A Fresh Look at the Mainframe

Mainframe Total Cost of Ownership Issues

Dr. John Shedletsky

VP IBM Software Competitive Technology Lab

Key Points – Distributed Costs

- **The cost of running additional workload on distributed servers goes up linearly**
  - Labor is now the highest cost element in distributed environments
  - Administrative staff costs increase in proportion to the number of servers
  - New workload requires additional servers
  - Cost of additional servers is linear
  - Cost of software licenses is linear
  - Electrical and air conditioning costs also increasing
- **Result – scale out strategies do not reduce the cost per unit of work as the workload grows**

Owing to the nature of individual contracts, some details of this pricing discussion may be at variance with specific instances

Key Points – Mainframe Costs

- **The cost of running incremental workload on the mainframe goes down as the total workload grows**
  - Labor costs hold steady as workload grows
  - IBM pricing policies designed to favor the addition of more workload
  - Special hardware pricing for new workload types
  - Lower software costs per transaction as workload grows
  - Lower electrical and air conditioning consumption than server farms
  - Trade-in value is recoverable for growth customers
- **Customers have learned that mainframes running high workloads are the most cost efficient platform**

Owing to the nature of individual contracts, some details of this pricing discussion may be at variance with specific instances

Mainframe Cost Per Unit of Work Goes Down as Workload Increases

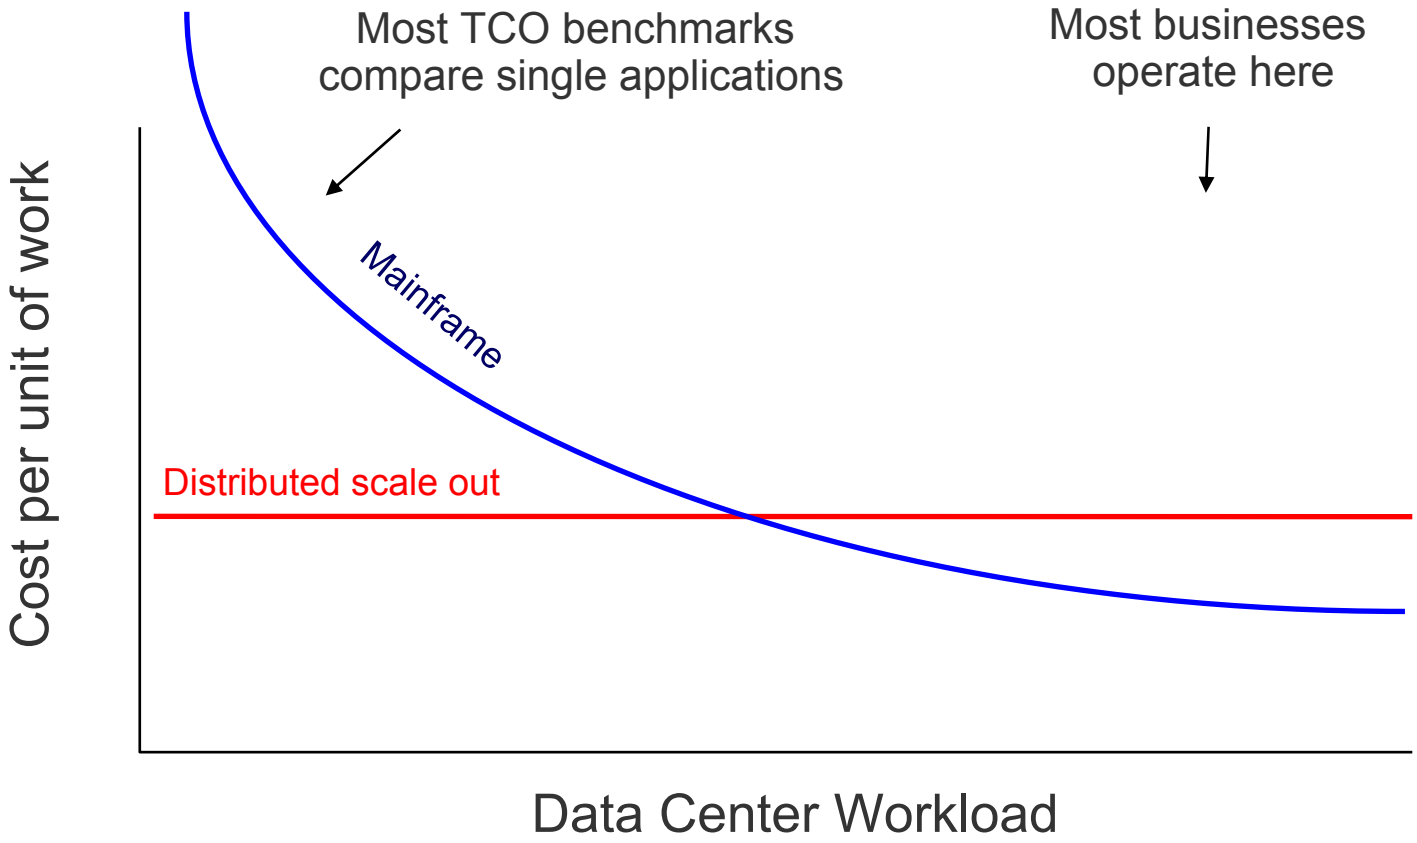

First National Bank of Omaha

First National Bank
Omaha

	Servers	*Reliability*	*Utilization*	*Staff*
First move: **Implemented distributed computing architecture that became** too difficult to monitor, maintain, upgrade and scale	▪ **30+ Sun Solaris servers** ▪ **560+ Intel servers**	**Un-acceptable**	**12%**	**24 people growing at 30% year**
Next move: **Consolidated back on the mainframe**	**z990**	**Much improved**	**84% with additional reserve capacity** *on-demand*	**Reduced to 8 people**

Where is the Cross Over Point?

It depends on your environment …

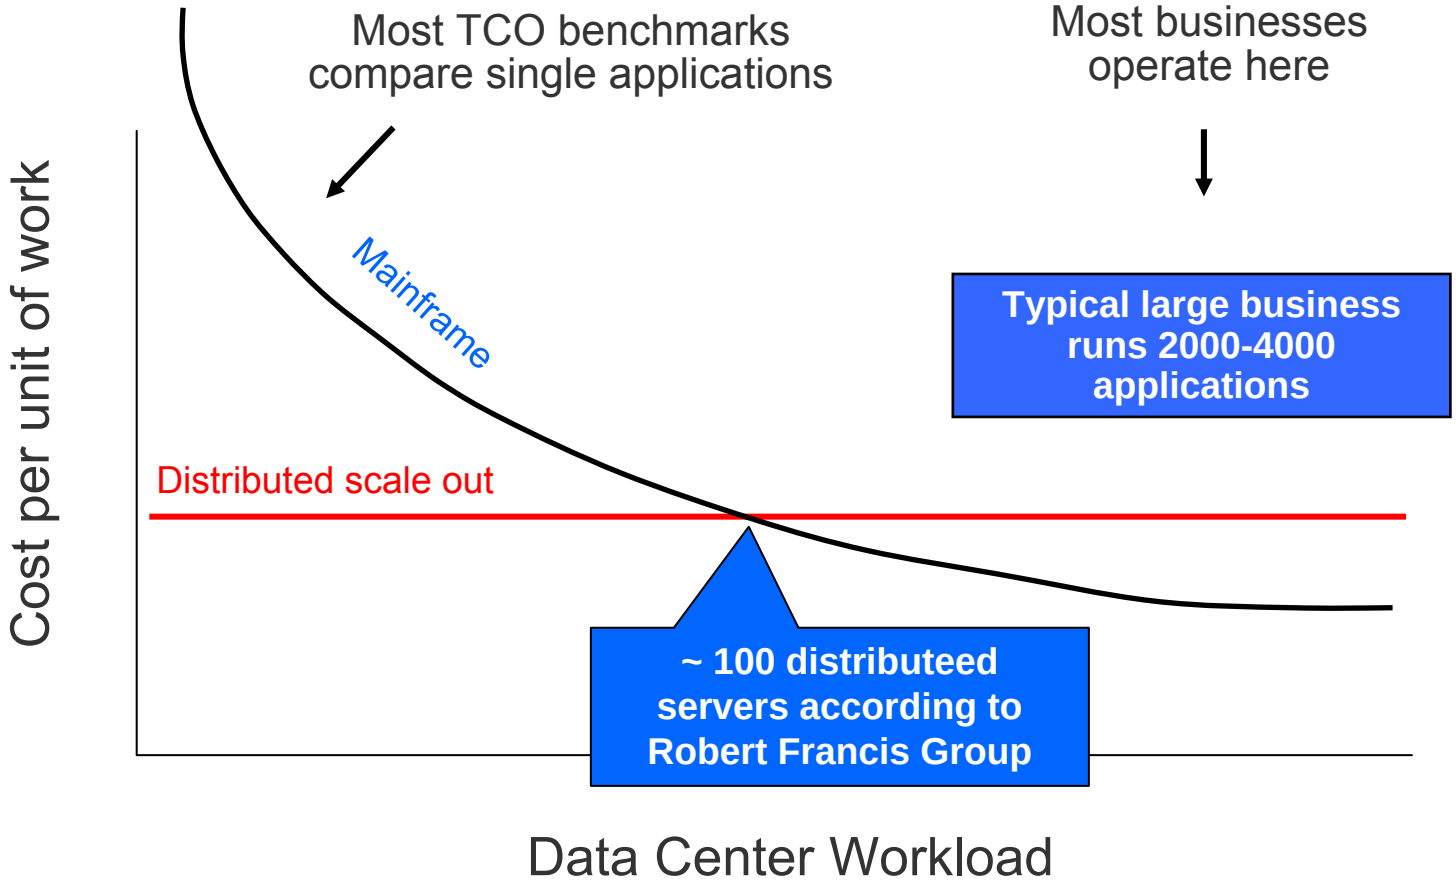

Here Are Some More Hints

- **Growth by Large Customers**
  - 95% of large mainframe customers (average installed MIPs from 13,000 to 15,000 have CAGR of installed MIPS by 21% to 31% since 2002
- **Growth by mid size customers**
  - 72% of mid size mainframe customers (average installed MIPS 1400) have CAGR of installed MIPS from 25% to 34% since 2002
- **Growth by small customers**
  - 70% of small mainframe customers (average installed MIPS 400 to 600) have CAGR of installed MIPS from 38% to 50% since 2002

Let's Break Down the Elements of Cost

Total Cost of Ownership =

Cost of hardware +

Cost of software +

Environmentals +

Cost of labor +

Financial terms

Mainframe Hardware Cost is Decreasing

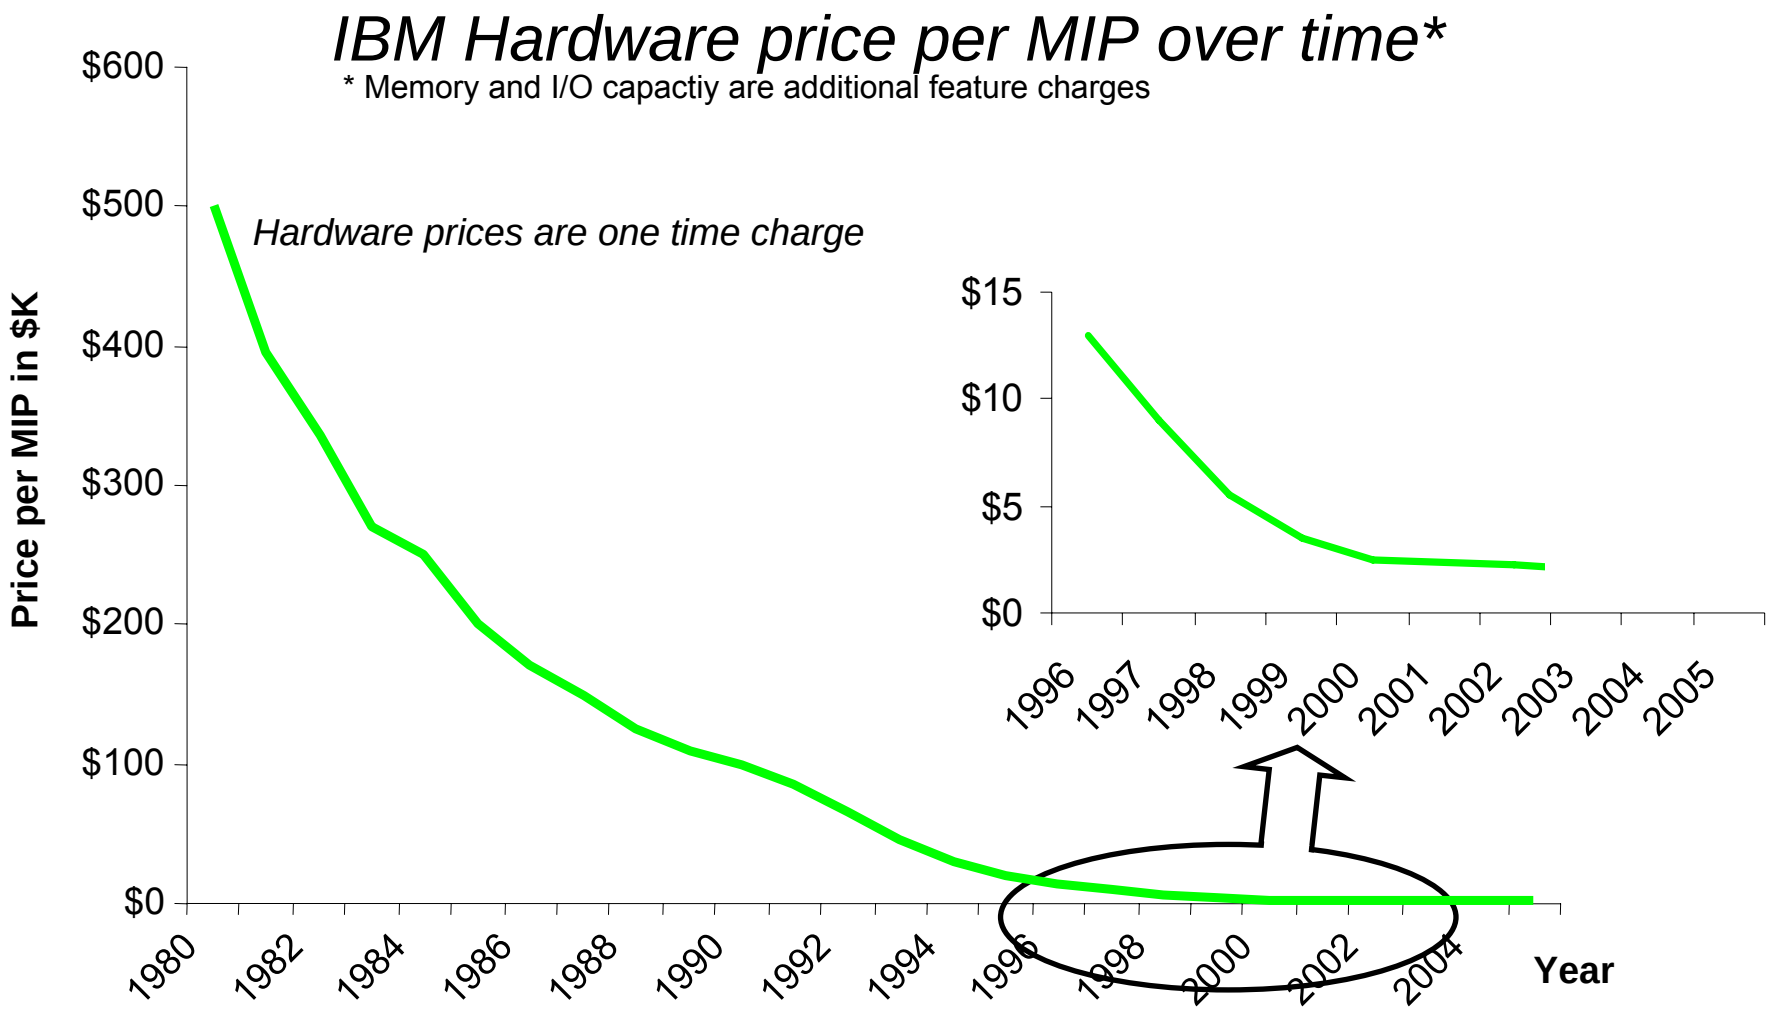

Source: Gartner

“Price Per MIP” Does Not Tell The Whole Story

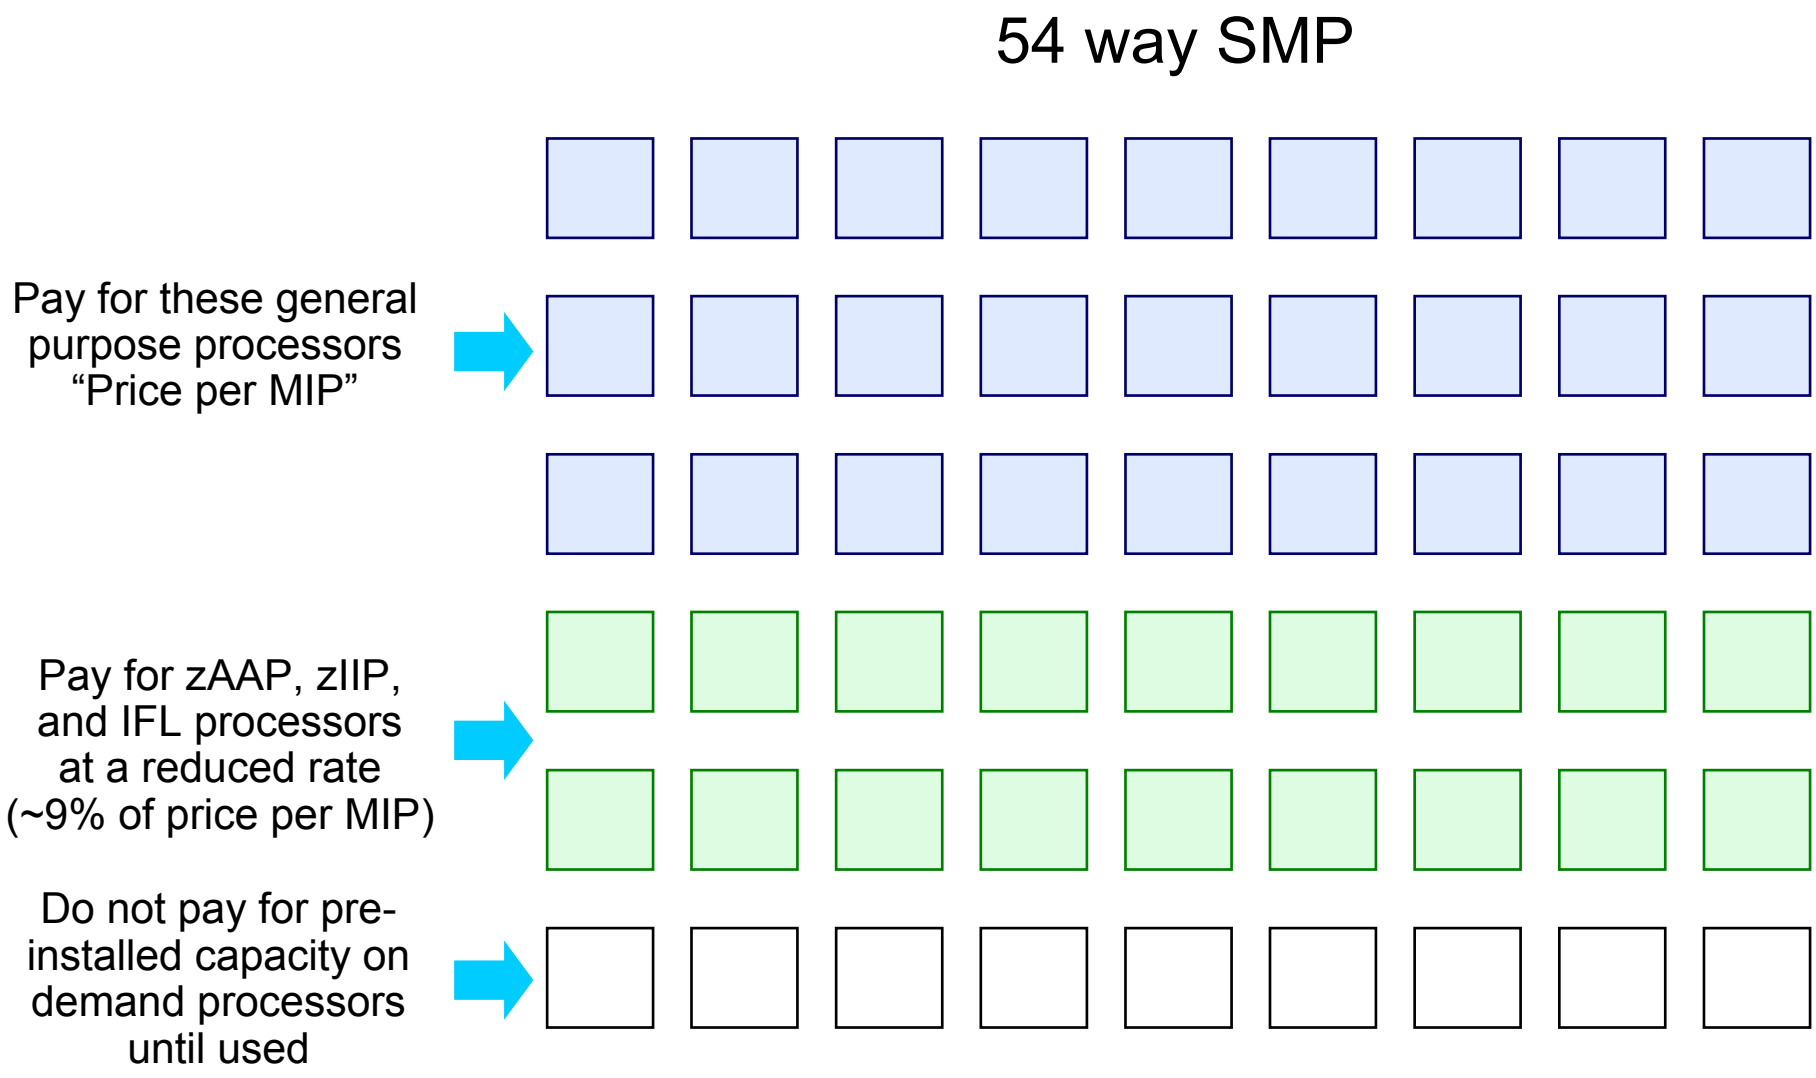

No Additional Charge For System Assist Processors and RAS Processors

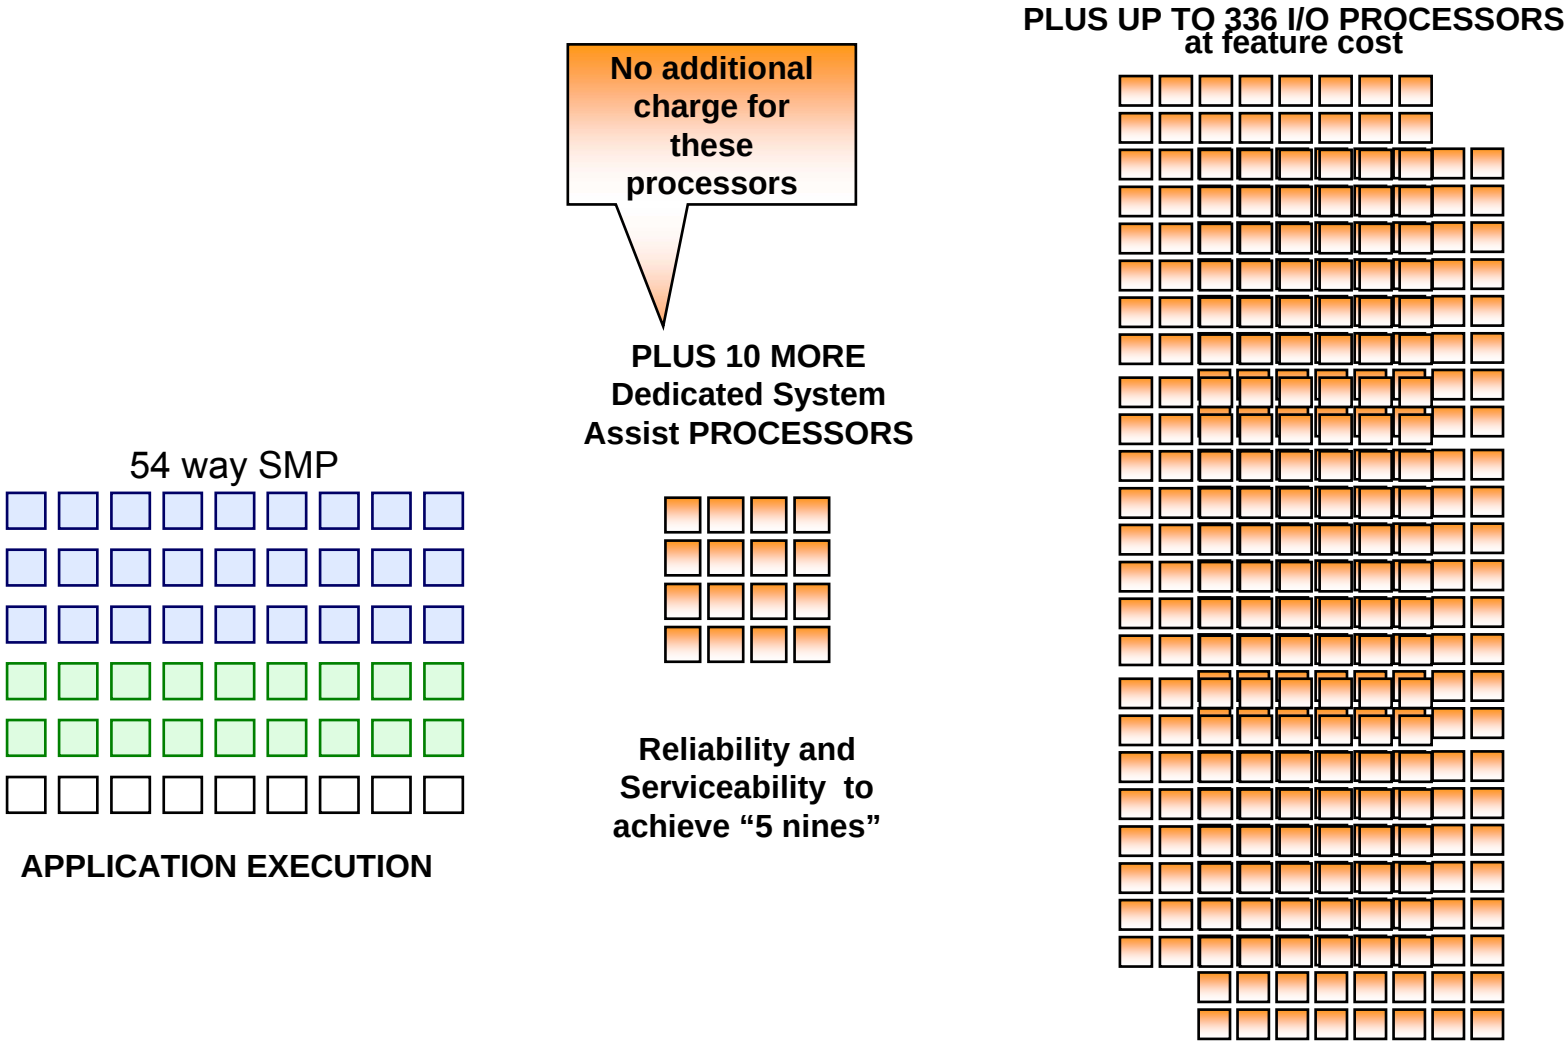

Disaster Recovery – Fast Failover For Less

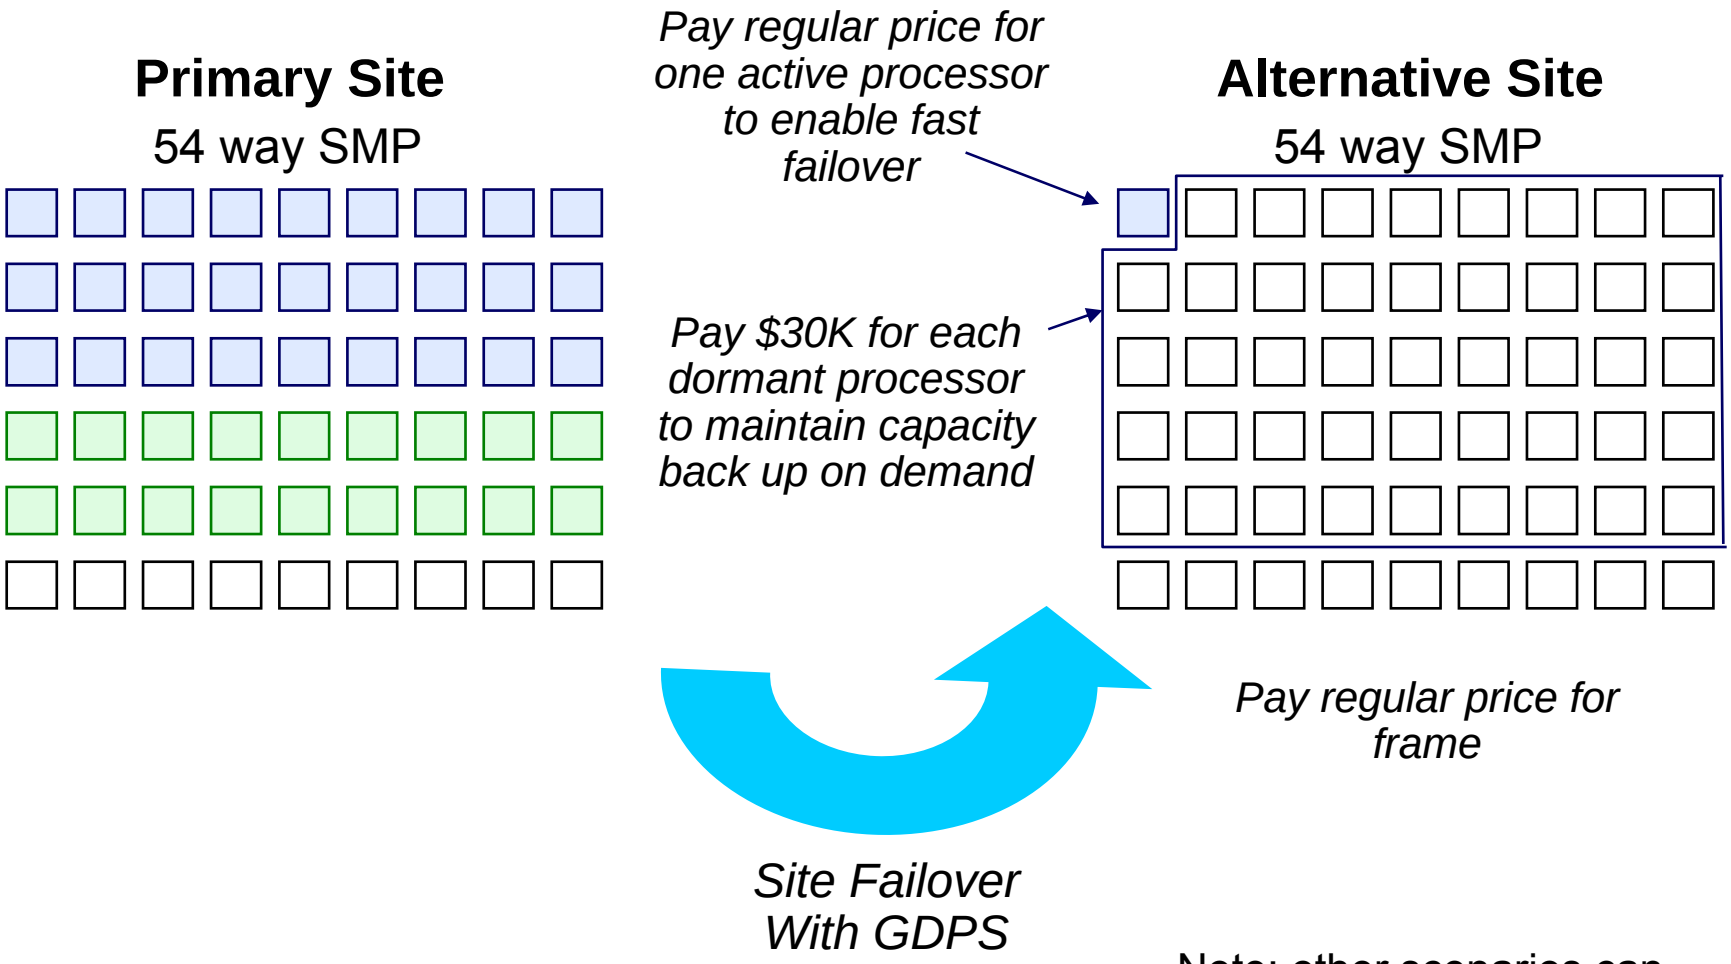

Note: other scenarios can reduce the price further

Datacenter in a Box

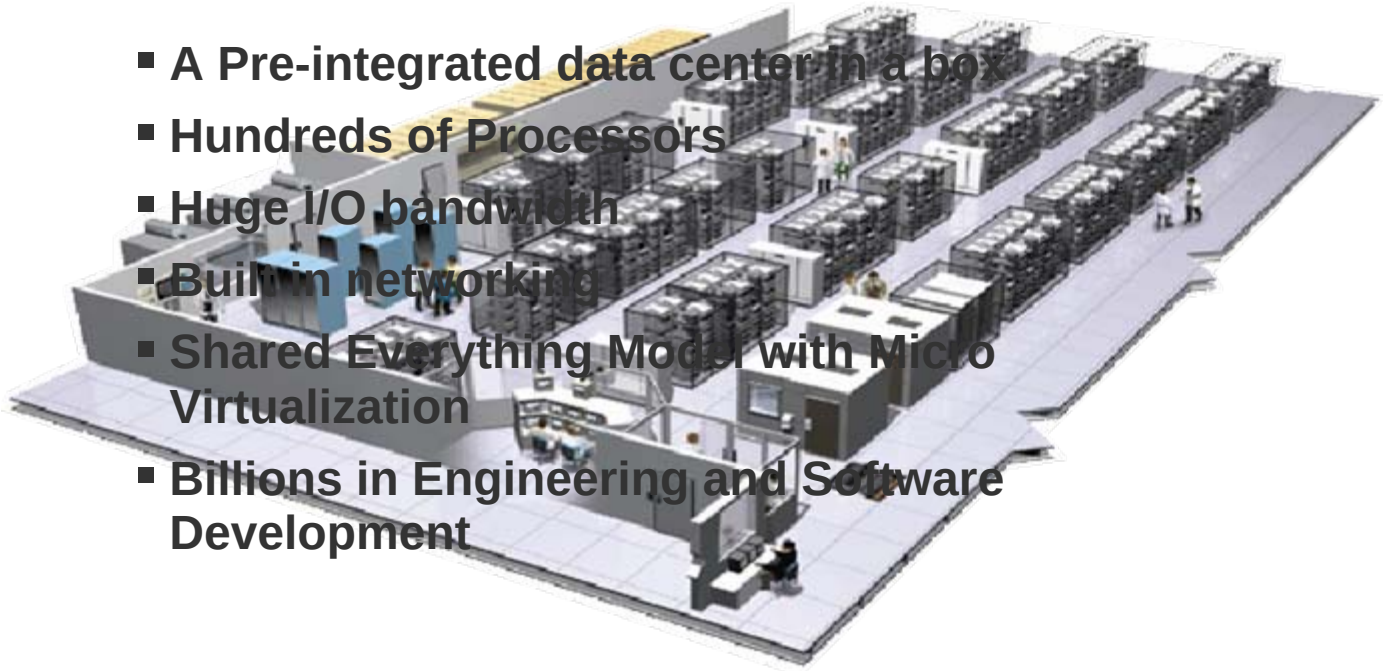

- A Pre-integrated data center in a box
- Hundreds of Processors
- Huge I/O bandwidth
- Built in networking
- Shared Everything Model with Micro Virtualization
- Billions in Engineering and Software Development

- Building your own datacenter is costly and complex
- Install and configure hundreds of devices
- Networking
- Data Silos and Synchronization
- Power consumption
- Linear Staffing Costs
- Frequent Outages

Resulting in tremendous efficiencies

No extra charge for this deep pre-integration!

Micro Virtualization – How it Looks in z Architecture

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

Workload Manager allocates resources as needed by service classes

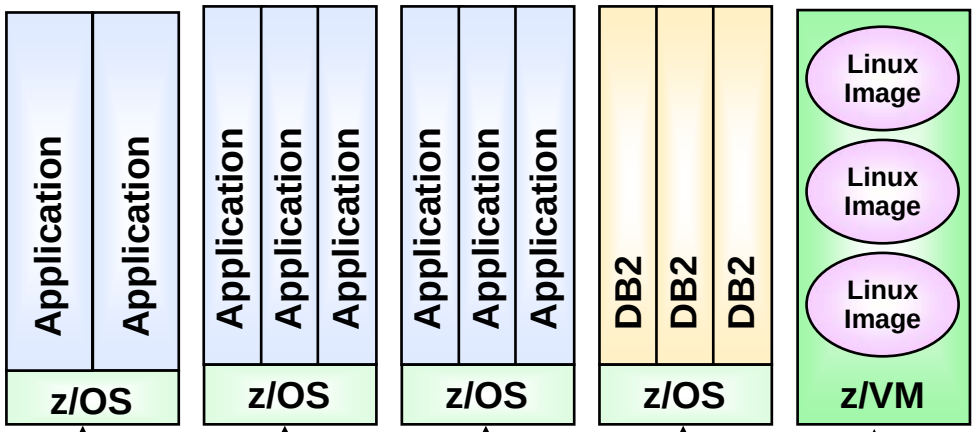

Internal networking via secure high speed Hipersockets

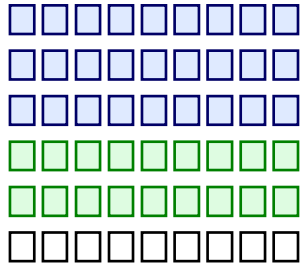

Intelligent Resource Director dynamically allocates processors to partitions

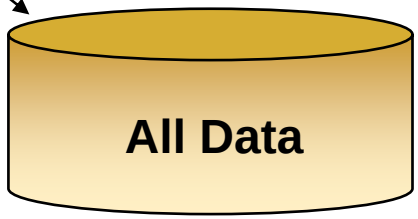

Shared access to all disk data and to external networks

Eligible workload automatically dispatched to zIIP and zAAP specialty processors

TCO Comparisons

Distributed

CONSOLIDATION

OFFLOAD

System Z

Asian Financial Services Customer Offload Project - Overall

3x HP 64-way Production Application and DB

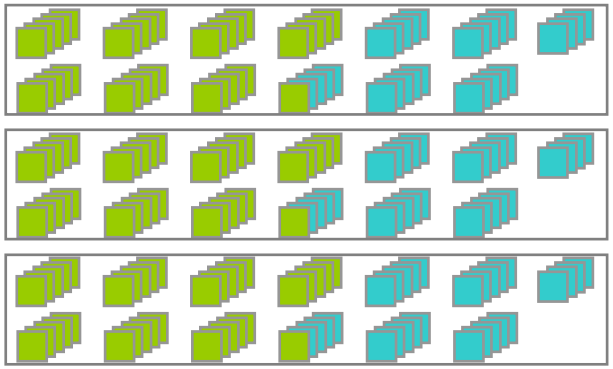

1x HP 64-way Dev&Test / Batch

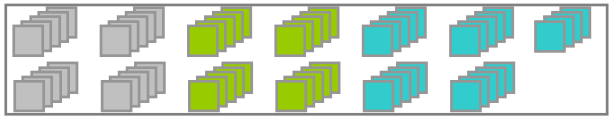

2x HP 32-way PL/1 (Mgmt, Dev&Test, and Batch)

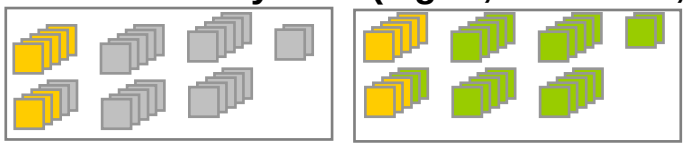

17 processors (6700 MIPS)

320 Unix processors (816,002 RPE's)

2x z990 5-way (production)

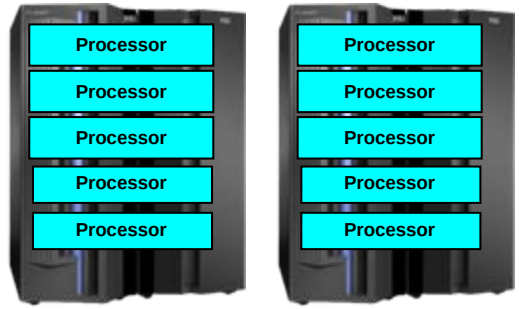

Z990 7-way (production + test)

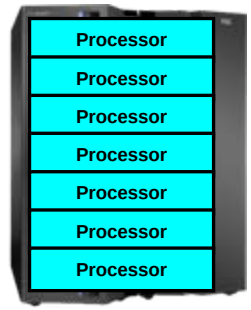

Plus:
2x HP 16-way servers : external, HP rx8620
3x IBM P570 servers : Web Appl server

122 RPE's per MIP

No disaster recovery

Asia Pacific Financial Services Customer Offload Project – Database

Production Oracle RAC cluster of 3 HP Superdome nodes (28 processors per node)

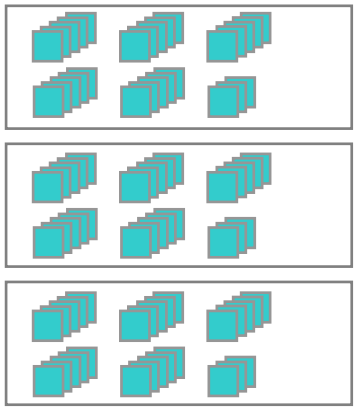

Batch (16 processors)

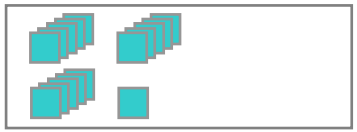

7 processors (2680 MIPS*)

z990 Processors for DB2 (production and development)

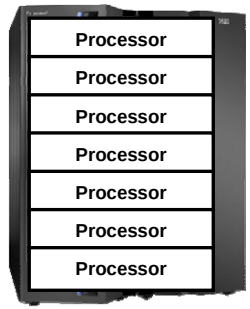

100 Unix processors Oracle RAC (233,510 RPE's)

87 RPE's per MIP

No disaster recovery

* DB2 is estimated to be 40% of total workload

European Banking Customer Study

TCO Analysis to Offload CICS Transaction Workload

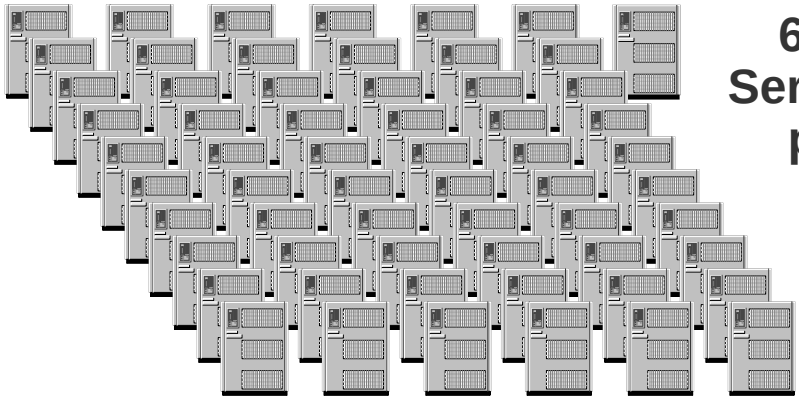

61 Windows Servers with 244 processors

To Windows

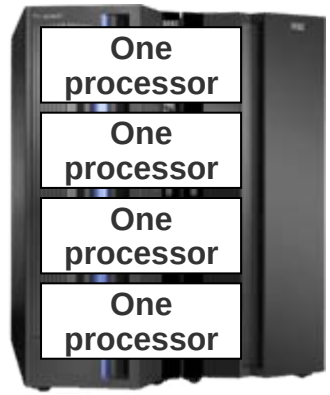

49 System P LPARS with 98 processors

To Unix

Conclusion: Same TCO with no benefit from additional migration cost

Why Do Servers Proliferate in Offload Scenarios?

- **The following considerations contribute to server proliferation**
  - De-multiplexing of applications to dedicated servers
    - One application workload per server group
    - Peak-to-average provisioning yields low utilization
    - Additional provision for expected growth in out years (no capacity on demand)
    - Batch workload may stress I/O capabilities
    - Separate servers for production, failover, development/test, disaster recovery
    - Infrastructure servers for systems management
  - Processing comparisons
    - Language expansion (CICS/COBOL path lengths are highly optimized)
    - Conversion factor (MIPS to TPM-C or RPE) worsens as I/O rates increase
    - Oracle RAC inefficiencies compared to DB2

- **Other TCO considerations**
  - 3 to 5 year lifetime for distributed servers requires repurchase
  - Dual environments during migration
  - Partial offloads eliminate the lowest cost MIPS first

Utilization of Distributed Servers

Server dedicated to one application

- Typical utilization of Windows Servers 5 – 10%
- Typical utilization of UNIX Servers 10 – 20%
- Typical utilization of System z Servers 85 – 100%

System z Virtualization, Workload Management, and Storage Bandwidth Achieve High Levels of Utilization

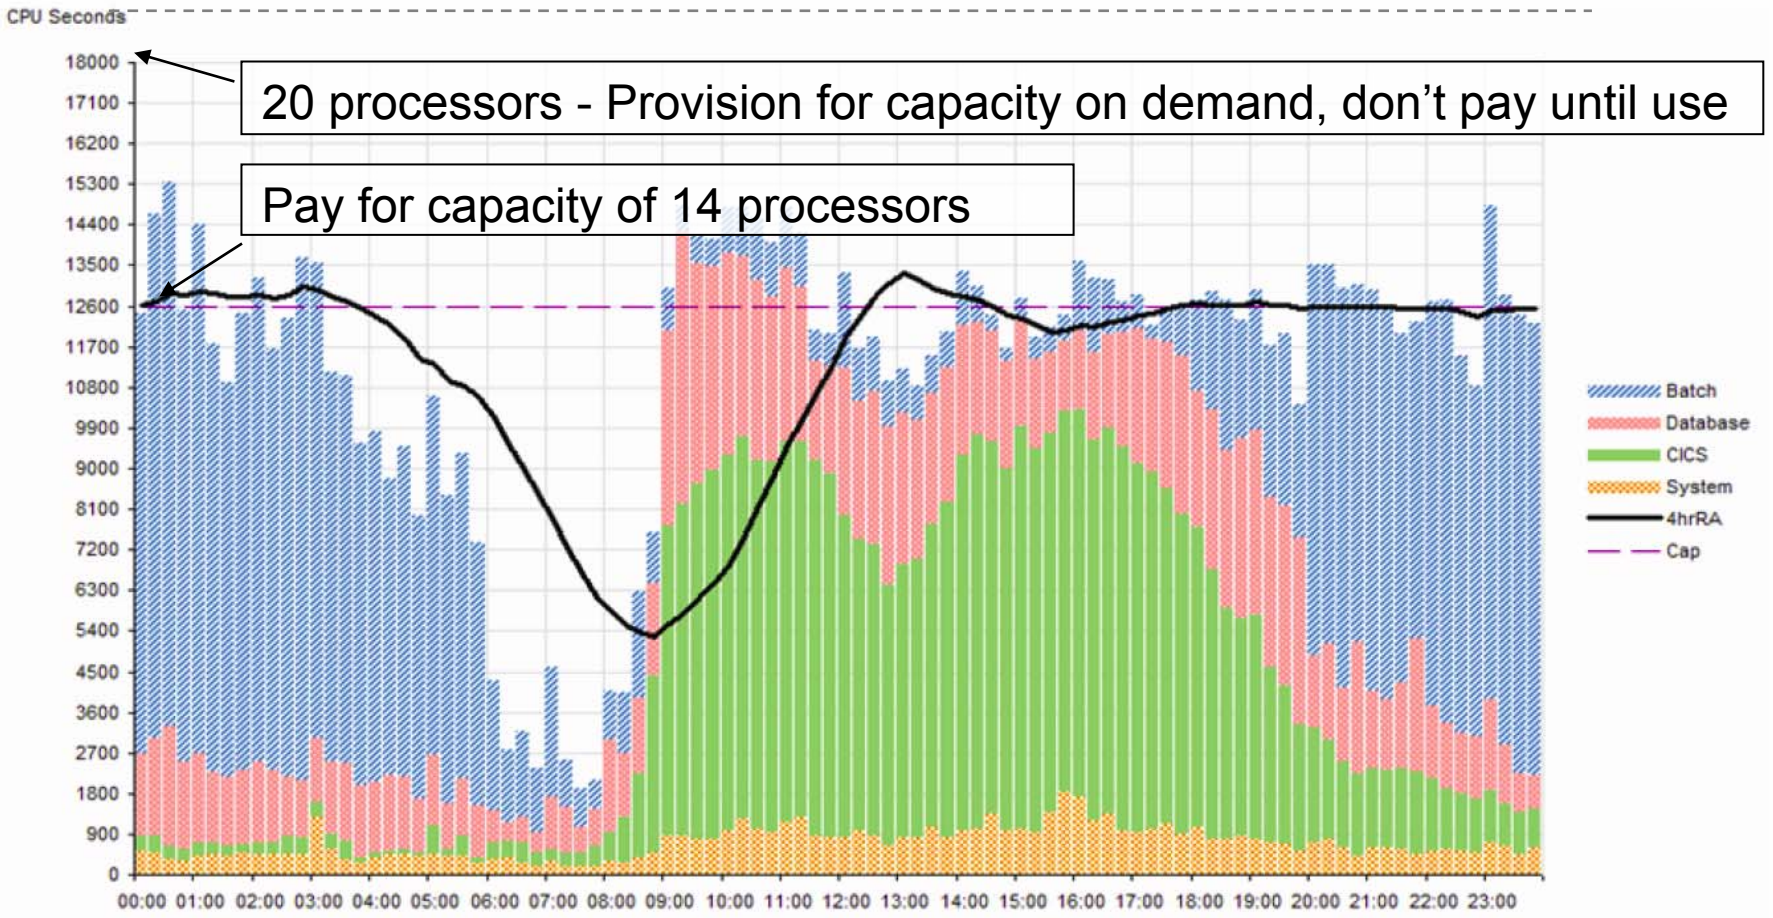

Note:

- *Each bar represents the amount of CPU seconds used in 15 minutes (= 900 seconds) with 2 10-way machines*
- *The way Workload Management controls the workload 4-hour rolling average to the Cap "high-water mark"*

System z Virtualization, Workload Management, and Storage Bandwidth Achieve High Levels of Utilization

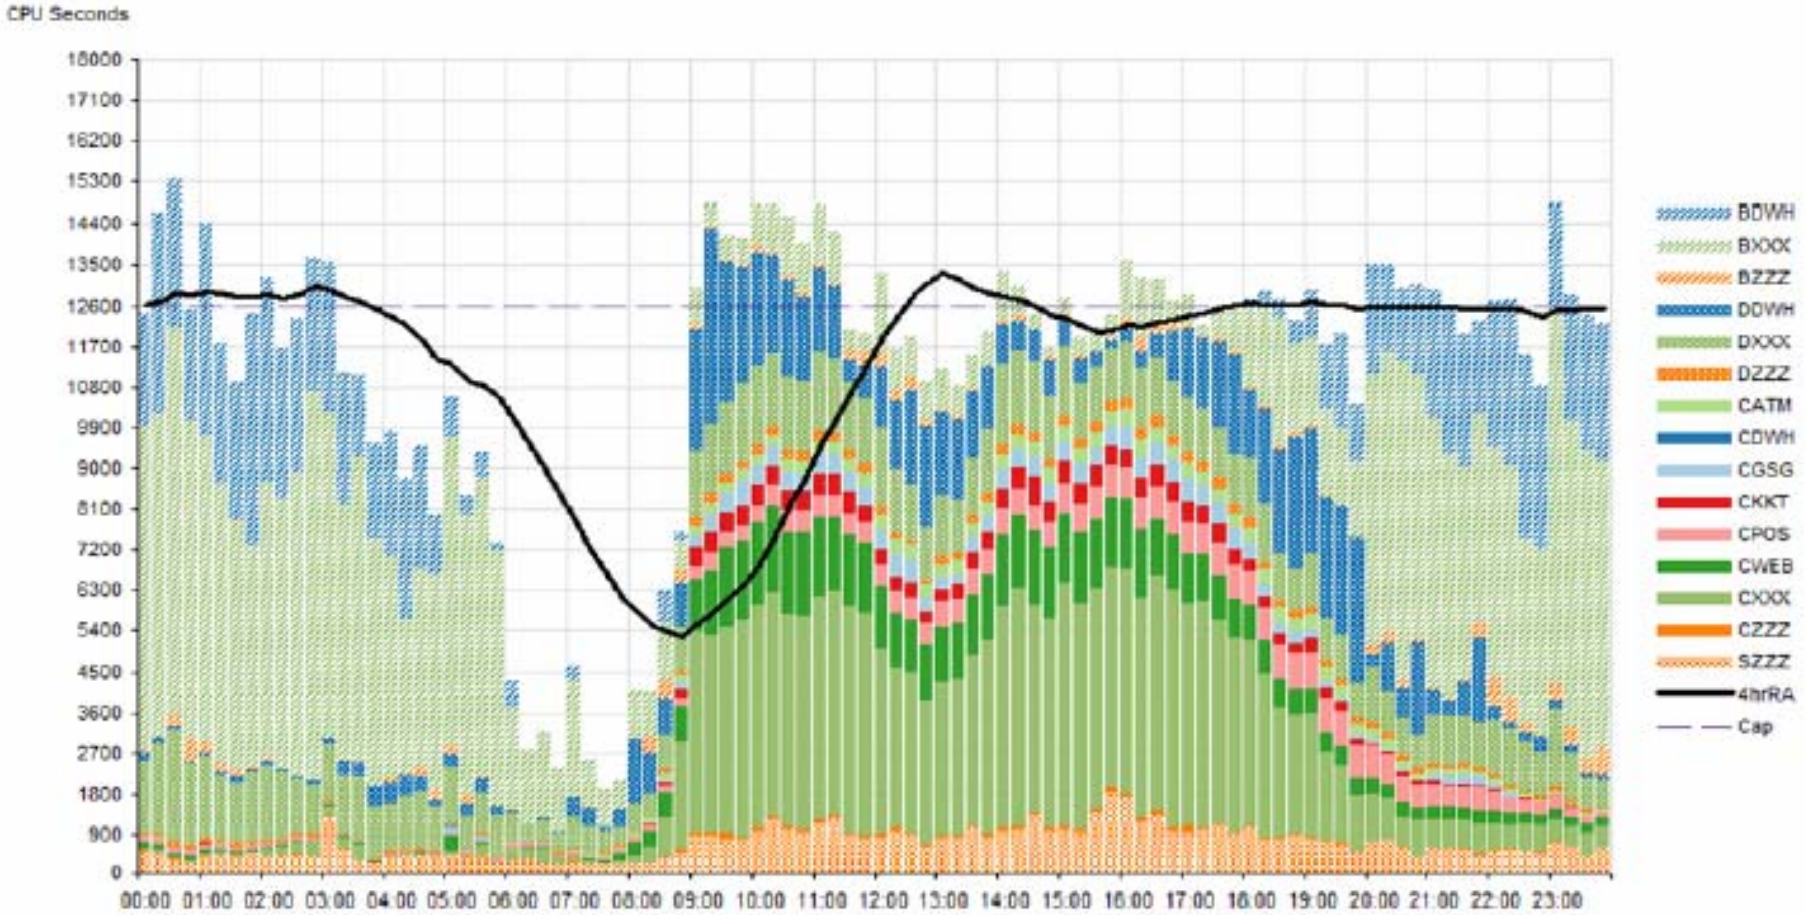

Distributed Systems Storage Utilization

- **The Total Cost of Storage is Typically Three Times More in Distributed Environments**
  - Application specific data silos tend to over-allocate
  - Storage utilization of 25-30% or less is typical in distributed environments
  - Mainframe fine grained allocation and data sharing yield typical storage utilizations of 80% +
  - Data copies are often used to separate “batch” style workloads from online
  - Cheap disks cannot be used by high RAS workloads in distributed environments

- **Management Headaches**
  - Disaster recovery of separated data silos
  - Synchronization, and transfer requirements

“Physically moving data” ranked highest in IT’s “pain index” In an IBM Storage study of over 200 companies,

IBM Software Price Per Transaction is Also Going Down

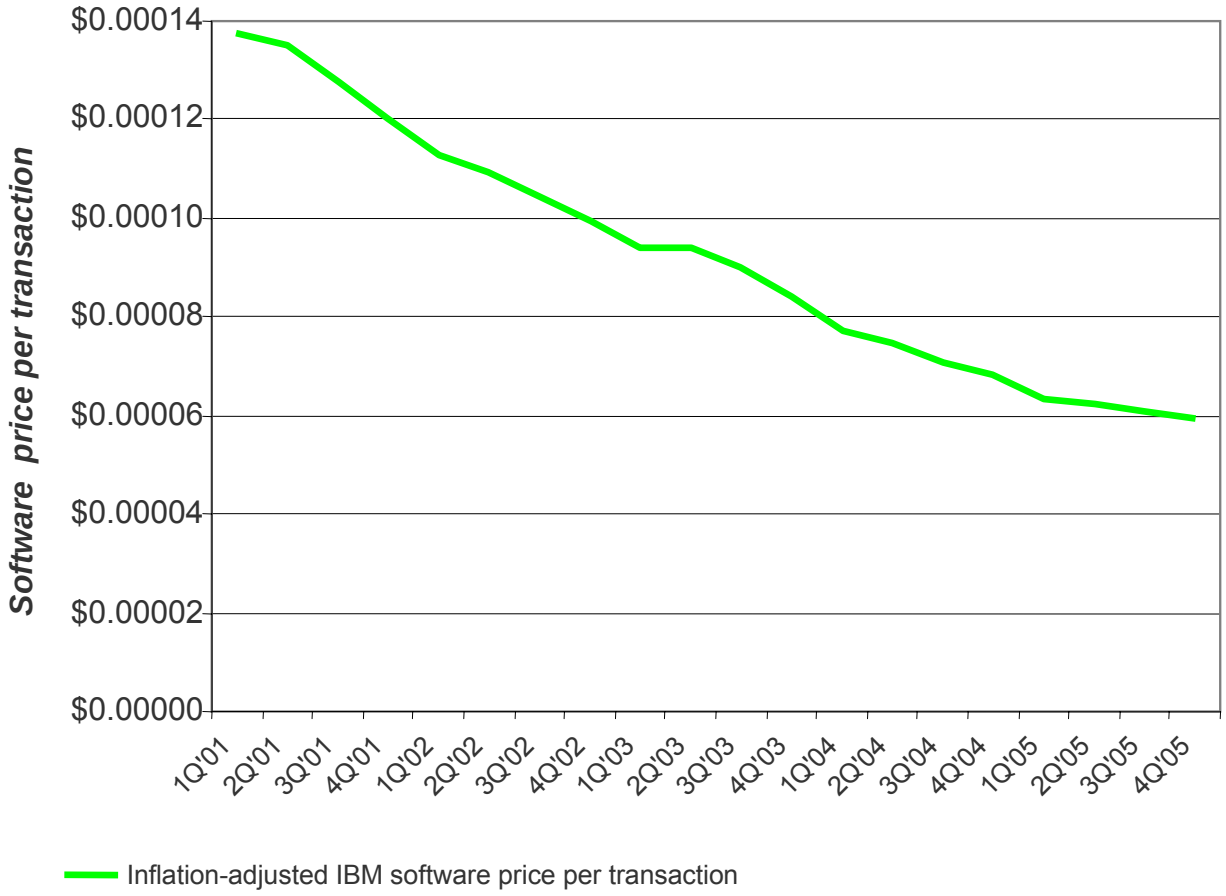

57% decrease in past 60 months

17% decrease per year

What makes the price go down?

Pricing curves favor growth

Specialty processors (zAAPs, zIIPs, IFLs)

Technical pricing allowances

Source: IBM SWG Finance
Data is WW customer revenue only (not IGS)
Data includes specialty engines
'Highway conditions .. mileage may vary'

Software Pricing Curves Favor Growth

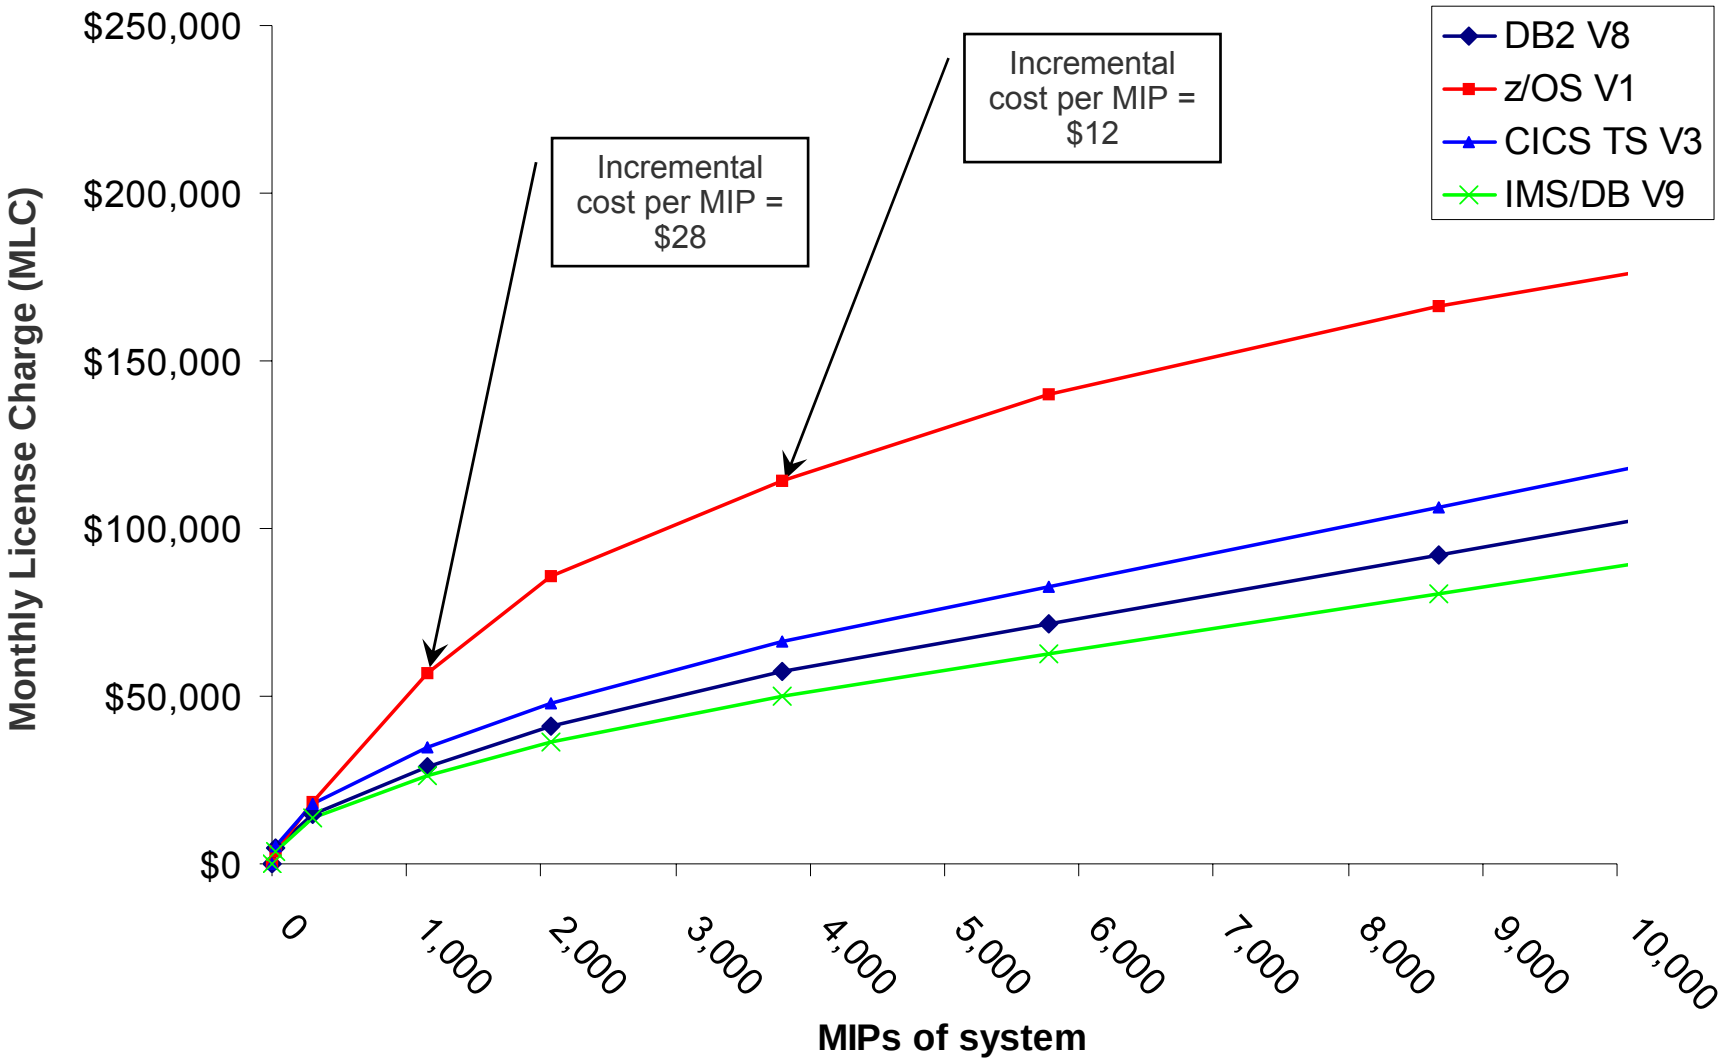

Overall Software Price Per MIP Decreases as System Size Increases

- Lower incremental costs for large users
- Server consolidation to fewer, bigger boxes
- Sysplex Aggregation allows customers to take advantage of these slopes

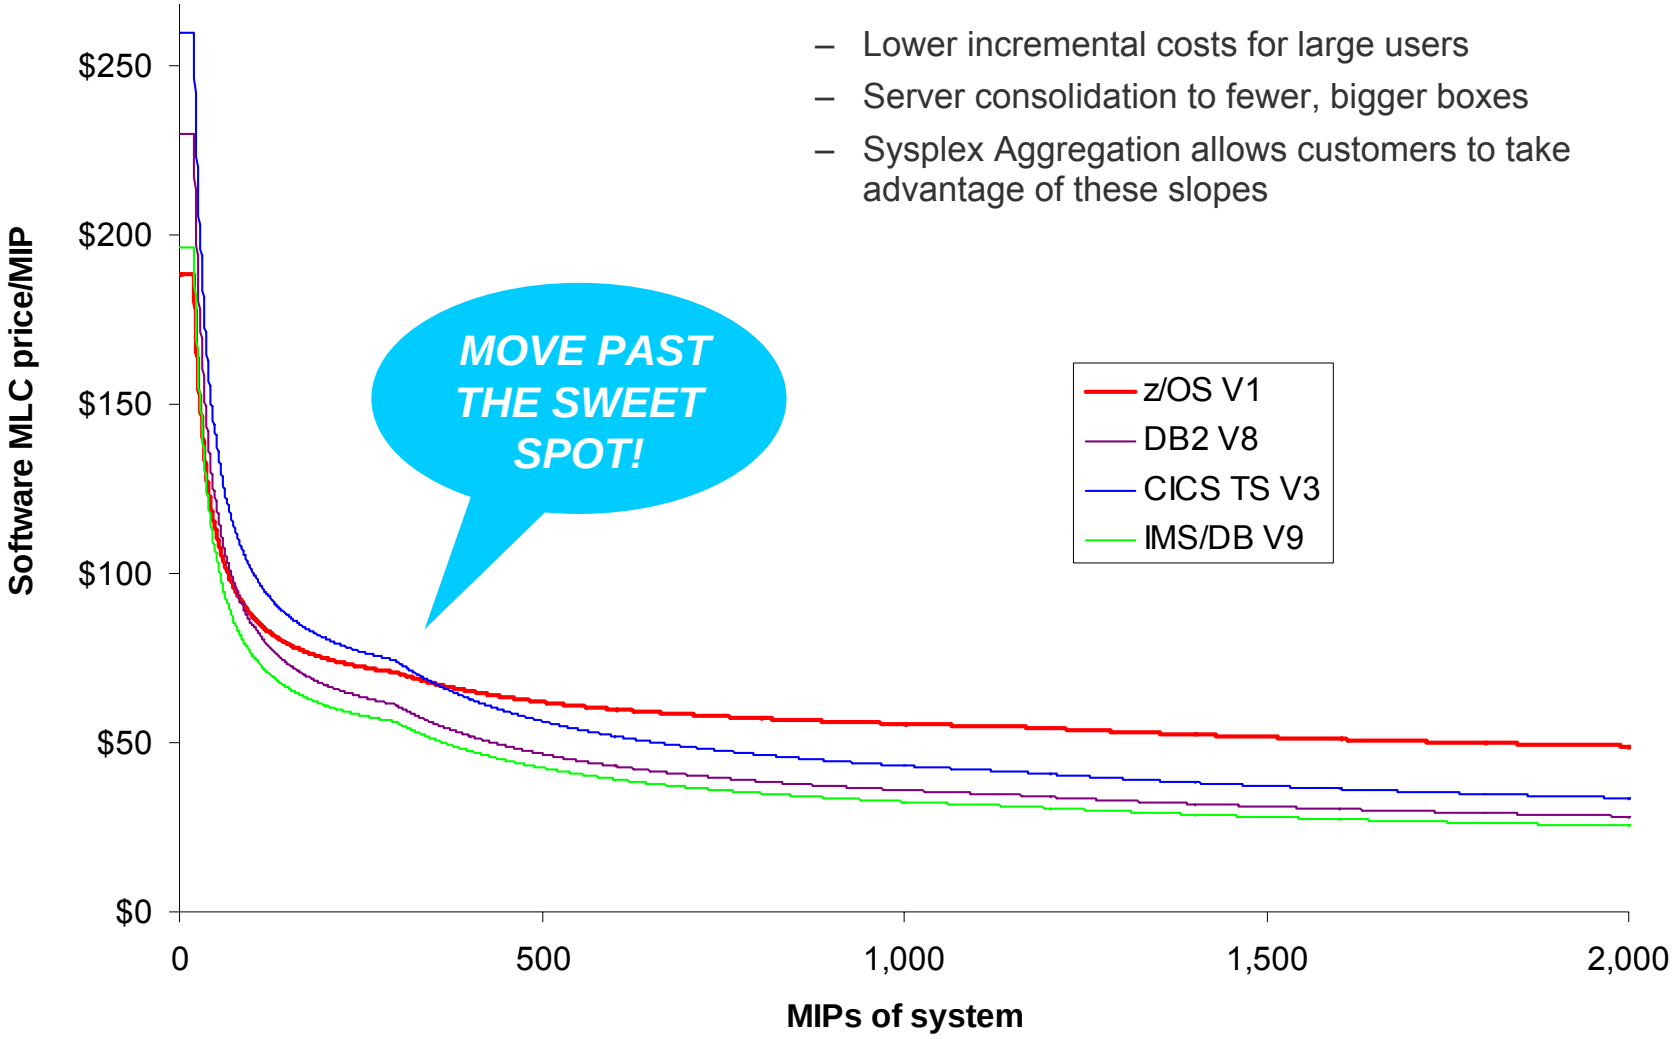

Let's Consider How this Works in Reverse … Lowest Cost MIPS are Offloaded First

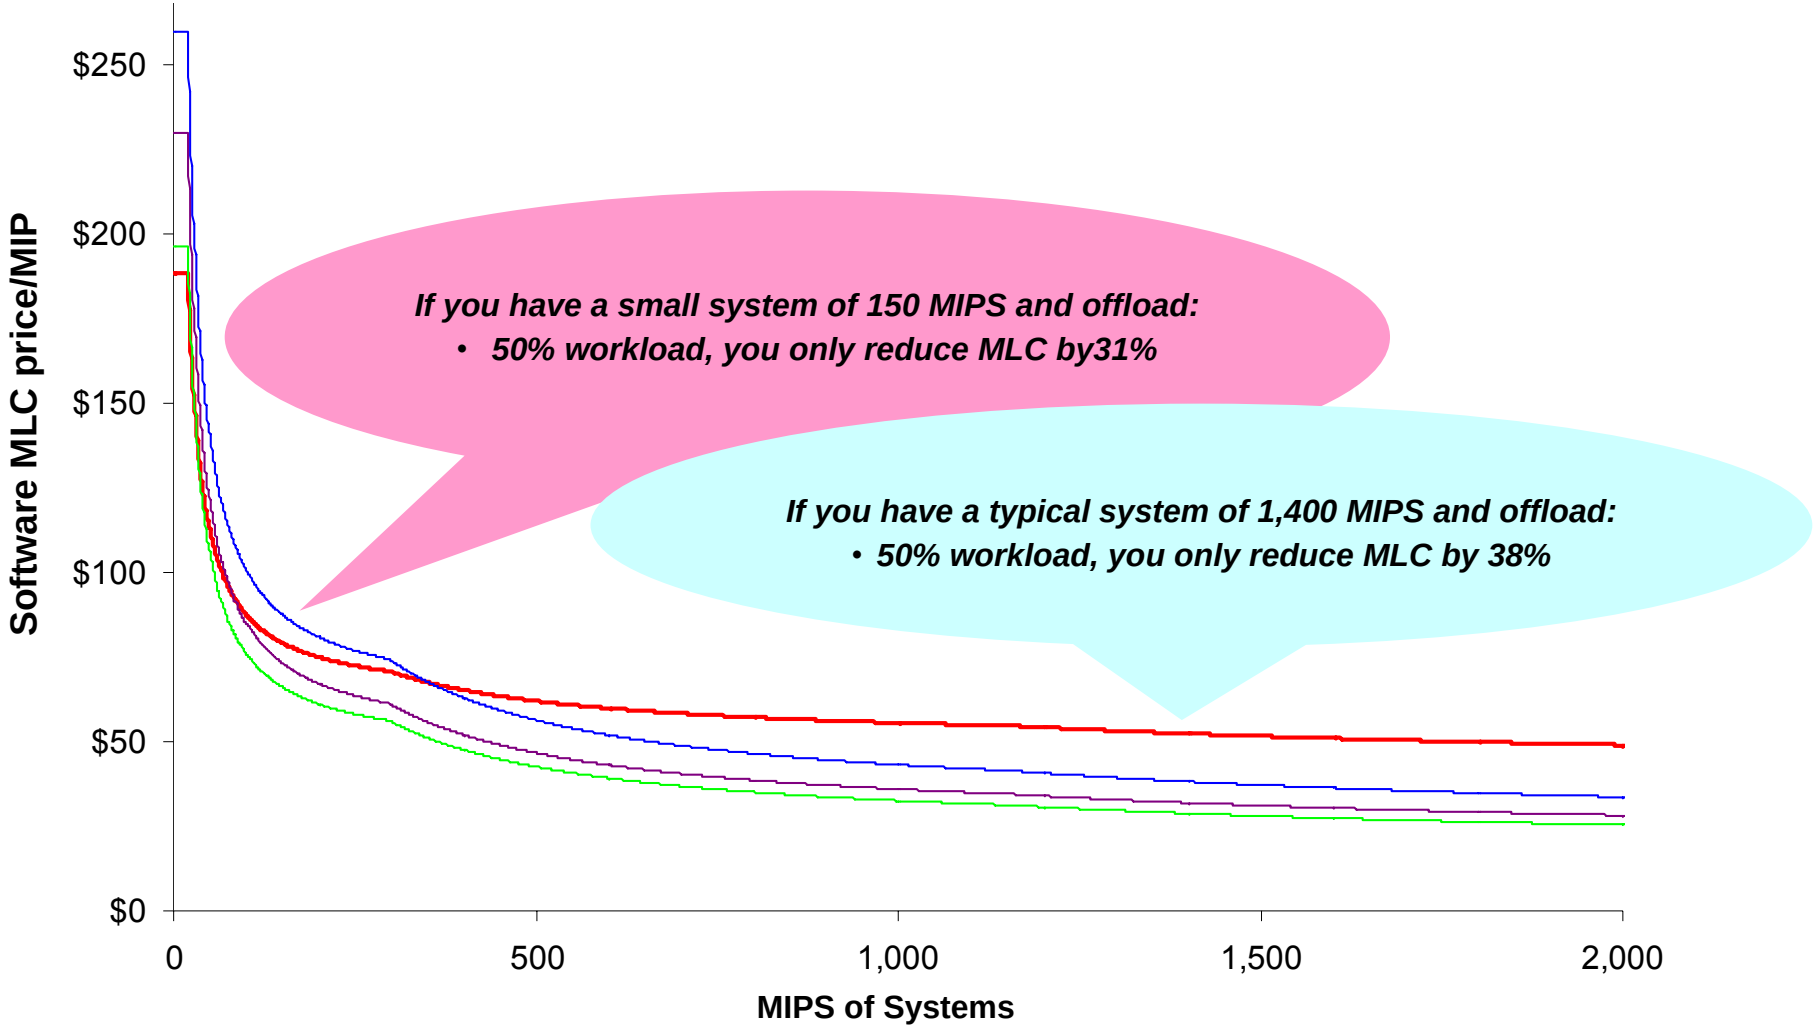

IBM Actually Charges on the Basis of MSUs (Millions of Service Units Per Hour)

- **1 MSU currently is equal to about 7.3 MIPS (for a z9 EC)**

  So…

- **Software for a 580 MIPS machine will be charged at a rate of 81 MSU's**

Various ratings online at:

MSUs http://ibm.com/zseries/library/swpriceinfo/hardware.html
LSPR MIPS http://ibm.com/zseries/lspr
Hardware SRM Constants http://ibm.com/zseries/srm

“Technology Dividend” = ~10% MSU Reduction Each Generation

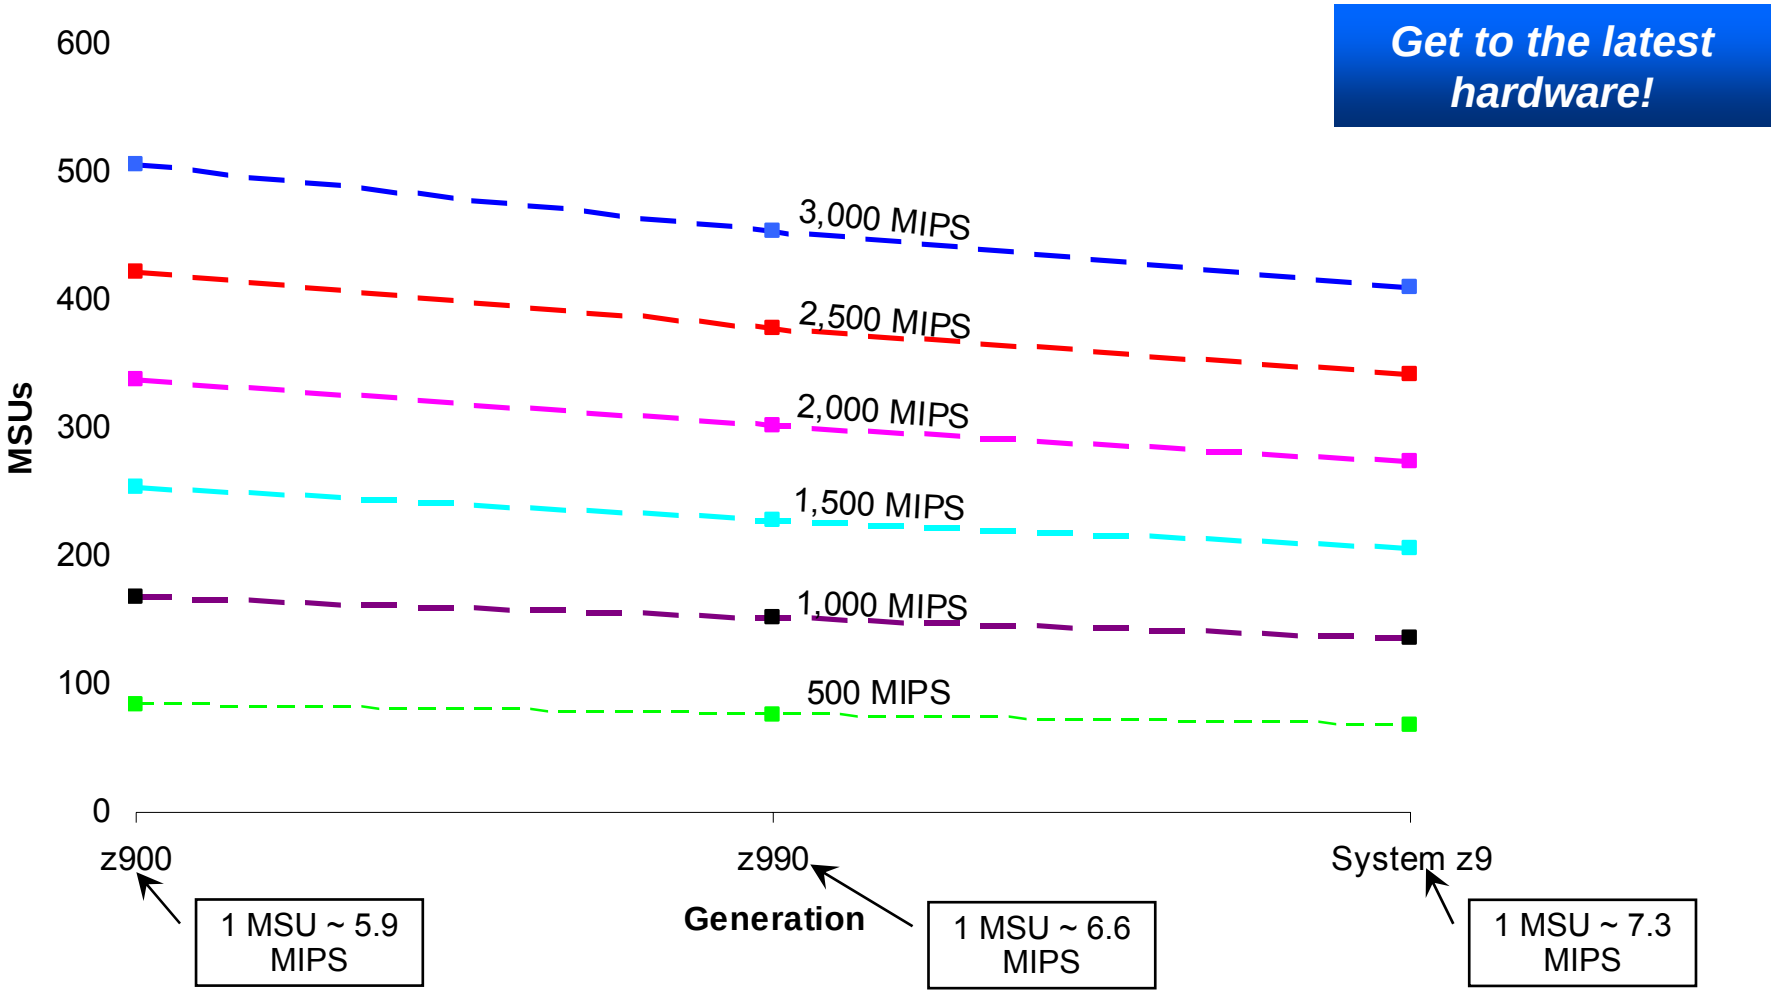

Example of Sub-Capacity Pricing , Saving ~ $16K MLC

FULL-CAPACITY PRICING

z900 Model 1C7

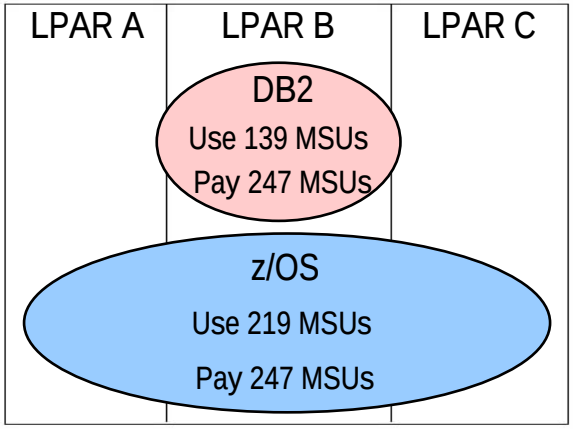

247 MSUs

Total cost = $106,915/mo

SUB-CAPACITY PRICING

z900 Model 1C7

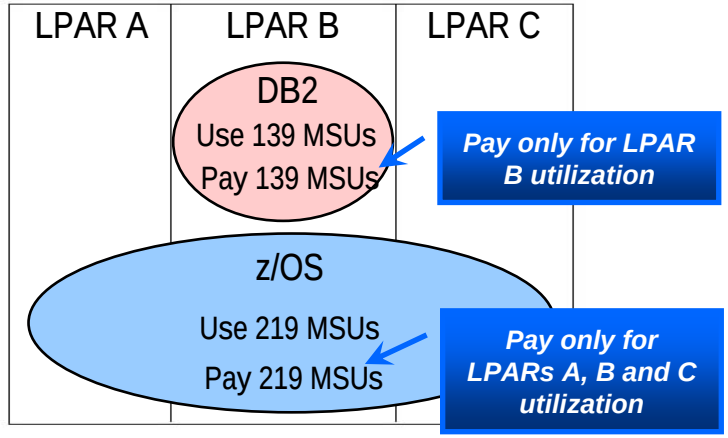

247 MSUs

Total cost = $91,011/mo

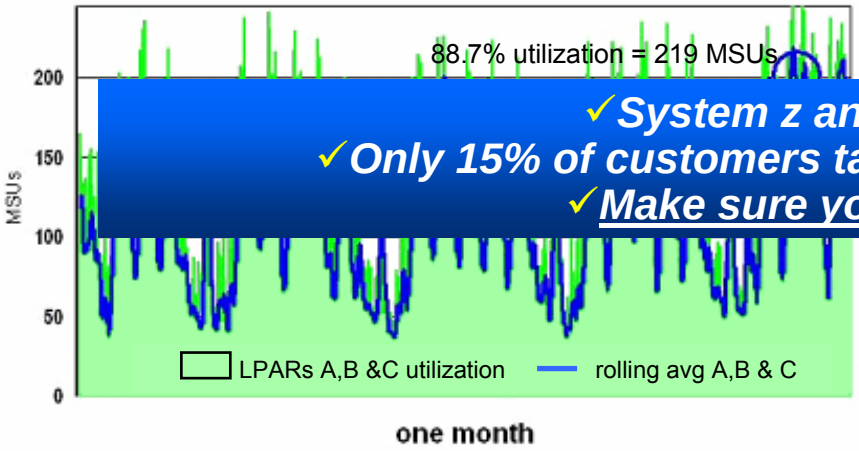

one month

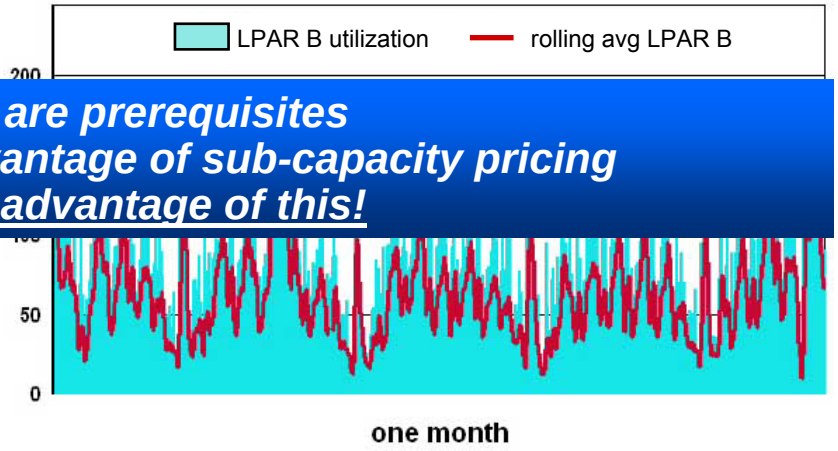

one month

✓*System z and z/OS are prerequisites*
✓*Only 15% of customers take advantage of sub-capacity pricing*
✓*Make sure you take advantage of this!*

Technology Dividend Helps Offset Software Upgrade Increases

G5 S/390 processor
9672-R56
Purchased in 3Q98
Back level 4 generations
540 MIPS (5 CPU's)
93 MSU's

Upgrade to
current generation
z9 Enterprise Class
2094-701

	MLC			Tech Div and sub capacity 540 MIPS (<1CPU) 76 MSU's MLC	If No Tech Div 540 MIPS (<1CPU) 93 MSU's MLC
Database (-3 Generations)	$15,378	*DB2 UDB v6*	→ *DB2 UDB v9*	$19,866	$21,383
Transaction Processing (-3 Generations)	$14,733	*CICS ESA v4*	→ *CICS TS v3.1*	$22,061	$21,914
Operating System (-2 Generations)	$46,485	*OS/390 Base*	→ *z/OS VI Base*	$27,633	$49,575
	$76,596			$69,560	$92,872

Example of Sysplex Aggregation, Saving >$82K MLC

NO AGGREGATION

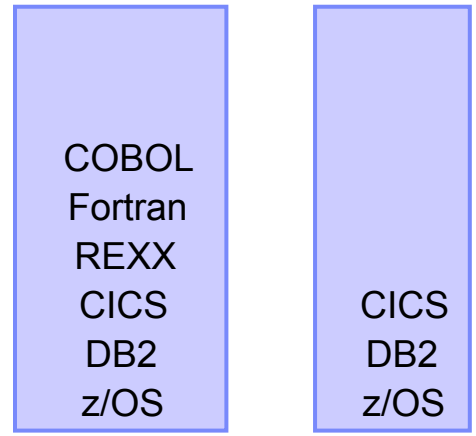

	302 MSUs		276 MSUs
COBOL	302 MSUs		
Fortran	302 MSUs		
REXX	302 MSUs		
CICS TS	302 MSUs	+	276 MSUs
DB2	302 MSUs	+	276 MSUs
z/OS	302 MSUs	+	276 MSUs

Total cost = $321,596/mo

AGGREGATION

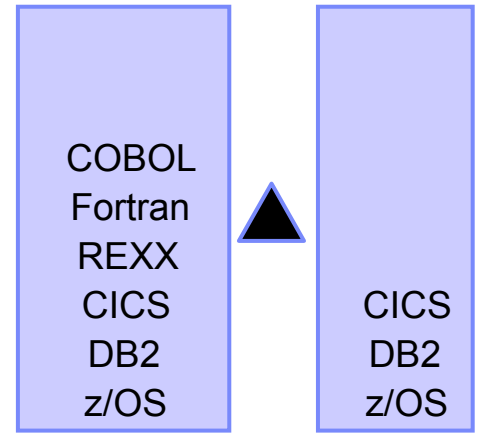

	302 MSUs	276 MSUs
COBOL	302 MSUs	
Fortran	302 MSUs	
REXX	302 MSUs	
CICS TS		578 MSUs
DB2		578 MSUs
z/OS		578 MSUs

Benefit of price slopes

Total cost = $239,090/mo

New! System z New Application Licensing Charge (zNALC)

- **New pricing model to encourage running new applications on z/OS**
- **z/OS MLC is discounted 80-90% depending on machine/LPAR size**
  - Examples:
    - System z Enterprise Class 710: $35,899 for z/OS using zNALC compared to $173,574 base price
    - System z Business Class S03: $6,294 for z/OS using zNALC compared to $44,707 base price
- **Middleware can use normal sub-capacity pricing**
- **Application must run in a separate LPAR(s) from current workload**
- **Application must be certified by IBM as a qualifying application**
  - Examples of qualifying applications WebSphere Application Server, Domino, SAP, Siebel, and PeopleSoft

The Economics of Workload Consolidation

- **Distributed servers typically run at utilization levels in the range of 5% to 20%**
  - Production servers, development servers, test servers
- **Virtualization and workload management enable consolidation on the mainframe**
  - Run multiple images on fewer processors
  - Achieve utilization levels of 85% or more
- **Mainframe "specialty engines" further improve consolidation economics**
  - WebSphere, Database, Linux

5% to 20% utilization

CONSOLIDATION

Full utilization

zAAP and zIIP “Specialty Engines”

- **Special assist processors**
  - For Java workloads (zAAP)
  - For selected DB2 workloads (zIIP)
- **Attractive pricing**
  - Hardware is $125K per processor one time charge
    - $125K for a 580 MIP processor
    - ~ 9% of the normal price
  - No charge for IBM software running on zAAP/zIIP
  - Free upgrade to next generation!
- **Requirements**
  - z9-109 hardware platform
  - Latest service levels
  - zAAP/zIIP hardware feature
  - Max number of zAAPs =< number of general purpose processors
  - Max number of zIIPs =< number of general purpose processors

How Much Workload is zAAP or zIIP- able?

- **How much DB2 workload can typically be run on a zIIP?**
  - DRDA Remote Access Protocol (**Database Server scenarios**)
    - Up to 55% of DRDA, 40% is typical
  - Parallel queries (**Data Warehouse scenario**)
    - Up to 80% of parallel queries
  - Most of **index maintenance** utilities
- **How much Java workload can typically be run on a zAAP?**
  - **WebSphere scenario**
    - Up to 85% of a WebSphere workload
- **Offloads to specialty processors reduce software load and charges on general purpose processors**
  - For sub capacity pricing, the offload must occur at a time that will reduce billable rolling average

Results may vary

Example: Consolidate New Data Warehouse Application on Mainframe

Existing Mainframe

Existing processors:
2 general purpose

Existing Disaster Recovery Site

Existing processors:
Pay for one general purpose processor until disaster switch over

Add 1 LPAR for New Data Warehouse w 42 TB Storage

Prod

1954 MIPS additional workload

Add four processors:
3 zIIP's
1464 MIPS (75%)
1 General purpose
489 MIPS (25%)

And Add Disaster Recovery

Prod

Pay for no additional processors until disaster switchover

3 year cost of acquisition $4.78M

Or add Superdome 9000 Server w 75 TB storage

Prod

*169,998 * RPE's*

Disaster Recovery typically not considered

3 year cost of acquisition $8.24M

** Production RPE's required = 1954 x 87 = 169,998*

zIIP Processors Lower the Cost of Acquisition

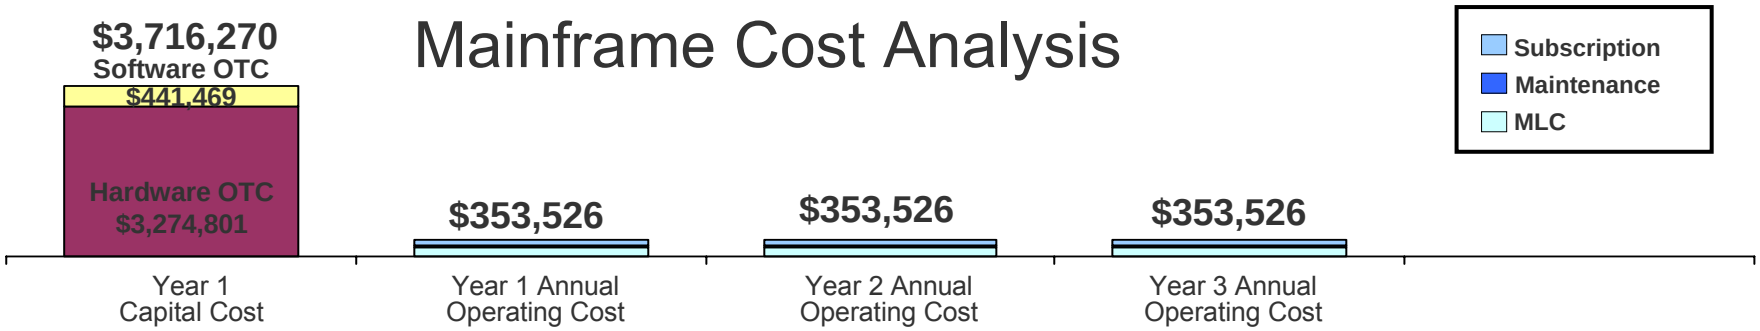

Total cost = $4,776,848

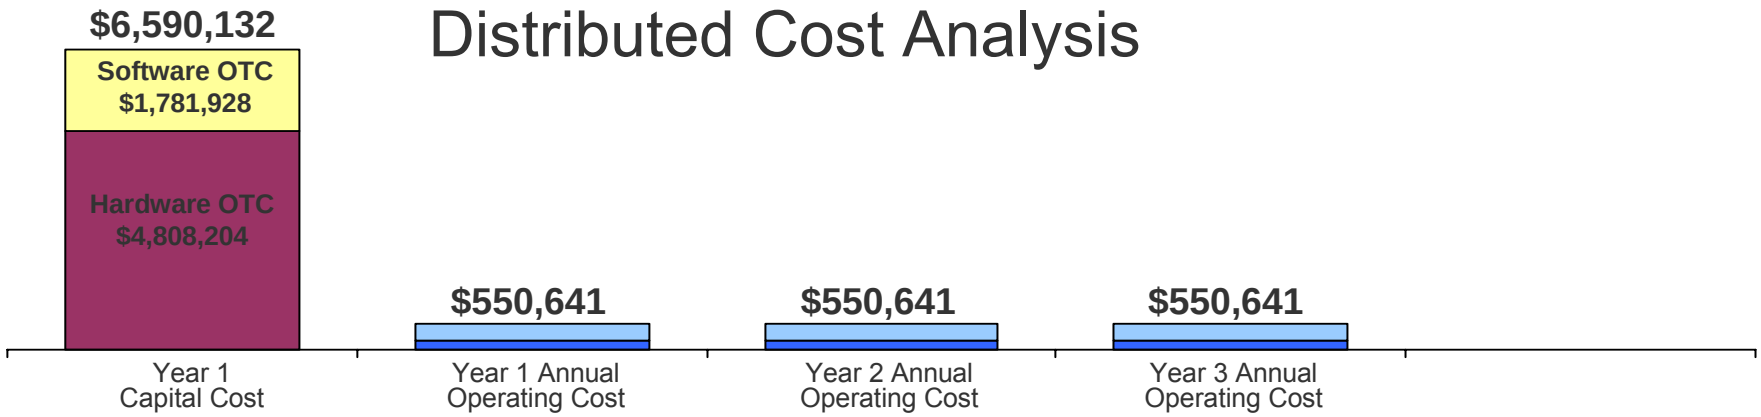

Total cost = $8,242,052

1.7 times more expensive

Data Warehouse Incremental Cost Breakdown

Mainframe Hardware

OTC		ANNUAL	
Z Processors	$1,825,000	Processor Maintenance	$123,540
IBM Storage (42TB)	$1,449,801	Storage Maintenance	0
TOTAL	$3,274,801	TOTAL	$123,540

Mainframe Software

OTC		ANNUAL	
Utilities	$441,469	Utilities S&S	$44,454
		DB2 MLC	$72,240
		QMF MLC	$34,716
		zOS MLC	$78,576
		SubTotal MLC	$185,532
TOTAL	$441,469	TOTAL	$229,986

Distributed Hardware

HP Processors	$1,700,735	Processor Maintenance	$164,044
HP storage (75TB)	$3,107,469	Storage Maintenance	$30,951
TOTAL	$4,808,204	TOTAL	$194,995

Distributed Software

OTC		ANNUAL	
Oracle EE & Utilities	$1,352,000	Oracle S&S	$297,440
Unix	$204,828	Unix S&S	$58,205
HP Storage SW	$225,100		
TOTAL	$1,781,928	TOTAL	$355,645

DB2 V9 Storage Compression Helped Reduce Storage Costs

- **Customers using beta versions of DB2 9 reported 50–80% storage savings**
- **Typical storage cost savings**
  - DB2 storage of 100 TB
  - 59% compression
  - 100 * .59 * $35K/TB = $2.03M
- **Oracle RAC software compression achieves 29%**

"With DB2 9, we're seeing compression rates up to 83% on the data warehouse. The projected cost savings are more than $2M initially with ongoing savings of $500,000 a year."

Michael Henson
DB2 Unix Team Lead, SunTrust Bank, Inc.

DB2 – Better Compression Ratio Than Oracle

- **TPC-H is a well known data warehouse benchmark**
  - Each vendor uses the same tables and same data
  - Oracle published their compression rates for TPC-H tables at the VLDB conference in 2003
  - IBM ran the same tests on the same tables
- **Test results**

Table	Compression Ratio	
	Oracle	DB2
LINEITEM	38%	58% (1.5x better)
ORDERS	18%	60% (3x better)
Entire Database	**29%**	**59% (2x better)**

Example: *Consolidate New WebSphere Application on Mainframe*

Existing Mainframe

Existing processors:
2 general purpose

Existing Disaster Recovery Site

Existing processors:
Pay for one general purpose processor until disaster switch over

Add 3 LPARs for New Web Application

Prod	Dev	QA
900 MIPS additional workload		

Add two processors:
one zAAP
510 MIPS WAS (85%)
one General Purpose
300 DB2 MIPS
90 WAS MIPS (15%)

And Add Disaster Recovery

Prod	Dev	QA

Pay for no additional processors until disaster switchover

3 year cost of acquisition $3.05M

Or add Superdome 9000 Servers

Prod

*82,531 * RPE's*

Dev and QA

82,531 RPE's

** Assume dev and QA is 25% of 900 MIPS total. Then production RPE's required = 900 x .75 x 122 = 82,350*

And Add Disaster Recovery

Prod

82,531 RPE's

3 year cost of acquisition $3.89M

zAAP Processor Lowers the Cost of Acquisition

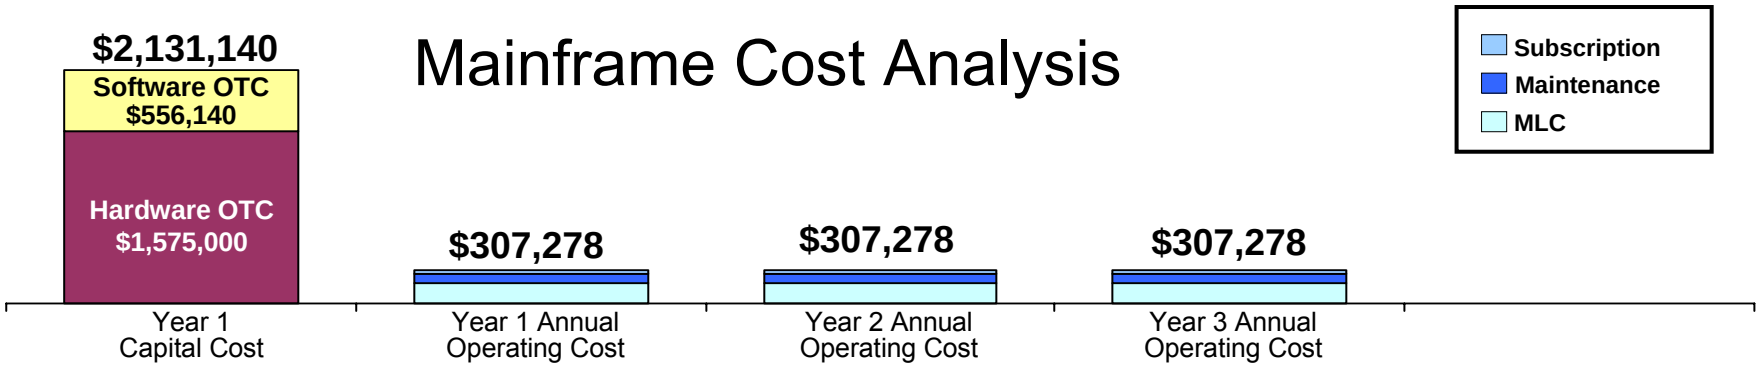

Total cost = $3,052,974

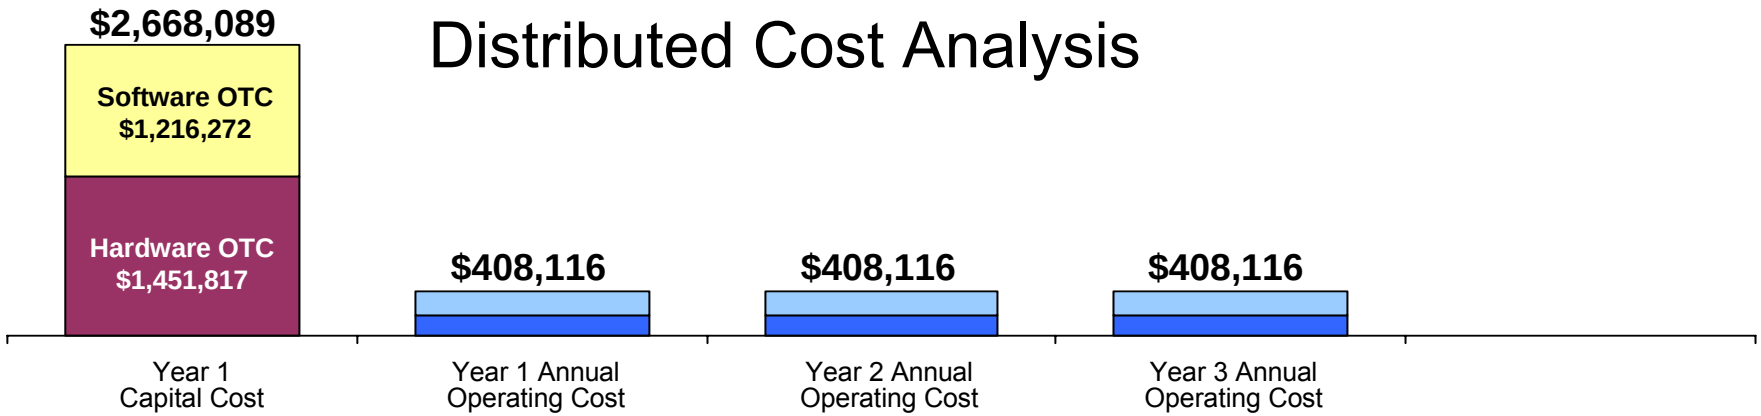

Total cost = $3,892,437

1.3 times more expensive

WebSphere Application Server IncrementalCost Breakdown

Mainframe Hardware

1 GP Processor	$1,450,000	Processor Maintenance	$88,500
zAAP	$125,000		
TOTAL	$1,575,000	TOTAL	$88,500

Mainframe Software

OTC		ANNUAL	
Utilities + WAS	$556,140	Utilities S&S	$44,454
		DB2 MLC	$72,240
		QMF MLC	$34,716
		zOS MLC	$67,368
		SubTotal MLC	$174,324
TOTAL	$556,140	TOTAL	$218,778

Distributed Hardware

OTC		ANNUAL	
3 16x32 Itanium Superdome Servers	$1,451,817	Servers Maintenance	$123,139
TOTAL	$1,451,817	TOTAL	$123,139

Distributed Software

OTC		ANNUAL	
Oracle EE & Utilities	$858,000	Oracle S&S	$188,760
WebSphere	$259,875	WS Maint	$51,975
Unix	$98,397	Unix S&S	$44,242
TOTAL	$1,216,272	TOTAL	$284,977

Integrated Facility for Linux (IFL) Makes Linux Consolidation Even More Attractive

- **Linux Consolidation**
  - Port multiple Linux systems to run as separate images under z/VM
  - Good economics due to typically low utilization of distributed servers
- **Use of IFL engine reduces the price further**
  - IFL specifically limited to Linux workloads
  - Hardware is $125,000 per processor
  - IBM and some other vendors charge software licenses per one IFL processor (not per image)
- **IFL Requirements**
  - z9-109, z990, z900, z890 or z800 hardware platform
  - No z/OS requirements
  - No limit on the number of IFLs

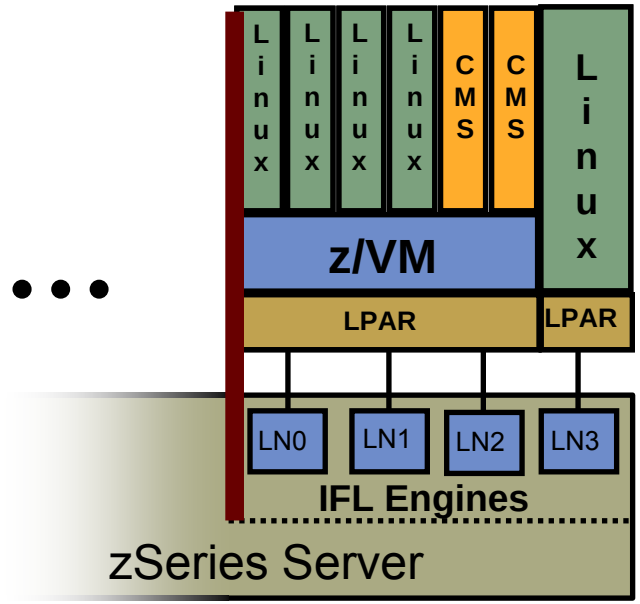

Example: IGS US Consolidated 62 Distributed Images to zLinux Images on an IFL

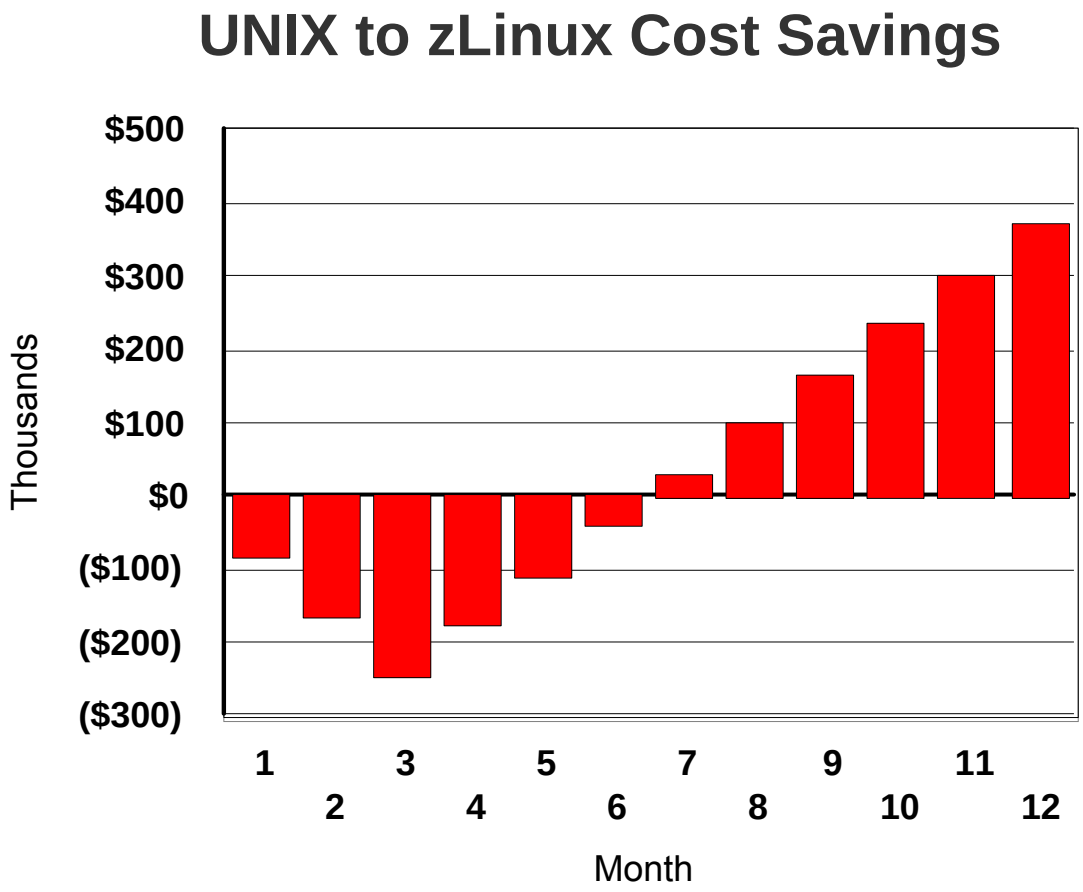

(Costs Savings are driven primarily by $89K monthly labor savings)

- ✓***Broke-even after 6 months***
- ✓***>$2.5M saved in 3 years***
- ✓***In Europe, €524K headcount saved in Year 1***

The Economics of Linux Workload Consolidation

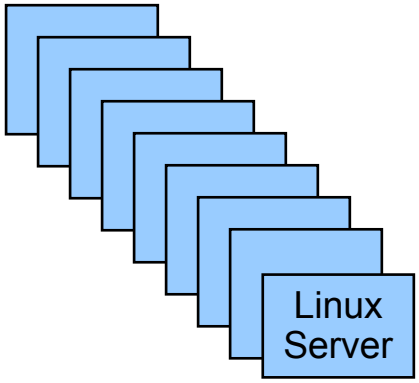

62 Linux servers with low utilization

62 @ $5,000 = $310,000

Plus 62 middleware licenses

Plus $6,500 x 62 = $403,000/yr labor

One IFL processor with high utilization

1 @ $125,000 = $125,000

Plus one middleware license

Little additional labor

Potential Savings from Linux Consolidation on System z

60 Linux Servers

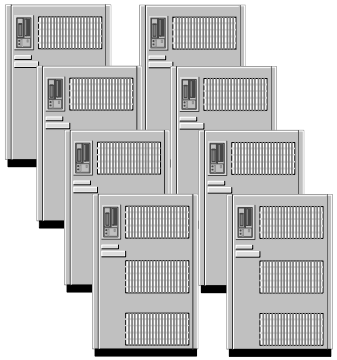

1 IFL

$1.8M saving over 3 years

	Distributed Linux/Intel @ low utilization				*Mainframe IFL @ high utilization*			
	Unit cost	*Quantity*	*Sub Total*	*3 year total*	*Unit cost*	*Quantity*	*Sub Total*	*3 year total*
Hardware & OS - every 3 years	**$4,000**	**60**	**$240,000**	**$240,000**	**$125,000**	**1**	**$125,000**	**$125,000**
HW Maintenance			**Included**		**$19,944**	**1**	**$19,944**	**$39,888**
VM virtualization			**N/A**		**$22,500**	**1**	**$22,500**	**$22,500**
VM S&S (25%)			**N/A**		**$5,625**	**1**	**$5,625**	**$16,875**
Annual Linux support	**$1,000**	**60**	**$60,000**	**$180,000**	**$14,000**	**1**	**$14,000**	**$42,000**
OTC Software license – WAS*	**$4,000**	**60**	**$240,000**	**$240,000**	**$4,000**	**1**	**$4,000**	**$4,000**
WAS S&S for 2 years	**$800**	**60**	**$48,000**	**$96,000**	**$800**	**1**	**$800**	**$1,600**
Annual labor for support	**$6,500**	**60**	**$390,000**	**$1,170,000**	**$6,500**	**1**	**$6,500**	**$19,500**
Annual power & cooling	**$920**	**60**	**$55,188**	**$165,564**	**$920**	**1**	**$920**	**$2,759**
Grand Total				$2,091,564				$274,122

** IBM WebSphere Application Server for Linux*

Saves $16+ Million with Linux on System z

Nationwide®
On Your Side™

Vastly improved TCO, Speed & Simplification

- **Problems:**
  - High TCO including data center power and floor space scarcity (new facility would cost $10M+)
  - Long server provisioning process

- **Solution:**
  - 350 servers virtualized with 15 z990 IFLs, supported by 3 staff
    - 12 mission critical applications with 100,000+ users/day
  - 50% reduction in Web hosting monthly costs, 80% reduction in floor space & power conservation
  - 50% reduction in hardware & OS support efforts; significant savings on middleware costs
  - Fast deployment (4 months)
  - Significantly faster provisioning speed (months → days)
  - Simple, robust mainframe high availability & disaster recovery

Québec Government Runs Oracle at IFL Prices

- **Consolidated 190 Oracle Databases (9i and 10g) onto a z9-EC with IFL's**
  - Reduced cost of hardware and software by 30%
  - Better database loading performance due to higher I/O bandwidth
  - Each administrator could manage 100 database instances
  - Easy migration
    - One migration per day
    - Create new Linux server in 30 min (vs 1 week – 3 months)
    - Clone Oracle DB instance in 30-45 min (vs 10 – 14 hours)
    - Unload/load
  - Inherit benefits of z platform – workload management, availability, disaster recovery
  - Expect to migrate 200 more Oracle databases per year

NEW! Execute .NET Code on the Mainframe at IFL Prices

Visual MainWin for J2EE

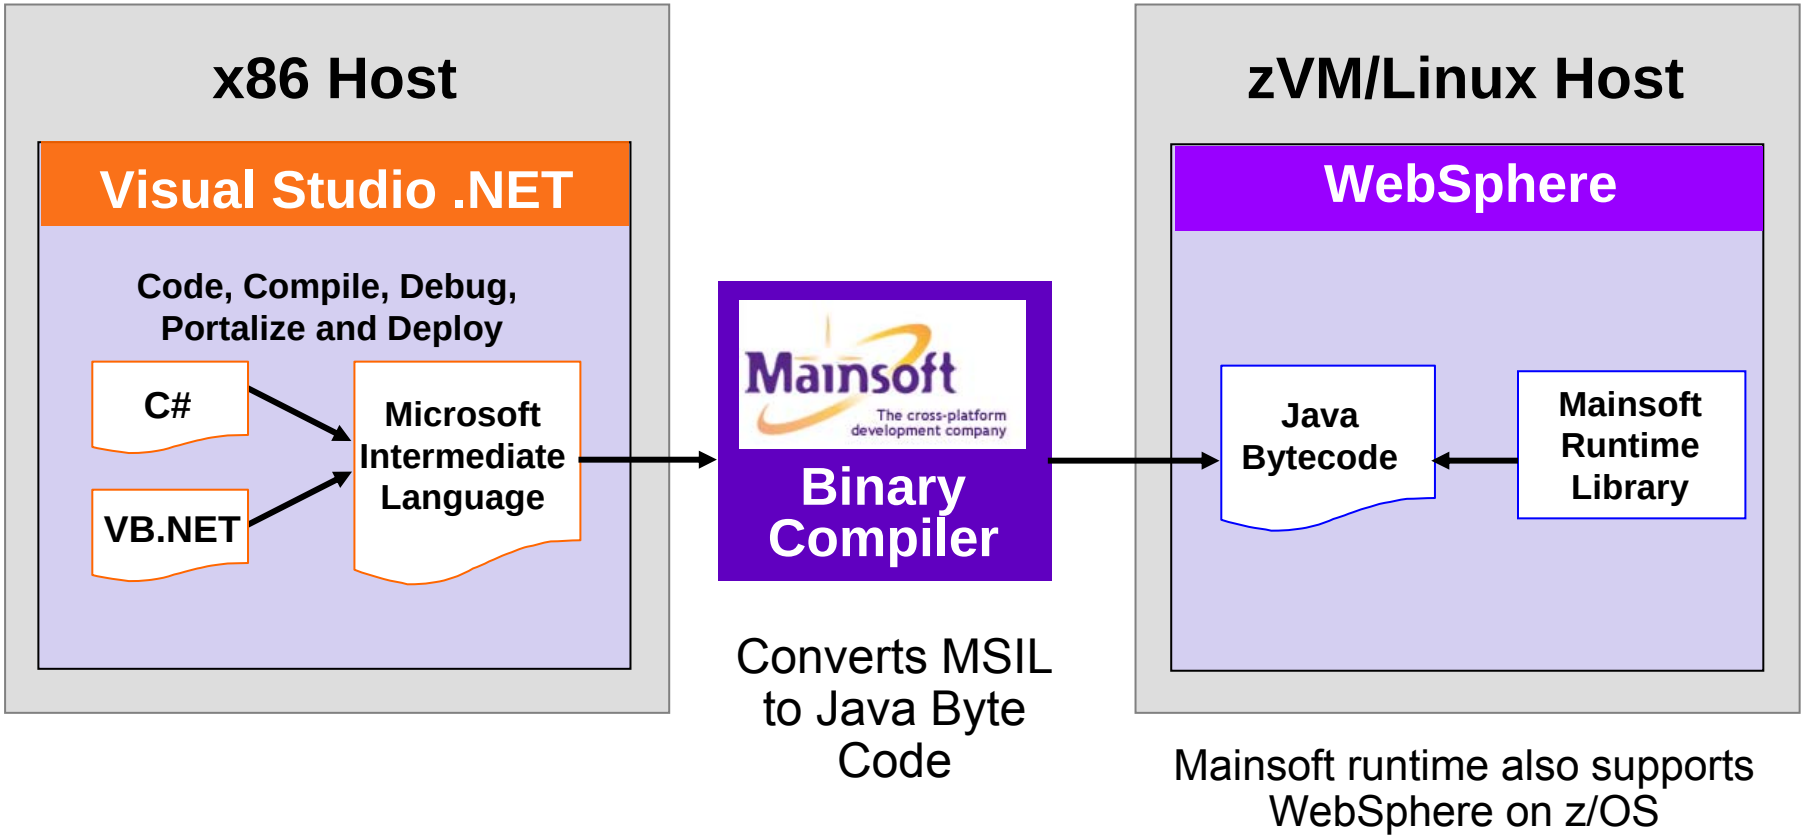

Contact: Ron Johnsen – VP WW Sales, ronj@mainsoft.com USA 408 200 4023

Distributed Power Costs Have Become a Major Issue

- **According to the Wall Street Journal, distributed server farms can generate as much as** 3,800 **watts per square foot**
  - In 1992 it was 250 watts/sq foot
  - By comparison, a System z9 consumes **107-312** watts per square foot – *one tenth or less the amount*
  - Turning on an IFL processor consumes 75 additional watts
  - Cooling cost is roughly an additional 60% of the power cost

- **More than half of all serious outages are now caused by power problems***
  - Room temperatures averaging 92ºF lead to erratic machine behavior
  - *"Power-related problems in 2005 will cause 4 of the 20 major failures, up from 2 of 20 last year"* (The Uptime Institute)

**Source: recent AFCOM survey of 200*

Mainframe Scale and Power Efficiency

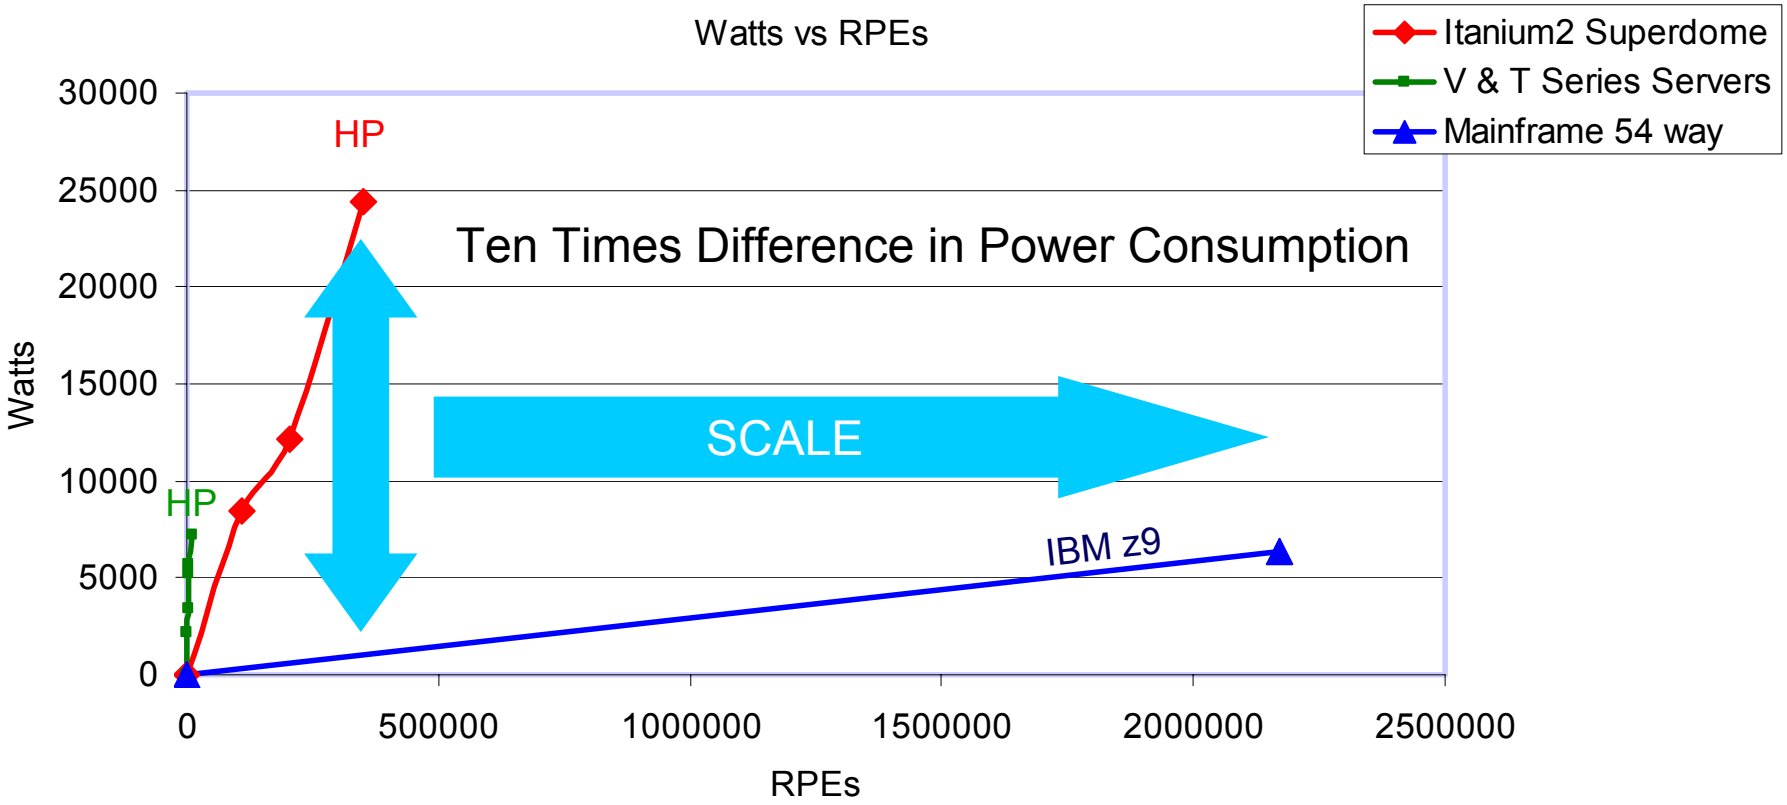

Source for HP Servers: Ideas International, Nov 06
Note: Uses equivalence ratio of 122 RPE's per MIP

Do the Math

- **HP Itanium 2 Superdome 9050 (64ch/128co) consumes a maximum of 24,382 watts**
  - 24,382 X .15 X 24 X 365 = **$32,038** per year for electricity
- **Mainframe with similar computing capacity consumes 2,500 watts**
  - **$3,204** per year for electricity
  - Power cost is $28,834 per year less
- **Similar savings on cooling capacity**
  - Cost of cooling is 60% to 80% the cost of power
  - Superdome total **$51,261** per year vs Mainframe **$5,126**

Fractional Availability Improvements Are Important

Example 1: Financial Services Company

- $300B assets, 2500+ branches, 15M customers
- Retail banking, loans, mortgages, wealth management, credit cards
- CRM System – branches, financial advisors, call centers, internet
- Number of users – 20,000+

	Unix/ Oracle	*zSeries/ DB2*
Availability %	**99.825%**	**99.975%**
Annual outage	**15h 20m**	**2h 11m**
Cost of Downtime	$45.188M	$3.591M

Sources: ITG Value Proposition for Siebel Enterprise Applications, Business case for IBM eServer zSeries, 2004 & Robert Frances Group, 2005

Financial Impact of Downtime Per Hour

Industry segment	*Cost*
Energy	**$2,818K**
Telecommunications	**$2,066K**
Manufacturing	**$1,611K**
Financial	**$1,495K**
Information Technology	**$1,345K**
Insurance	**$1,202K**
Retail	**$1,107K**
Pharmaceuticals	**$1,082K**
Banking	**$997K**
Consumer Products	**$786K**
Chemicals	**$704K**
Transportation	**$669K**

Replace Third Party Tools to Reduce Costs

- **LabCorp**
  - 35 products replaced - includes RMM, TWS, SCLM and DB tools and AD tools
  - About 700 MIPS
  - $12M saved
- **Putnam Investments**
  - Over 20 products replaced at 2 sites - includes RACF, RMM, TWS, SCLM, SA390, GRS and DB2 Suite of tools
  - About 1500 MIPS
  - $Millions saved
- **Hennepin County**
  - Products replaced - includes RACF, TWS, SA390,DB2PM, TDS
  - About 1100 MIPS
  - $3M in savings
- **Major automotive manufacturer**
  - Doubled MIPS from 600 to 1200
  - Annual savings of $1.8M
- **Typically 30-50% lower run-rate after initial ROI period**

A typical customer engagement replacing BMC tools

Original Product	*IBM Replacement*
Mainview for z/OS	IBM Tivoli OMEGAMON XE for z/OS
Mainview Stop X37	IBM Tivoli Allocation Optimizer
Mainview Explorer and Alarm Manager	IBM Tivoli OMEGAMON DE on z/OS
Mainview EasyHSM and StorageGuard	IBM Tivoli Storage Optimizer
Control-M & R	IBM Tivoli Workload Scheduler
Auto Operator	IBM Tivoli System Automation
Mainview for CICS	IBM Tivoli OMEGAMON XE for CICS / CICS PA
Mainview for DB2	IBM Tivoli OMEGAMON XE for DB2
Mainview for IMS	IBM Tivoli OMEGAMON XE for IMS
UltraOpt for IMS	IMS NETWORK COMPESS FACILITY V1
Image Copy Plus for IMS	IMS HIGH PERFORMANCE IMAGE COPY
Unload Plus for IMS	IMS HIGH PERFORMANCE UNLOAD V1
Prefix Resolution Plus for IMS	IMS HIGH PERFORMANCE PREFIX RESOLUTION V3
Load Plus for IMS	IMS HIGH PERFORMANCE LOAD
Secondary Index Utility / EP	IMS INDEX BUILDER V2.3

People Expense has Become the Dominant Component of TCO

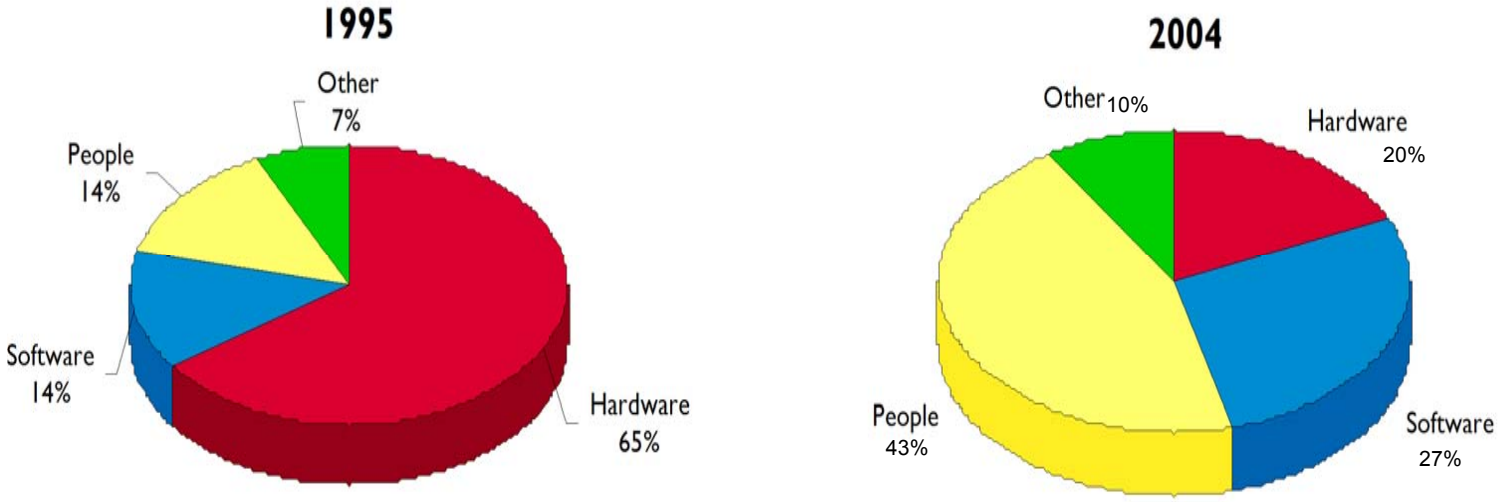

Based on IBM Scorpion customer analyses

Since 2000, Labor Costs Have Exceeded the Cost of All Servers … and are *Still* Growing

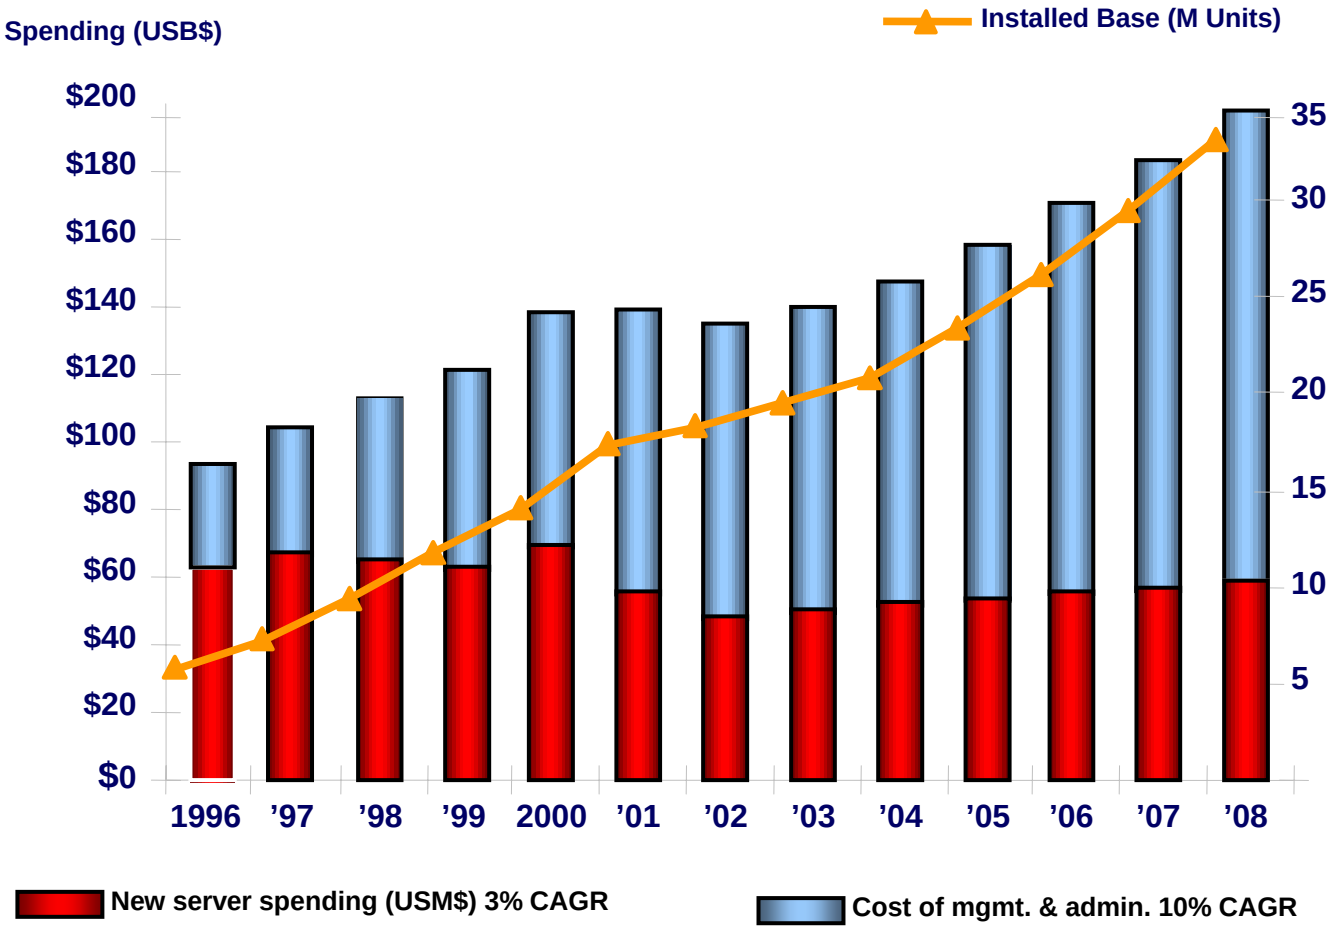

The cost of labor scales linearly with the number of distributed servers

Distributed scale-out is costly!

Source: IDC

Data Center Staffing Levels for System z Have Not Increased Despite Large Increase in MIPS

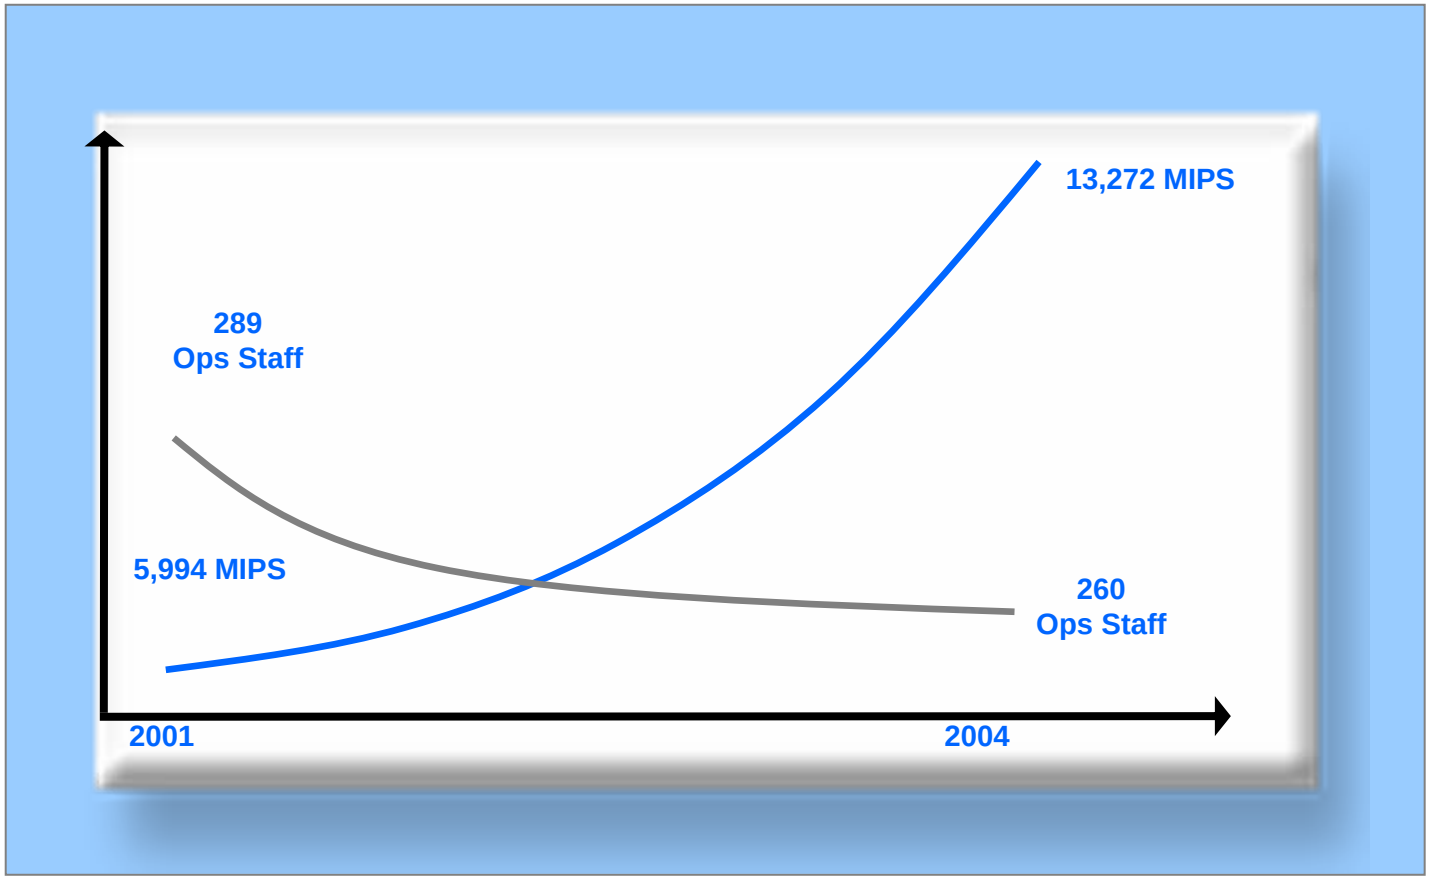

$\frac{5,994}{289}$ = 21 MIPs/HC

$\frac{13,272}{260}$ = 51 MIPs/HC

Source: Gartner

A Comparison of Labor Costs for Two Environments That Execute Roughly Equivalent Workloads

Topic	System z- 3,192 MIPS	900 Distributed Servers
Operations	$105K ……10% of 6 FTEs	none
Customer Engineers	$52K……. 0.3 FTEs $50K LAN charges $35K z- charges	$400K SUN charges $300K LAN charges $40K p- charges $100K HP charges
Systems Engineers	$551K ….. 3.15 FTEs	$5,250K…..30 FTEs (Operations in the Systems charge)
Security Admin	None	$600K
Total	$793K	$6,690K

In this case, System z requires 1/8 the labor costs of the distributed environment

Source: IBM SWG Data Center

Labor Cost Per Transaction on System z is Decreasing

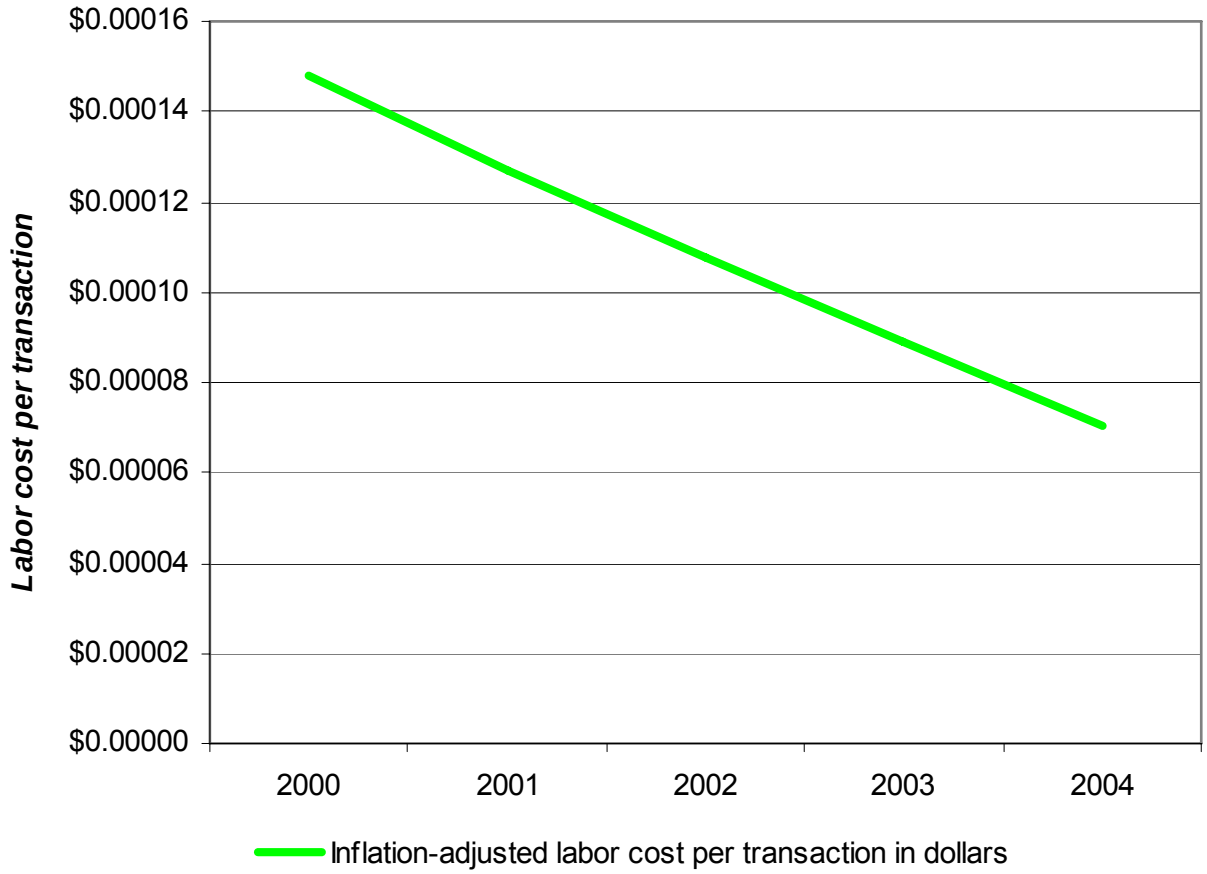

Source: IBM Global Services UK

16.9% decrease per year

What makes the price go down?

Increasing workloads
Data-center-in-a-box design reduces need for labor
Scalability of the mainframe
Ease of incremental upgrade
Inherent reliability of the mainframe
Fewer repairs and patches
Intelligent Workload Management
Including CICSPlexSM
Minimal security risks & breaches
IBM integration, testing & support

Trade-In Value Reduces Mainframe Net Present Value Costs

- **Upgrade to next generation mainframe**
  - Specialty processors are upgraded to next generation free of charge
  - Growing and current customers are typically charged a discounted upgrade price for processors even if the upgrade is to a new generation
  - Same financial effect as trade-in value on existing MIPS
- **Upgrade to next generation distributed systems**
  - Life time of 3 to 5 years
  - Must repurchase existing processor capacity plus any growth
- **Long term TCO implications can be important**

Java Application Hosting Example Considering 50% Trade in on General Purpose Processor

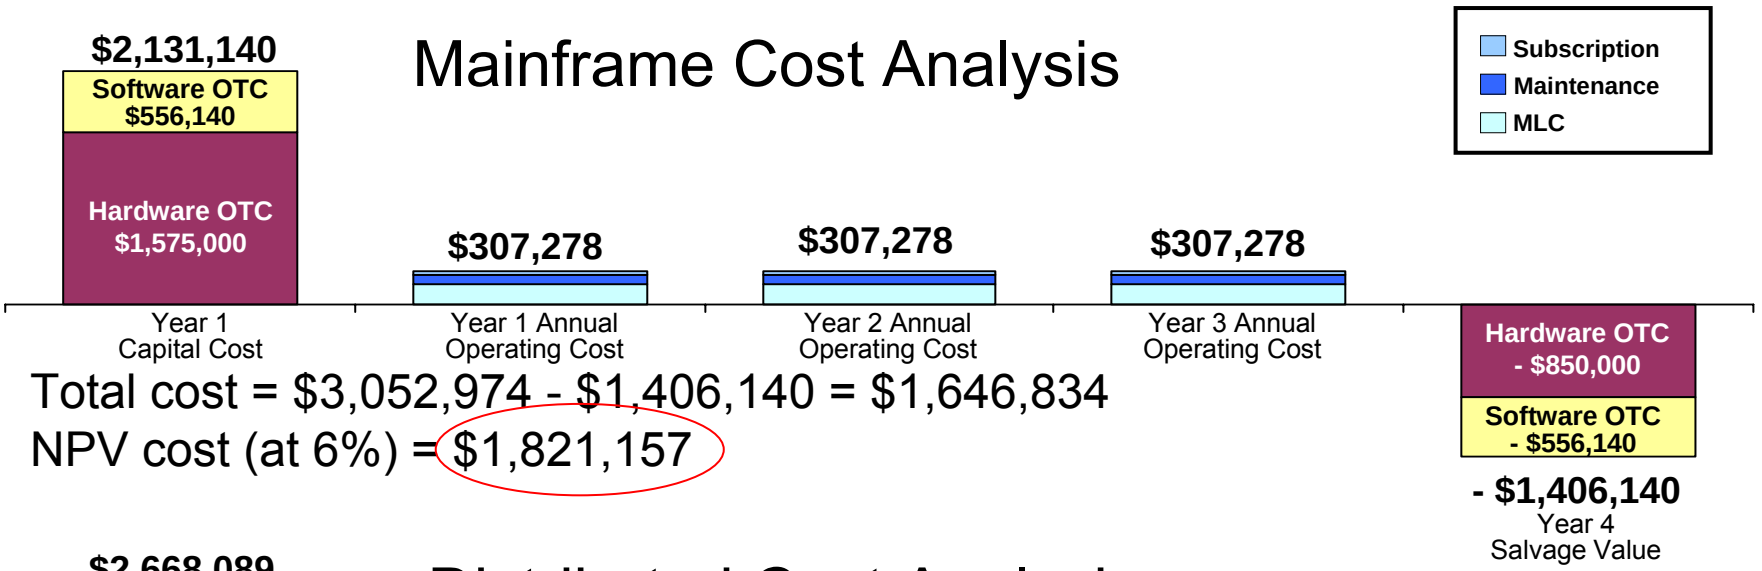

Total cost = $3,052,974 - $1,406,140 = $1,646,834

NPV cost (at 6%) = $1,821,157

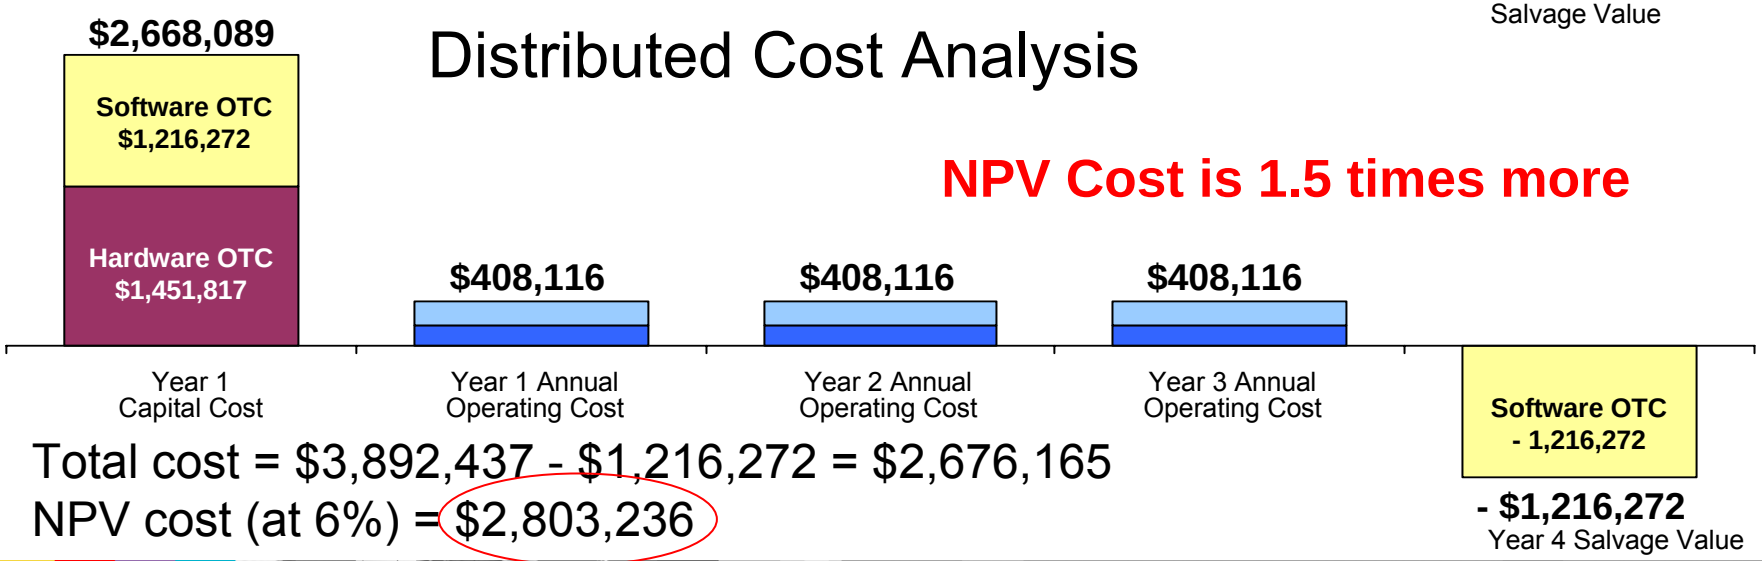

NPV Cost is 1.5 times more

Total cost = $3,892,437 - $1,216,272 = $2,676,165

NPV cost (at 6%) = $2,803,236

Conclusion: Total Mainframe Hardware, Software & Labor Costs Have Been Cut in Half in 4 Years

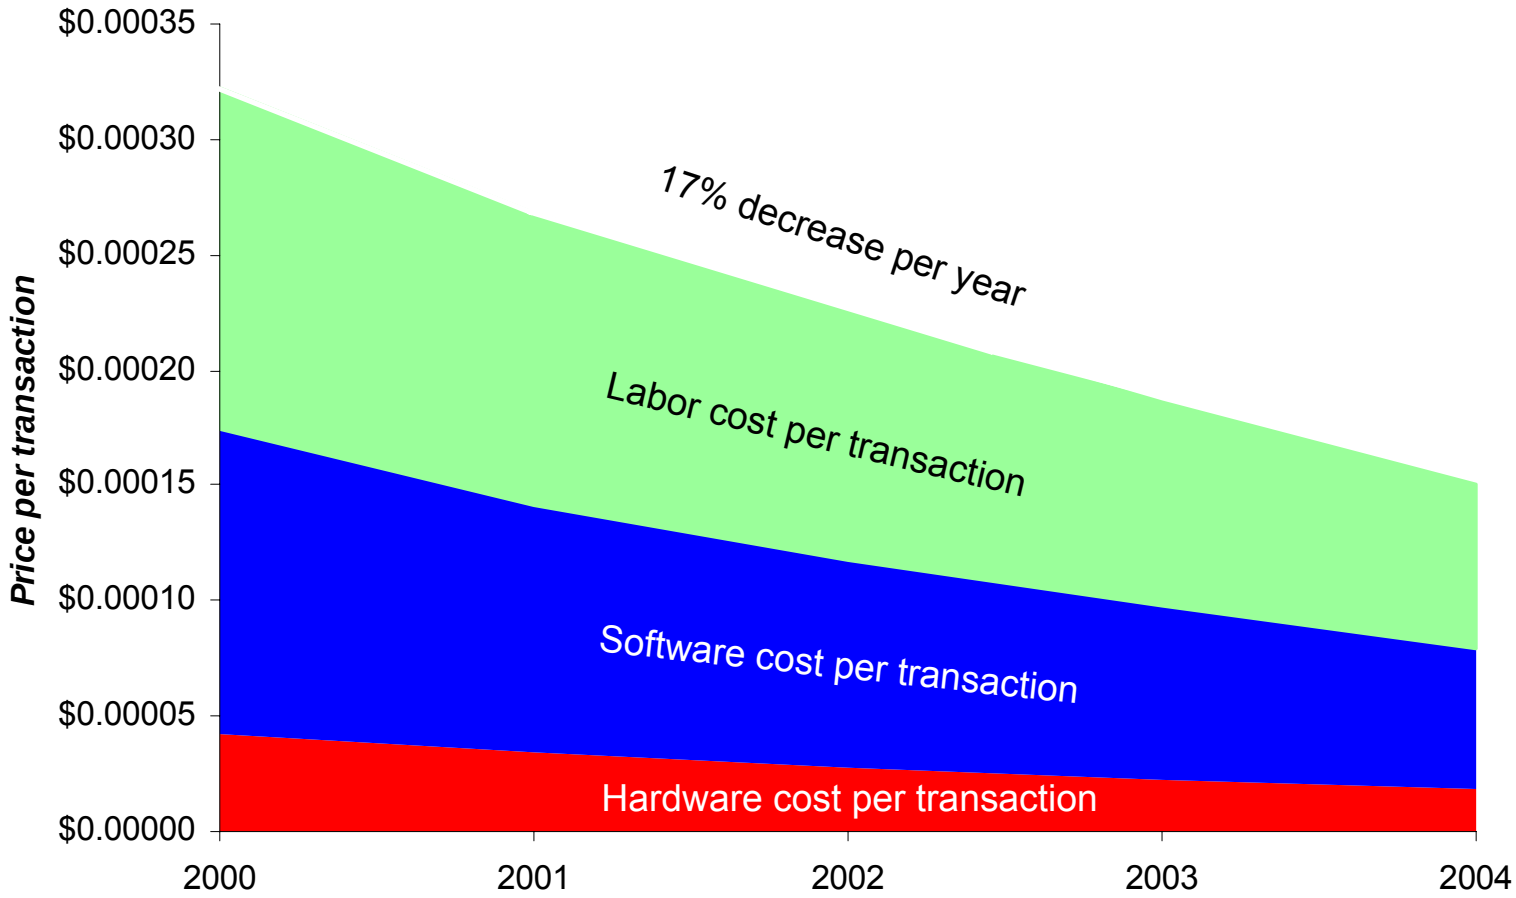

GAD TCO Study – a WebSphere Banking Example

Distributed servers have higher service, monitoring and support costs and cost more to develop and implement

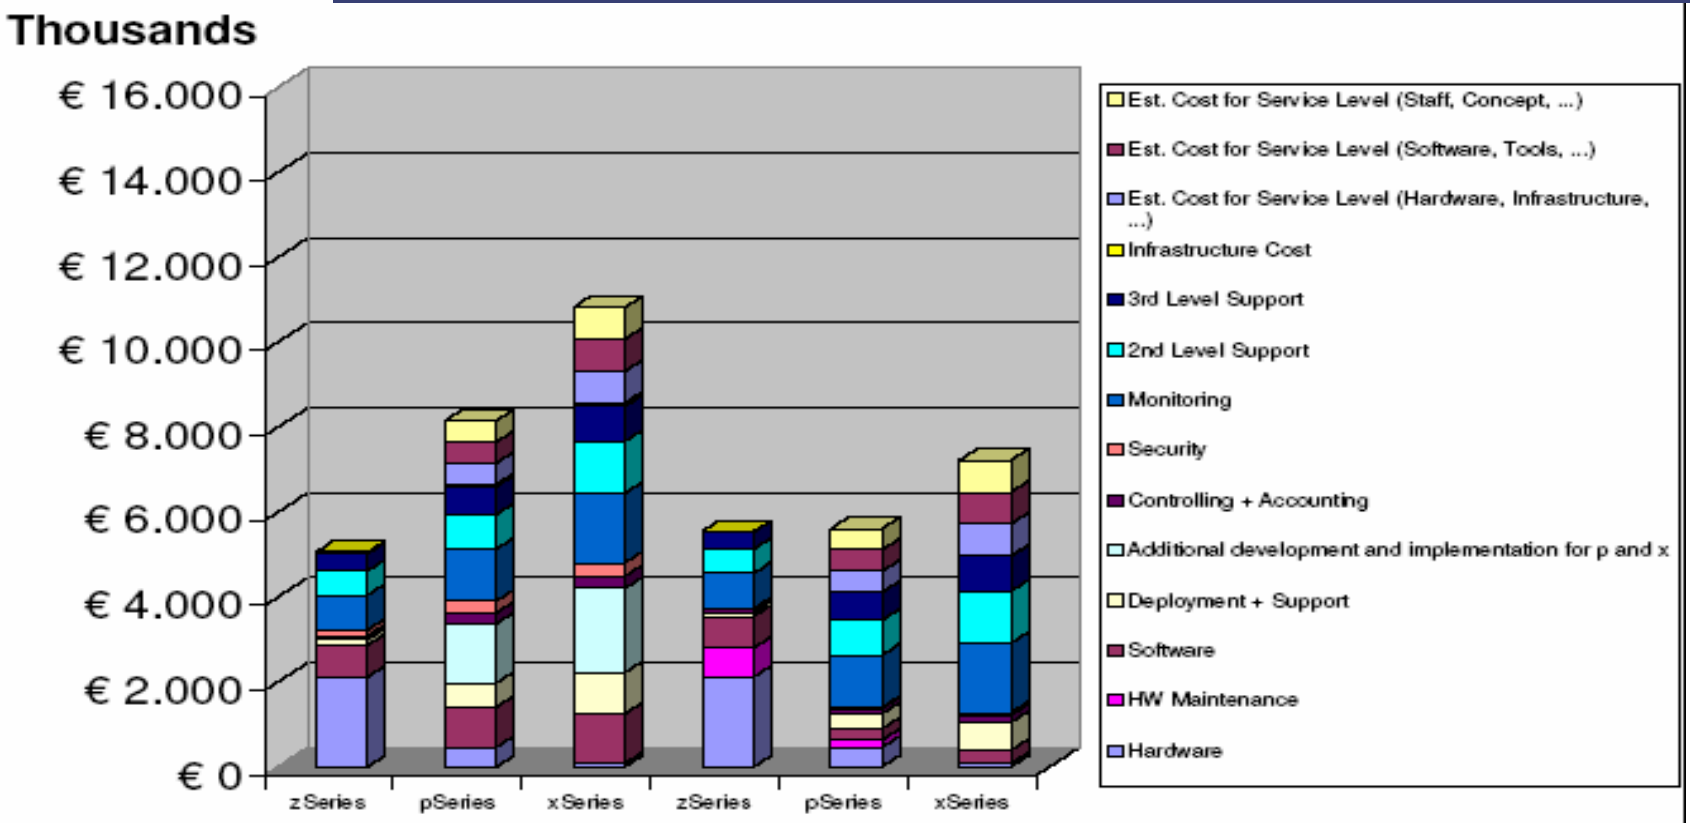

Col's 1-3: Year 1 Col's 4-6 : Year 2 -4

Tale of Two Customers

	Baldor	Welch's
Supplier	IBM	Dell
Moved From....	3 Mainframes and 8 Unix Servers	S/390 and AS/400
Moved to...	1 z990 System z Server	100 Intel Servers
Virtualization	z/VM	VMWare
Decision to Completion Time	Approximately 6 months	Started sometime before June 2005 "...project will continue into 2007"
IT Staff	Down to 38	50
IT Spending	1.2% of Sales (and still declining....now down to 0.9%)	About 2.5% of Sales
Max Power consumption	15.8 kW	48.4 kW

Three years ago, Baldor's IT director had investigated migrating to a Windows server environment with cluster fail-over. *"We thought we were going to save a ton of money,"* but the systems crashed all the time, he noted, and the idea was quickly abandoned.

"We have a very stringent requirement of being up all the time … Weighing heavily in support of the mainframe was its track record. There hadn't been any mainframe downtime since 1997"

When Does 30% Incremental Workload Growth Cost Less?

3 Year TCO with HI-RAS requirement

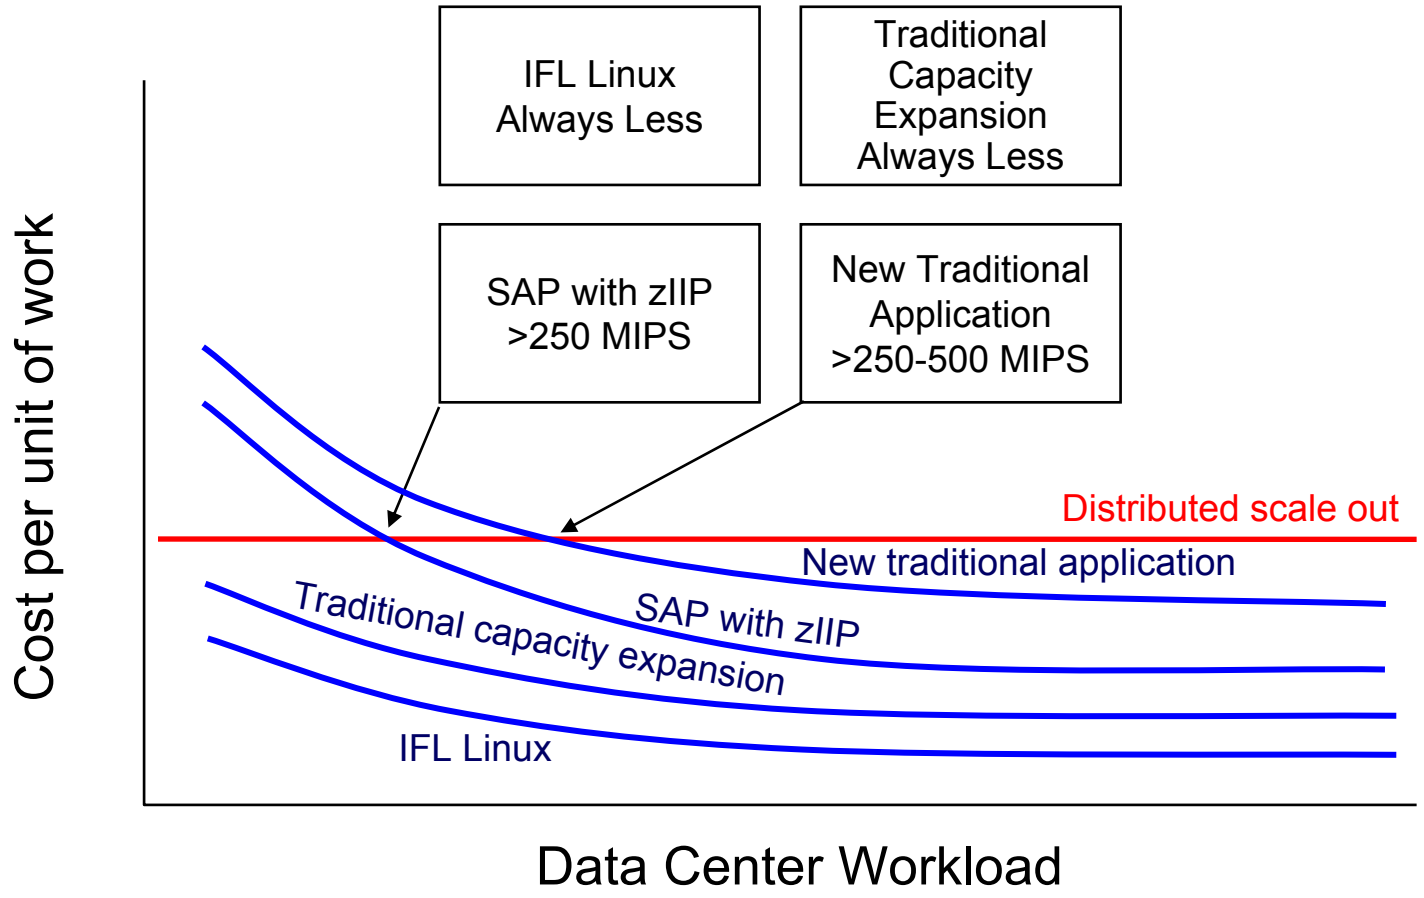

Source: Eric Kutcher, McKinsey Analysis

Summary

- The proper comparison between mainframe and distributed is not a single application benchmark
- The proper comparison is a distributed data center versus a mainframe, running high volume mixed workloads
- Under this comparison, mainframes have significant cost advantages

How Customers Can Get the Lowest TCO on the Mainframe

1. **Move past the "sweet-spot" to realize lower prices**
   - Grow core-business MIPS
2. **Use the latest technology and pricing models**
   - Upgrade to System z
   - Utilize specialty processors
   - Utilize sysplex aggregation
   - Exploit sub-capacity pricing
   - Execute an ELA
3. **Maximize utilization**
   - Drive mainframes at 90+% utilization, 24 hours by 7 days
   - Consolidate workload onto System z
4. **Minimize other costs**
   - Minimize software tool costs
   - Minimize outages and security breaches…
5. **Stop spiraling labor costs**