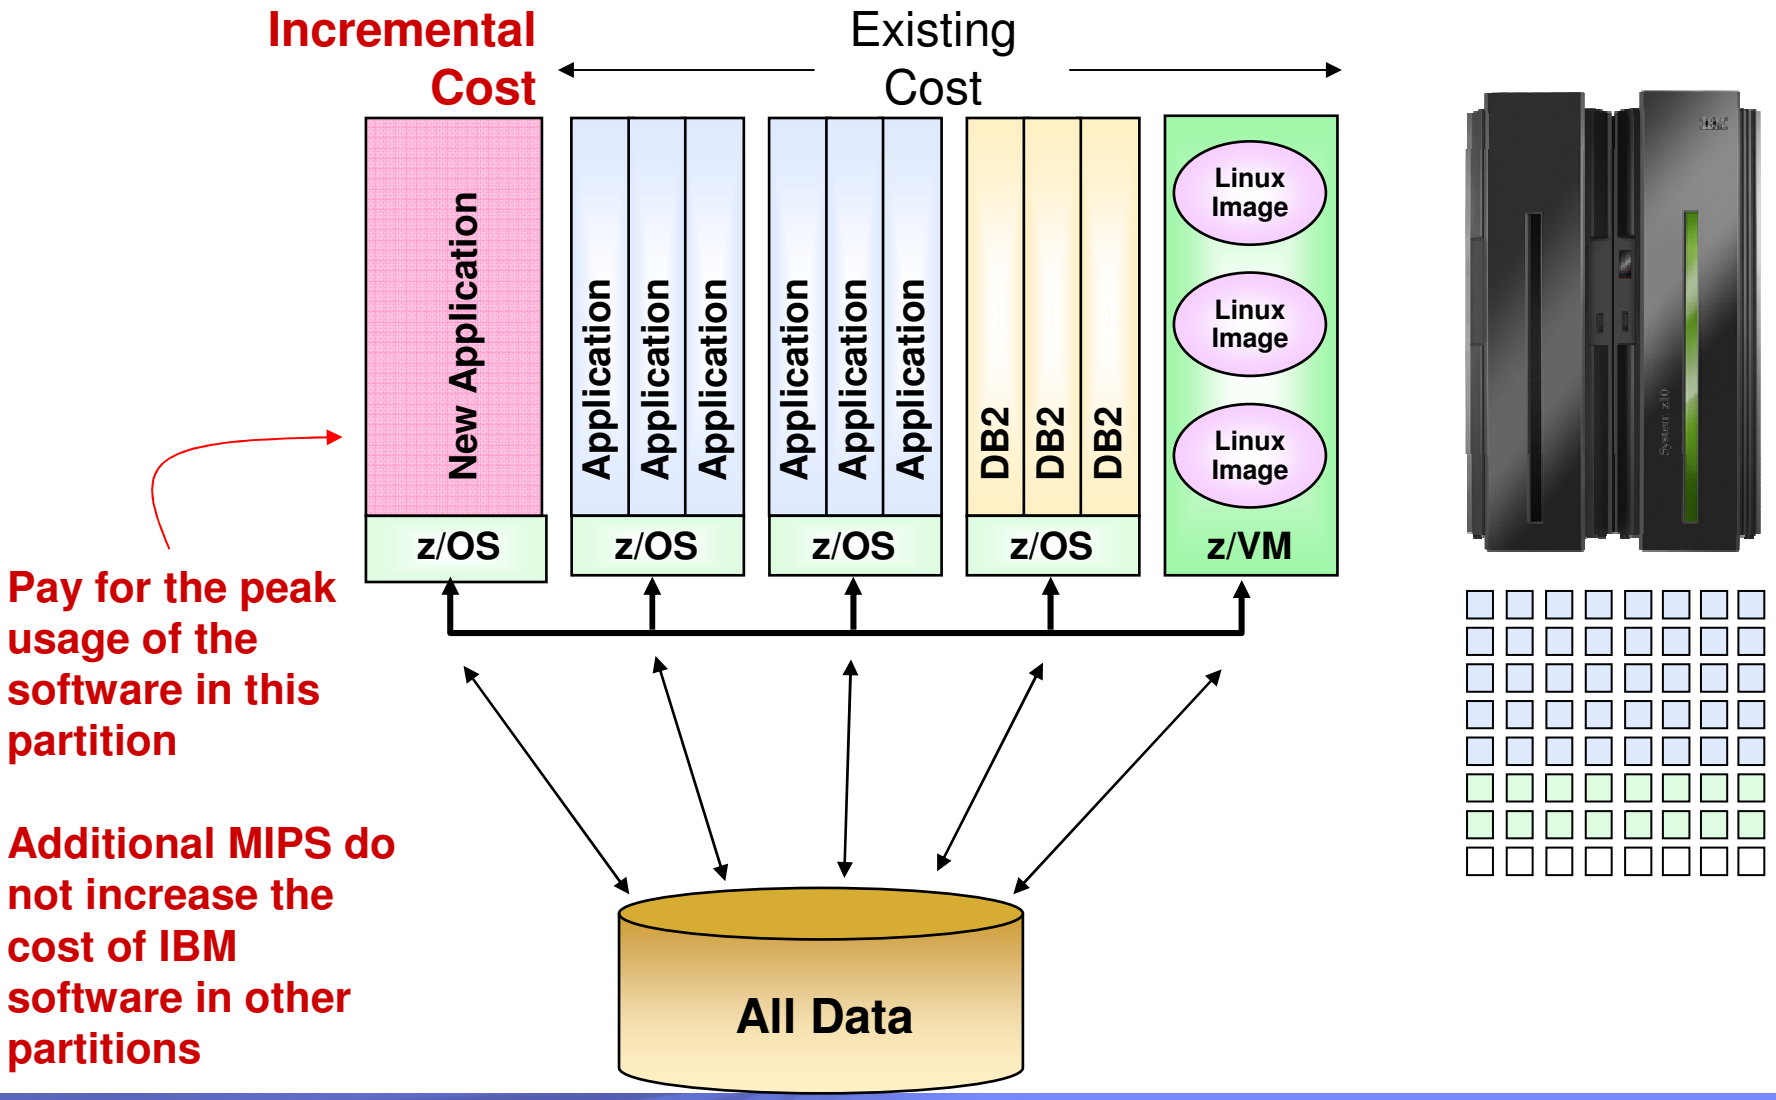

…Or On zLinux With IFL Pricing

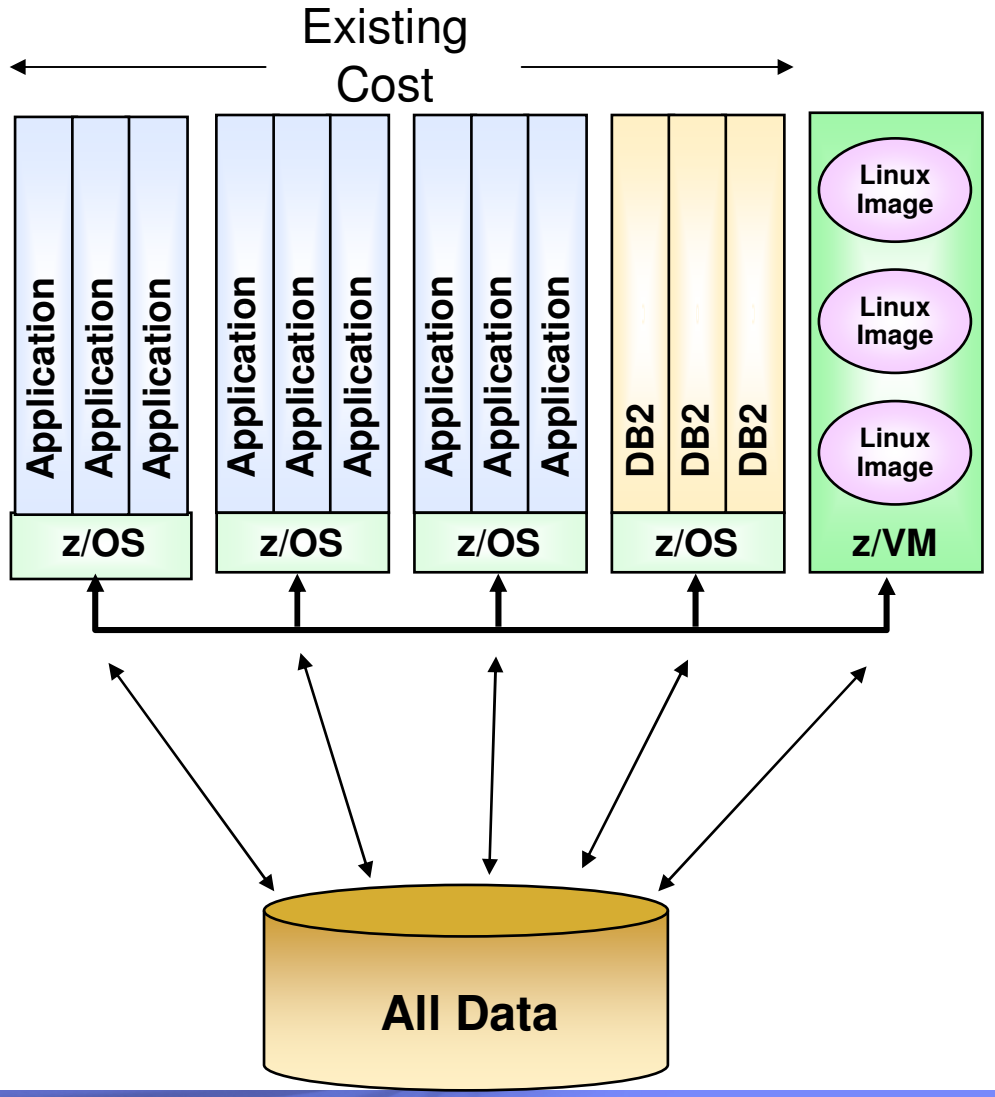

Incremental OTC payment per IFL core

Does not increase the cost of IBM software in other partitions