



The Modern Mainframe... At the Heart of Your Business

**Delivering Next Generation Solutions at the
Lowest Cost**



© 2006 IBM Corporation

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



	Servers	Reliability	Utilization	Staff
First move: Implemented distributed computing architecture that became too difficult to monitor, maintain, upgrade and scale	<ul style="list-style-type: none"> ■ 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 Acquisition =
Cost of hardware +
Cost of software +


Let's Break Down the Elements of Cost

Total Cost of Ownership =
Cost of hardware +
Cost of software +
Environmentals +
Cost of labor +
Financial terms

10 - Delivering Next Gen Solutions at the Lowest Cost v1.5.ppt

10

Mainframe Hardware Cost is Decreasing




10 - Delivering Next Gen Solutions at the Lowest Cost v1.5.ppt

11


“Price Per MIP” Does Not Tell The Whole Story

54 way SMP


Pay for these general purpose processors
“Price per MIP” →



Pay for zAAP, zIIP, and IFL processors
at a reduced rate
(~9% of price per MIP) →



Do not pay for pre-installed capacity on
demand processors
until used →



No Additional Charge For System Assist Processors and RAS Processors



PLUS UP TO 336 I/O PROCESSORS
at feature cost

No additional charge for these processors

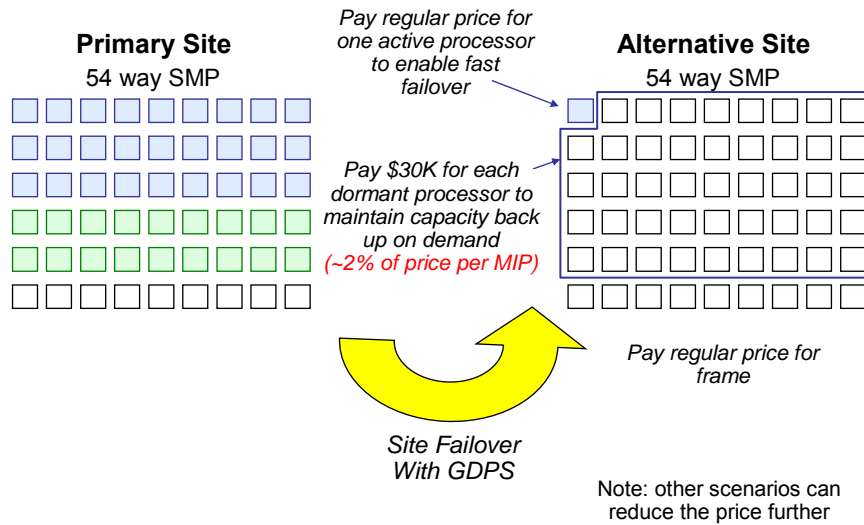
PLUS 10 MORE Dedicated System Assist PROCESSORS



Plus Reliability and Serviceability Processors



Disaster Recovery – Fast Failover For Less



10 - Delivering Next Gen Solutions at the Lowest Cost v1.5.ppt

14

Data Center 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

Resulting in tremendous efficiencies

- 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

No extra charge for this deep pre-integration!

10 - Delivering Next Gen Solutions at the Lowest Cost v1.5.ppt

15

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

— Inflation-adjusted IBM software price per transaction

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



10 - Delivering Next Gen Solutions at the Lowest Cost v1.5.ppt

18

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

10 - Delivering Next Gen Solutions at the Lowest Cost v1.5.ppt

19

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

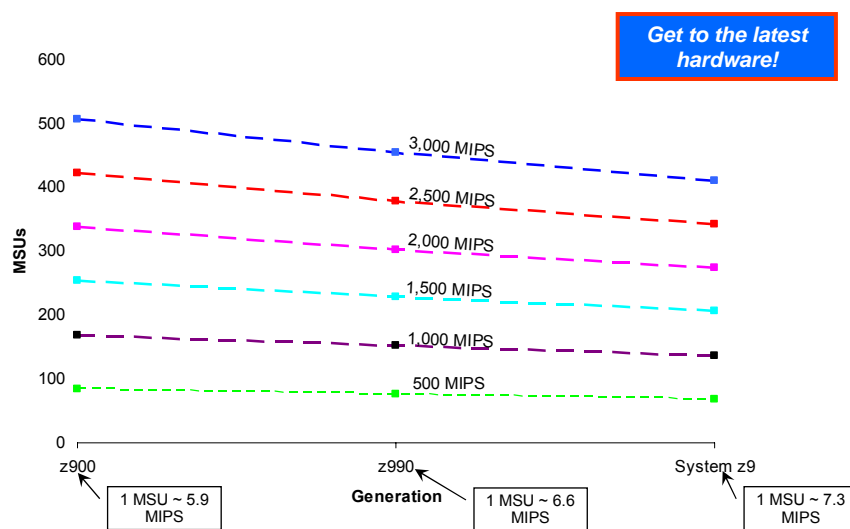
Various ratings online at:

MSUs <http://ibm.com/zseries/library/swpriceinfo/hardware.html>

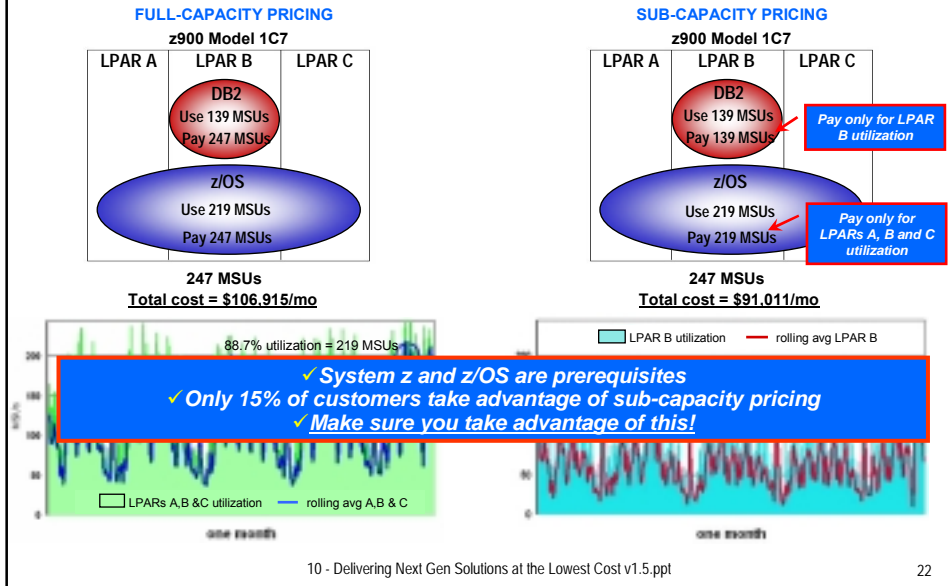
LSPR MIPS <http://ibm.com/zseries/lspr>

Hardware SRM Constants <http://ibm.com/zseries/srm>

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




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



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

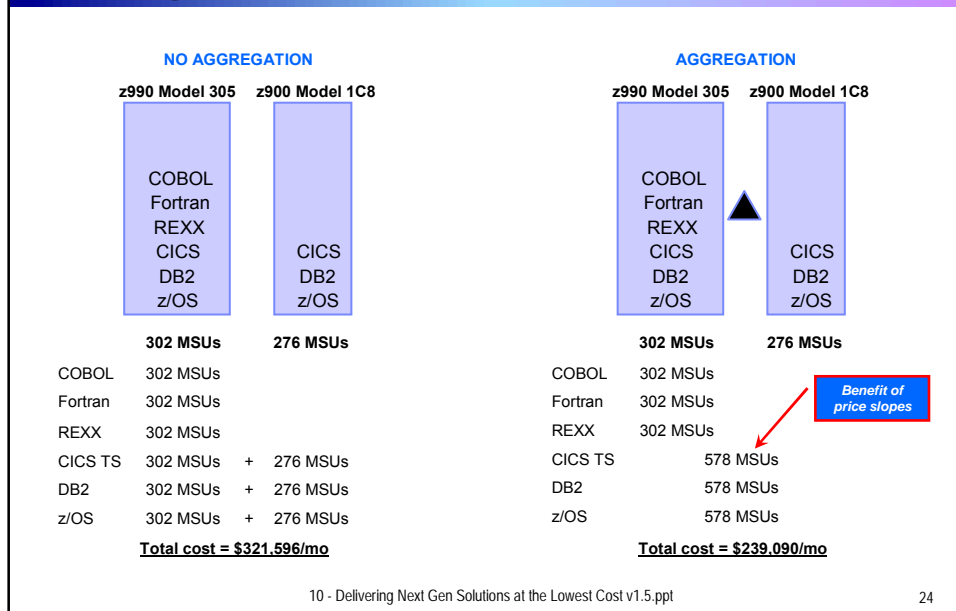


Upgrade to
current generation
z9 Enterprise Class
2094-701

	MLC (PSLC)			MLC (VWLC)	If No Tech Div
Database (-3 Generations)	\$15,378	DB2 UDB V6	→	DB2 9	\$19,868
Transaction Processing (-3 Generations)	\$14,733	CICS ESA V4	→	CICS TS V3.1	\$22,061
Operating System (-2 Generations)	\$46,485	OS/390 Base	→	z/OS VI Base	\$27,633
	\$76,596			\$69,562	\$78,810

* PSLC – Parallel Sysplex Licensing Charge, VWLC – Variable Workload Licensing Charge

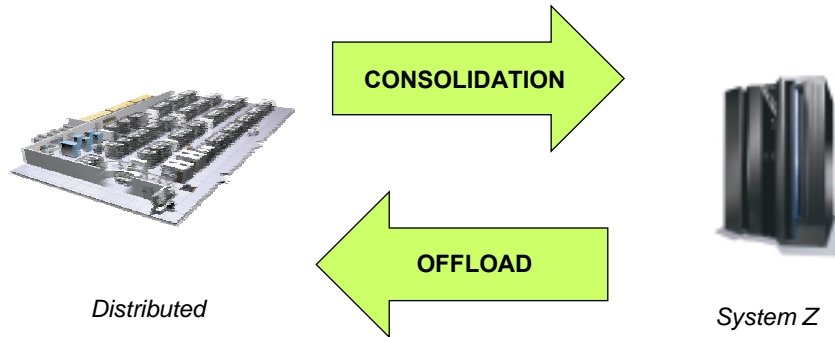
Example of Sysplex Aggregation, Saving >\$82K MLC



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

TCO Comparisons




10 - Delivering Next Gen Solutions at the Lowest Cost v1.5.ppt

26

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" zIIP, zAAP, and IFL further improve consolidation economics
 - ▶ Database, WebSphere, Linux



10 - Delivering Next Gen Solutions at the Lowest Cost v1.5.ppt

27

How Much Workload is zAAP or zIIP- able?

- How much DB2 workload can typically be run on a zIIP?
 - ▶ Queries received via DRDA Remote Access Protocol (**Database Server scenarios**)
 - 40% is typical
 - ▶ Parallel queries (**Data Warehouse scenario**)
 - Up to 80%
 - ▶ Some of **index maintenance** utilities
- How much Java workload can typically be run on a zAAP?
 - ▶ **WebSphere scenario**
 - Up to 85% of a WebSphere workload
- How much Linux workload can typically be run on an IFL?
 - 100%
- 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

TCO Comparisons

Earlier today we saw that strategic hosting of incremental growth on System z with zIIP costs less than HP/Oracle for

- ▶ Data Warehouse
- ▶ SAP Data Server

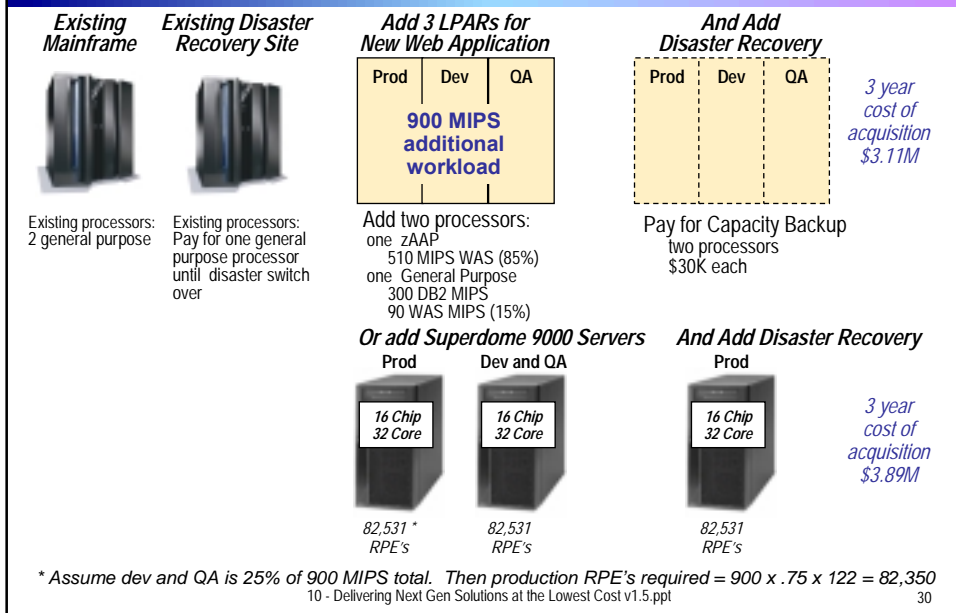
We also saw examples of Linux roll-up consolidation reducing cost

Let's consider a web application that can exploit zAAP




IBM

Example: Consolidate New WebSphere Application on Mainframe



zAAP Processor Lowers the Cost of Acquisition



WebSphere Application Server Incremental Cost Breakdown

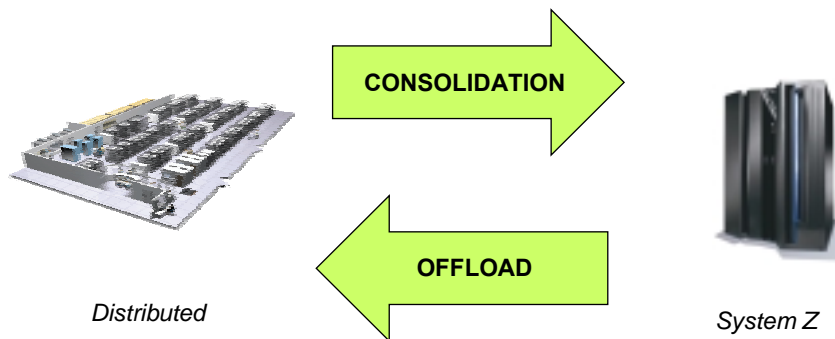
Mainframe Hardware				Mainframe Software			
OTC		ANNUAL		OTC		ANNUAL	
1 GP Processor	\$1,450,000	Processor Maintenance	\$88,500	Utilities + WAS	\$556,140	Utilities S&S	\$44,454
zAAP	\$125,000			DB2 MLC	\$72,240	QMF MLC	\$34,716
2 DR Processors	\$60,000			zOS MLC	\$67,368	SubTotal MLC	\$174,324
TOTAL	\$1,635,000	TOTAL	\$88,500	TOTAL	\$556,140	TOTAL	\$218,778

Distributed Hardware				Distributed Software			
OTC		ANNUAL		OTC		ANNUAL	
3 16x32 Itanium Superdome Servers	\$1,451,817	Servers Maintenance	\$123,139	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	\$1,451,817	TOTAL	\$123,139	TOTAL	\$1,216,272	TOTAL	\$284,977

10 - Delivering Next Gen Solutions at the Lowest Cost v1.5.ppt

32

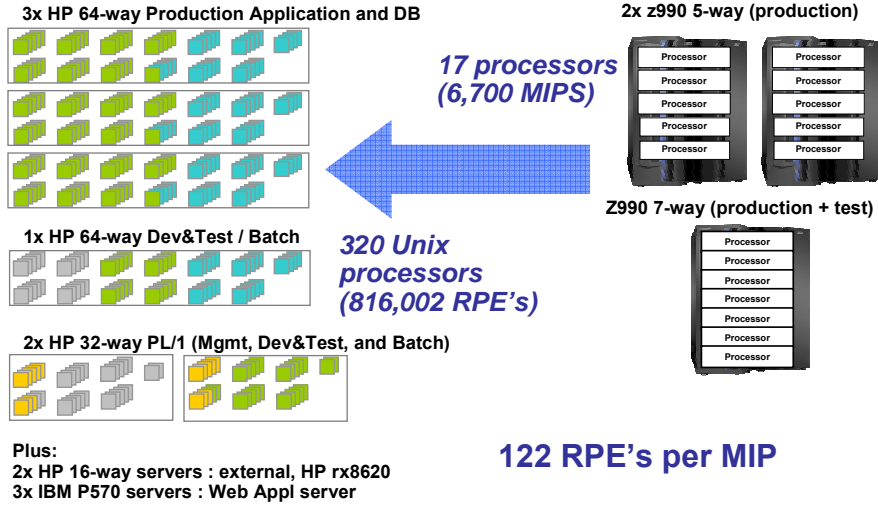
TCO Comparisons



10 - Delivering Next Gen Solutions at the Lowest Cost v1.5.ppt

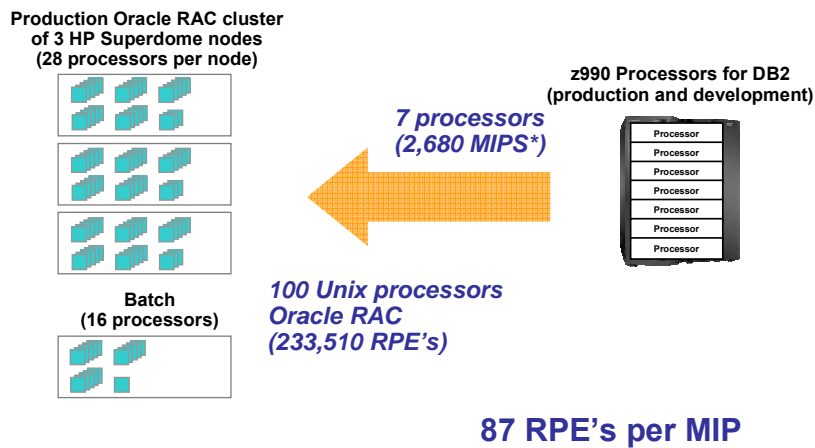
35

Asian Financial Services Customer Offload Project - Overall



No disaster recovery

Asia Pacific Financial Services Customer Offload Project – Database




No disaster recovery

* DB2 is estimated to be 40% of total workload


Did They Save Money by Offloading?

- Competitors told the customer they would save \$22M over 4 years
- IBM analysis determined that the offload solution would actually cost **\$66M more** than growing on z

Competitor's 4 Year Cost / Savings Prediction



IBM Analysis



Investment Maint MLC Operations Disaster Recovery

38

Lessons Learned About the Promises Made by the Competitors

- They **over-estimated** the mainframe costs
 - Over-provisioned too early $\Delta\$3.6M$
 - Used highest hardware purchase & maintenance list prices $\Delta\$9.4M$
 - Continued using older software; no sub-cap pricing $\Delta\$2.7M$

OVERESTIMATED BY $\Delta\$15.7M$
- They **under-estimated** the offload costs
 - Forgot about mainframe coexistence during migration $\Delta\$9.5M$
 - Forgot about high cost of power & cooling $\Delta\$1.1M$
 - Forgot about the financing charges $\Delta\$2.5M$
 - Added a test server $\Delta\$2.1M$
 - Under-provisioned batch processing (15 % growth case) $\Delta\$6.3M$
 - Failed to take into account technology updates $\Delta\$14.6M$
 - Did not provide Disaster Recovery $\Delta\$40.6M$

UNDERESTIMATED BY $\Delta\$76.7M$

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

Let's Consider The Other Elements of Cost

Total Cost of Ownership =
Cost of hardware +
Cost of software +
Environmentals +
Cost of labor +
Financial terms

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

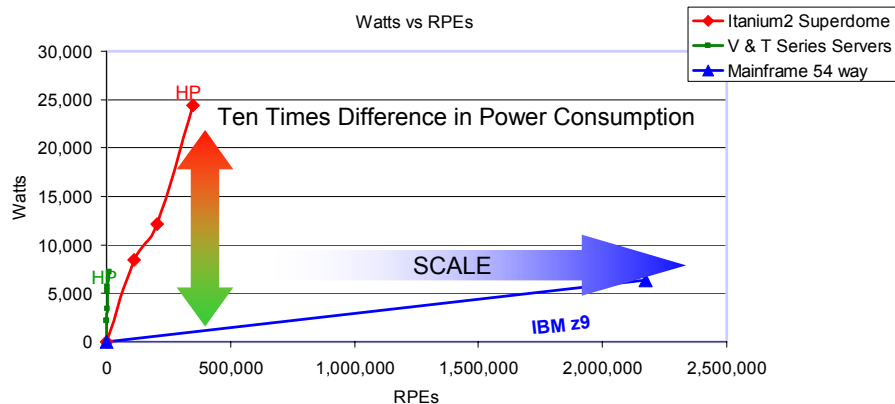
- More than half of all serious outages are now caused by power problems*
 - ▶ Room temperatures averaging 92°F lead to erratic machine behavior
 - ▶ *“Power-related problems in 2005 will cause 4 of the 20 major failures, up from 2 of 20 last year”* (The Uptime Institute)

*Source: recent AFCOM survey of 200

10 - Delivering Next Gen Solutions at the Lowest Cost v1.5.ppt

43

Mainframe Scale and Power Efficiency



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


10 - Delivering Next Gen Solutions at the Lowest Cost v1.5.ppt

44

Do the Math


- HP Itanium 2 Superdome 9050 (64ch/128co) consumes a maximum of 24,382 watts
 - ▶ $24.382 \times .08 \times 24 \times 365 = \$17,087$ per year for electricity
- Mainframe with similar computing capacity consumes 2,500 watts
 - ▶ **\$1,752** per year for electricity
 - ▶ Power cost is \$15,335 per year less
- Similar savings on cooling capacity
 - ▶ Cost of cooling is 60% to 80% the cost of power
 - ▶ Superdome total **\$27,339** per year vs Mainframe **\$2,803**

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



10 - Delivering Next Gen Solutions at the Lowest Cost v1.5.ppt

47

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



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

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

Source: Gartner

10 - Delivering Next Gen Solutions at the Lowest Cost v1.5.ppt

48

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	<u>\$793K</u>	<u>\$6,690K</u>

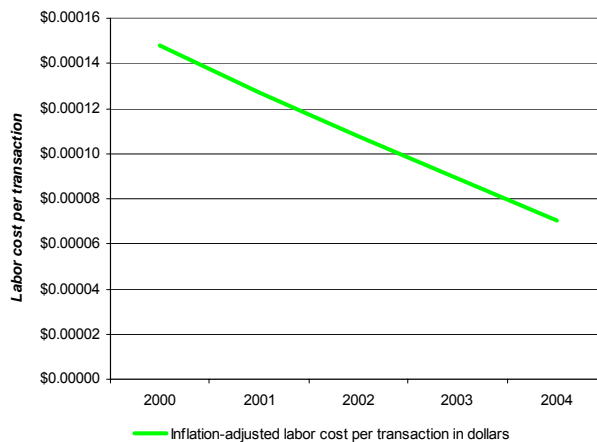
In this case, System z requires 1/8 the labor costs of the distributed environment

Source: IBM SWG Data Center

10 - Delivering Next Gen Solutions at the Lowest Cost v1.5.ppt

49

Labor Cost Per Transaction on System z is Decreasing



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

Source: IBM Global Services UK

10 - Delivering Next Gen Solutions at the Lowest Cost v1.5.ppt

50


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 customers typically receive credit for existing MIPS investment when upgrading to new generation
 - ▶ Full **trade-in value** applied to upgrade and growth 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 Example Considering 100% Trade in on System z Processors



Conclusion: Total Mainframe Hardware, Software & Labor Costs Reduced by 62% in 5 Years



10 - Delivering Next Gen Solutions at the Lowest Cost v1.5.ppt

53

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"

10 - Delivering Next Gen Solutions at the Lowest Cost v1.5.ppt

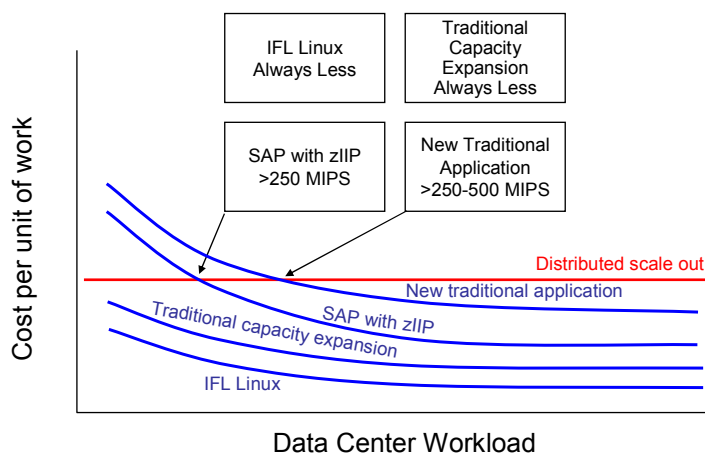
54

TCO – What Have We Seen Today?

- Incremental Data Warehouse workload on System z cost **less** than Oracle RAC on HP Superdome
- Incremental Data Server on System z cost **less** than Oracle HP Superdome
- Incremental WebSphere workload on System z costs **less** than distributed deployment
- Consolidation of Linux servers onto System z **saves** big money
- System z uses less power and requires fewer operational staff

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

