TCO Lessons Learned From Customer Engagements

The New zEnterprise – A Cost-Busting Platform

Eagle TCO Engagements

- **Free of Charge** total cost of ownership study that helps customers evaluate the lowest cost option among alternative approaches. The study usually requires one day for an on-site visit and is **specifically tailored to a customer's enterprise.**
- The study can be focused on at least one of the areas below :

- We conduct Eagle studies for System z, POWER, and PureSystems accounts
- Over 300 customer studies since the formation of the TCO Eagle team in 2007
- **Engage our Eagle-Eyed TCO Experts!**
  - Start by sending an email to: eagletco@us.ibm.com

TCO: Understand The Complete Picture

Competitors Tell Stories

Your competitor says I can save money by moving workloads off System z

CIO

Our competitor's claims are often false….We've shown several cases where System z was the lowest cost platform for core business workloads

IBM

Is There A Cross-Over Point?

Is a 500 MIPS workload small enough to offload from System z?

CIO

It depends on the ***nature*** of the workload, rather than the ***size***!

IBM

Core Proliferation For A Small Offload Project

2x 16-way Production / Dev / Test / Education
App, DB, Security, Print and Monitoring
4x 1-way Admin / Provisioning / Batch Scheduling

36 Unix processors
(222,292 Performance Units)

$17.9M (4 yr. TCO)

z890 2-way Production / Dev / Test / Education
App, DB, Security, Print, Admin & Monitoring

0.88 processors
(332 MIPS)

$4.9M (4 yr. TCO)

41x more cores

Almost 5 Year Migration

670 Performance Units per MIPS

Plus:
2x HP SAN Servers (existing)
Many (existing) Windows servers

No Disaster Recovery

Core Proliferation For A Smaller Offload Project

z890 Production / Test

4x p550 (1ch/2co)
Application and DB

0.24 processors
(88 MIPS)

8 Unix processors
(43,884 Performance Units)

$8.1M (5 yr. TCO)

33x more cores

$4.7M (5 yr. TCO)

3 Year Migration

499 Performance Units per MIPS

Status Of An In-Flight Example

- Financial customer publically claims huge offload and savings with 2 year ROI
  - Fairly typical when it comes to claims we have seen
  - Under the headlines it becomes clear this is a projection of the results for a multi-stage project which has not yet started

- A couple of years later IBM examined the situation at the customer again
  - The total MIPS on the mainframe was exactly the same
    - 20% of the MIPS now zIIP rather than GP MIPS so some software savings
    - Customer now has many tens of Intel servers / cores / software licenses AS WELL
  - The customer has spent $34 million so far (2/3 through the multi-stage project)
  - The total number of operational FTEs appears to have increased by 4 so far
  - The project is overrunning by about 18 months so far
  - Migrating DB2 to the competing database is proving difficult
  - The original project sponsor has left their job very suddenly
  - Publicized savings are from changing outsourcer to a non-profit organization
    - Most customers don’t have this option (large outsourcing savings)

So, What About “Failed” Rehosts?

- Not surprisingly, nobody wants to talk about them
- Sometimes you can infer failures
  - Search the web for announcements of huge mainframe rehosts (majority of these are statements of intent)
  - Now follow up each one to determine what they achieved
  - Admittedly, it’s easier to do this if you are IBM and still receiving mainframe revenues from the customers!
- Based on our own personal experience, we can describe a couple of representative real examples…

2005 Offload Failure Example

Lombard Canada Ltd. partnered with Micro Focus to replace their old mainframe

- 200 MIPS
- CICS, COBOL, VSAM, DB2

"We estimate this project will save us in excess of $1 million a year, but more importantly, it will enable us to become more competitive in our industry both today and in the future." *

VP of IT
Lombard Canada Ltd.

BUT one year after starting, the project was abandoned

- System Integrator and Micro Focus did not have the skills
- Lombard spent millions on conversion with no results
- VP lost his position

Today, Lombard continues as a System z customer, moving to z114

*Source: http://www.finextra.com/news/Announcement.aspx?pressreleaseid=4858

2011 Offload Failure Example

- Customer objective: Offload 3,500 MIPS with Micro Focus $10M budget and 1 year schedule
- 18 months later they had spent $25M and moved only a 10% of their MIPS
- Additional costs came from
  - Internal staff to cover the overrun
  - Substantial manual steps replaced mainframe automation
  - Needed many additional software products
  - Ended up acquiring additional distributed capacity over initial prediction (just to support the 10% they actually offloaded)
  - Extending the dual-running period of the rehost
- Executive sponsor lost their job

Projects Delayed And Expensive

- Many large projects overrun, and rehosting projects tend to be larger than anticipated
- Most real cases start with 1 year projections which turn into 2 or 3 years
- They also show terrible TCO cases, even where they can achieve annual savings after migration (sometimes they can't even do that)

- A couple of typical examples…

	European Government	*US County Government*
Project Completion	+1 year delay	+1 year and 8 months delay
Migration Cost (inc. hardware and software)	$19.6M	$6M
Payback Period	>29 years	>20 years

Question: Are Rehosting Vendor Estimates Based On Real Migration Experiences?

- Typical rehosting vendor estimates are ridiculously low

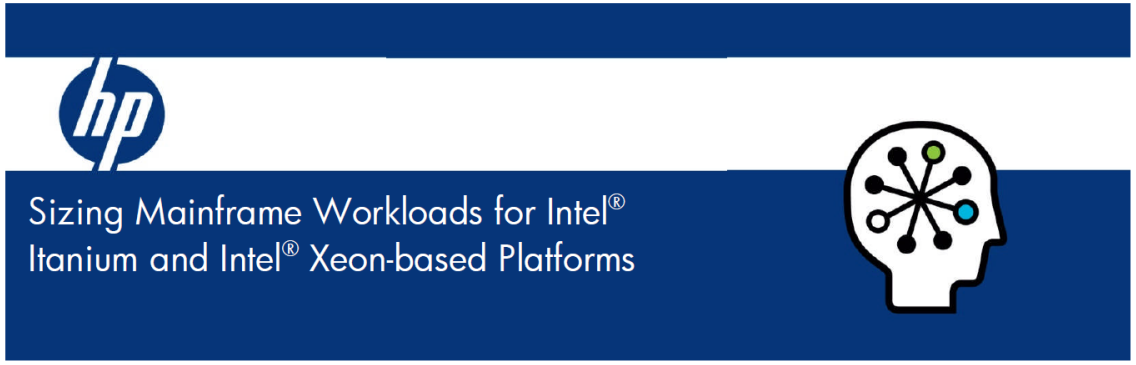

Core proliferation of less than 3x understates observed migration results

System	CPU	CPU cores	zOS MIPS	zOS MIPS/core CPU Sockets
X86-based Blade server	3.0 GHz Quad core x86	8	1858	**232**
HP rx8640	1.6 GHz Intel® Itanium® Dual-Core Processor 9100 Series	16	4046	**253**

- Compare to our most efficient real example (29x): Instead of 176 cores, they should have been able to rehost with just 7 cores according to HP!

Lessons Learned Can Be Grouped Into Three Broad Categories

- Always compare to an optimum System z environment
- Look for not-so-obvious distributed platform costs to avoid
- Consider additional platform differences that affect cost

All examples discussed are from actual Eagle Team customer studies

(1) Always Compare To An Optimum System z Environment

- Updating hardware and software reduces cost
- Specialty processors reduce mainframe cost
- Sub-capacity may produce free workloads
- Replace ISV software with IBM software
- System z Linux consolidation saves money
- Changing database can impact capacity requirements
- Don’t forget Solution Edition!

Currency Reduces Cost – Hardware

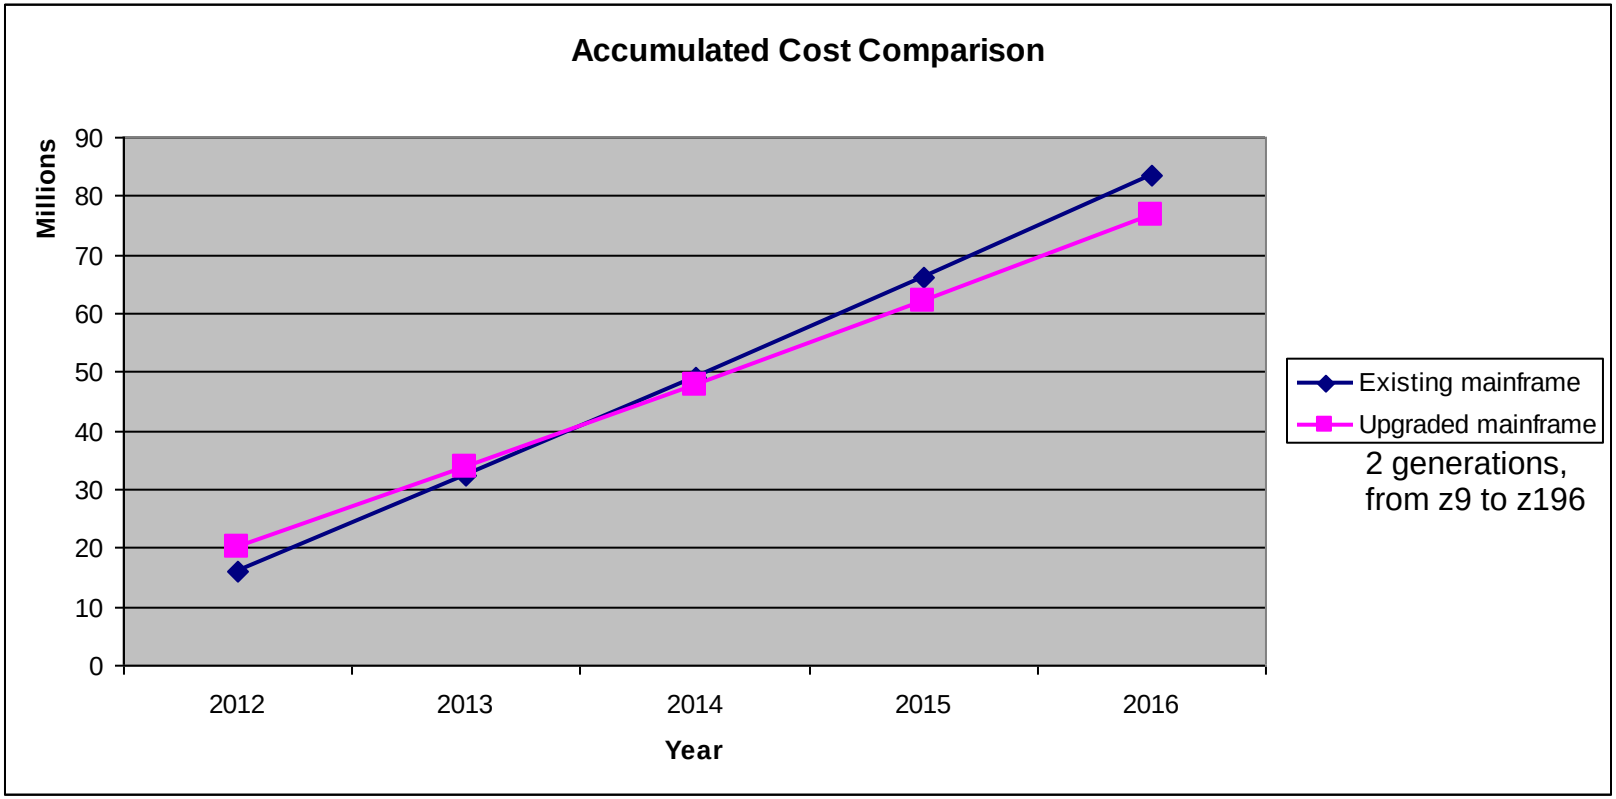

- Typical customer hardware refresh scenario
  - 2M investment pays back >1M savings every year – most cases positive in a 3 year period
  - Savings from technology dividends and specialty processor offload
- Comparing latest technology servers to old mainframes is unfair but often done

Performance Improvements Means MLC Costs Go Down And Frees Up Hardware Capacity

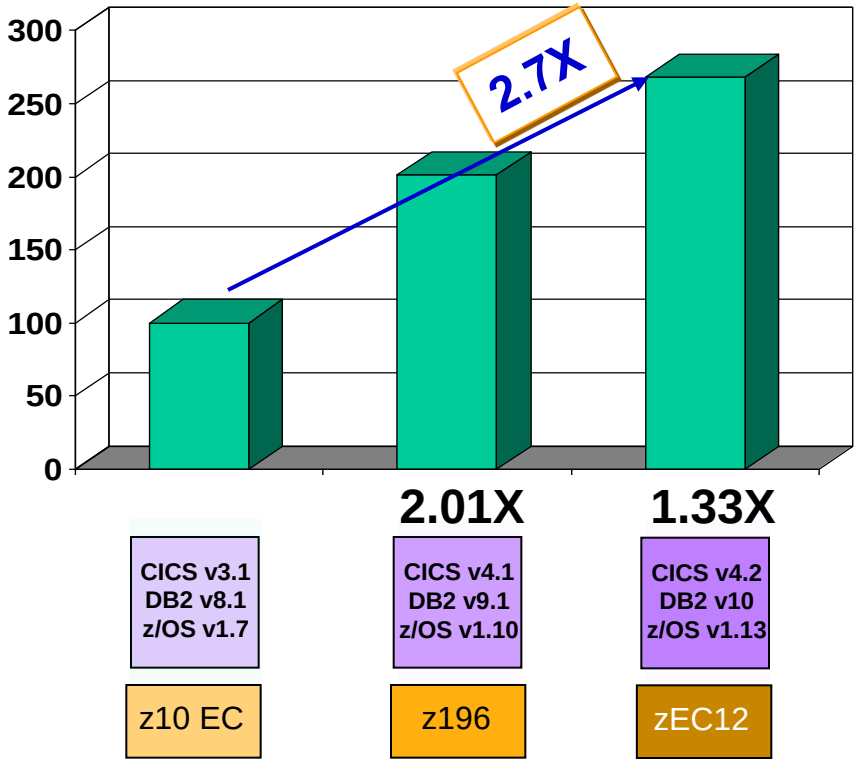

Customer examples:

(1) Large MEA bank

- Delayed upgrade from z/OS 1.6 because of cost concerns
- When finally did upgrade to z/OS 1.8
  - Reduced each LPAR's MIPS by 5%
  - Monthly software cost savings paid for the upgrade almost immediately

(2) BMW Autos

- Upgraded to DB2 10
- Realized 38% pathlength reduction for their heavy insert workload
  - Other DB2 10 users saw 5-10% CPU reduction for traditional workloads

Additionally, save costs by moving to newer compilers and tuning

IBM internal core banking workload (Friendly Bank). Results may vary.

US Financial Company Doubles Performance After Upgrading To DB2 10 And Tuning

Tests showed **2.15x** boost in performance of business intelligence application

- First computed 42 operational BI reports serially
- Then database software upgraded to DB2 10
  - Performed tuning such as computing additional indexes, collecting additional statistics and pre-computing global Temp tables
- Results showed **54%** reduction in response time

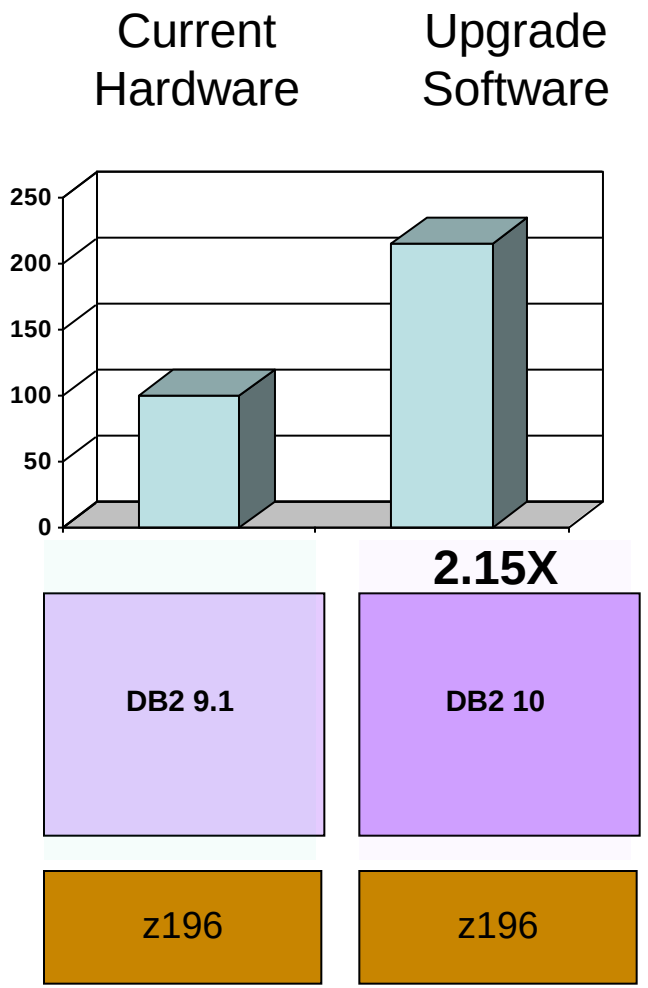

Results may vary

zEC12 Continues A History Of Performance Improvements For Java

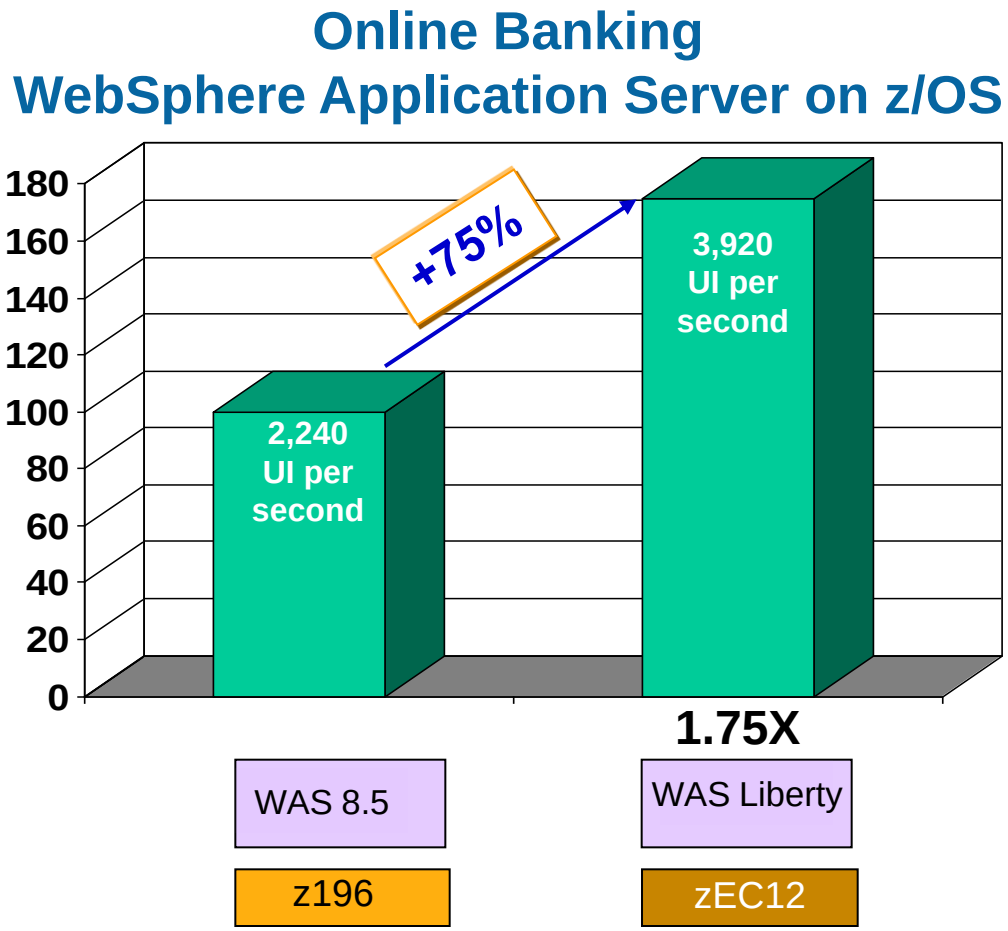

IBM internal Java workload

UI per second = User interactions per second

Specialty Engines Continue To Revolutionize Mainframe Costs

- Specialty processors – zAAP, zIIP, IFL
- Very attractive per processor cost (OTC)
  - Considerably less cost than a general purpose z/OS processor
  - No charge for IBM software running on zAAP/zIIP
  - IBM software on an IFL costs 120 PVU's (less than an Intel dual core)
- Customers typically offloading between 40%-95% of workload

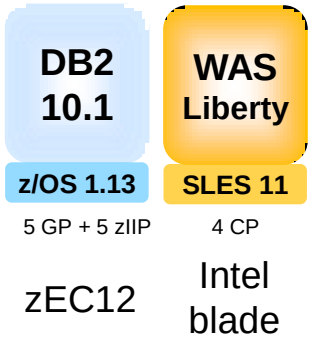

$102 / UI per sec

New WAS middleware now provides as much as 97% zAAP offload!

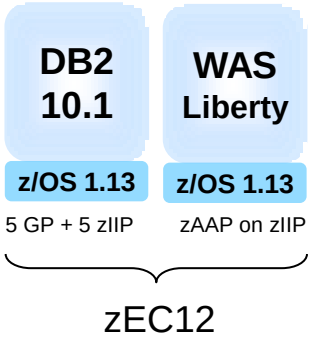

$46 / UI per sec

IBM internal banking workload.

Sub-Capacity May Produce Free Workloads

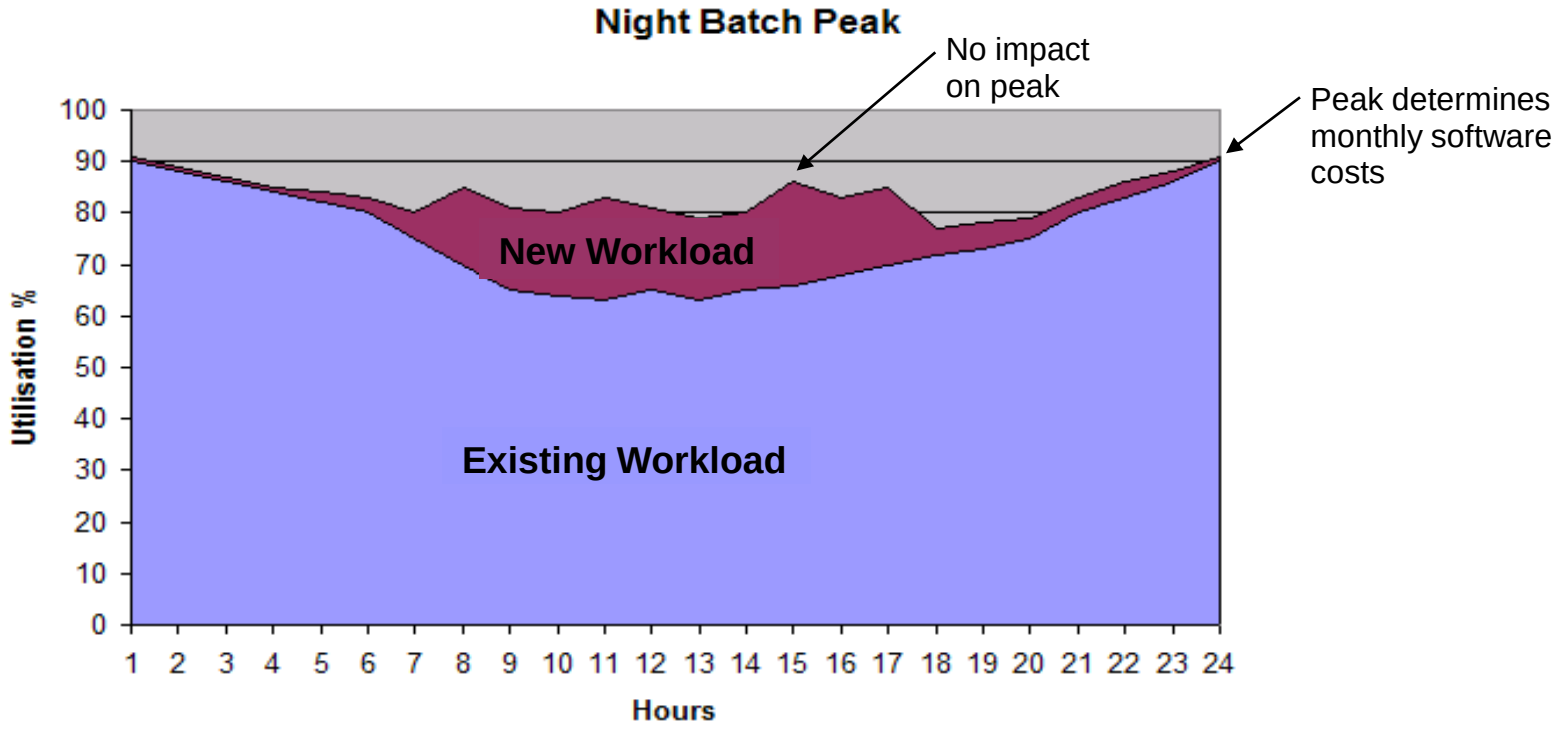

- Standard “overnight batch peak” profile – drives monthly software costs
- Hardware and software are free for new workloads using the same middleware (e.g. DB2, CICS, IMS, WAS, etc.)
- Ensure you exploit any free workload opportunities, and conversely, avoid offloading free applications!

Replace ISV Software With IBM Software

Medium-sized European financial company:

Typical mixed software profile (ISV and IBM)

Average Profile (BEFORE)		
Weighted MIPS		8,800
Cost Per MIPS per Year		Profile
IBM Software	1,000.00	24.72%
		0.00%
ISV Software	1,540.00	38.07%
TOTAL SW	2,540.00	

New profile after migration to IBM

Actuals (AFTER)		
Weighted MIPS		8,900
Cost Per MIPS per Year		Profile
IBM OTC	376.09	13.66%
IBM MLC	1,023.77	37.20%
ISV Software	136.09	4.94%
TOTAL SW	1,535.95	

- IBM software costs now slightly higher, but ISV software costs dramatically lower
- Result: $1000/MIPS per year savings

Replacing ISV Software With IBM Is Also A Better Deal Than Offloading

	Mainframe Offload	*Move to IBM Tooling*
Investment $, Period	$54M over 2 years	$3M over 1 year
Predicted Annual Cost Savings	$13M from year 3	$6M from year 2
5 Year TCO, Breakeven	$140M, year 7	$101M, year 2
Level of Risk	Very High	Very Low

- Finance perspective
  - Large, risky and expensive project with distant payback under unlikely assumptions, versus
  - A small, low risk and cheap project with instant payback

System z Linux Consolidation Saves Money

- Large financial services company
  - Mixed Oracle and WAS environments on Intel
  - Scaling out rapidly – up to 172 images on 836 cores

	Refresh existing x86	*Consolidate on zLinux*	*Observation*
Software	9.76M	3.70M	Software costs down 62%
Hardware	2.32M	5.22M	Hardware costs up 125%
Labor/Facilities	3.83M	0.69M	Charge to department down 82%
Migration	0.18M	0.41M	
Cost Avoidance	-	-0.61M	
Total	**16.1M**	**9.41M**	

- 5 Year Savings: 6.6M
- 26:1 core consolidation from virtualized x86 to zLinux

Changing Database Can Have Dramatic Capacity Impacts

- IMS DB is a very efficient hierarchical non SQL data store
- Attempts to replace with an RDBMS are suboptimal
  - 2X – 3X MIPS used by DB2 vs. IMS DB
  - Degraded response time

- European customer is trying this, while still running the business
  - New data still being entered daily into IMS, then batch-replicated to DB2, then more ETL to islands of Oracle databases for other applications
- No new function, and unable to keep up with competition
- Only 30% converted in 4 years, 500M€ spent so far

New Mainframe Workloads With Unbeatable Price Points Via Solution Edition

- Bundle of System z **hardware, software** and **maintenance**
  - 3 or 5 year **Best Price**
- Focus: new System z workload opportunities
  - Not for existing workloads
- Solution Editions usually include:
  - System z hardware (new footprint or incremental)
  - Prepaid hardware maintenance
  - Comprehensive middleware stack (including subscription and support)
  - Storage as an option for all Solution Editions
  - Services for some Solution Editions

(2) Look For Not-so-obvious Distributed Platform Costs To Avoid

- Distributed servers refresh every 3 to 5 years
- Distributed server disaster recovery is typically at 100%
- Customers often overlook significant tools replacement costs

Distributed Servers Need To Be Replaced Every 3 To 5 Years

- IT equipment refreshed 2 – 7 year intervals, normally 3 or 4 years
- Distributed servers re-purchased each time
  - Normally with some additional growth capacity (CPU, memory, I/O and other specialty cards like cryptographic offloads)
- With a growing mainframe, customers normally only have to purchase the additional (new) MIPS capacity
  - Existing MIPS are often carried over to the new hardware
  - Existing memory, I/O facilities and specialty processors / cards are also normally carried over to the new hardware
- Five year studies show this effect, short time periods do not

Distributed Servers Need To Be Replaced Every 3 To 5 Years (2)

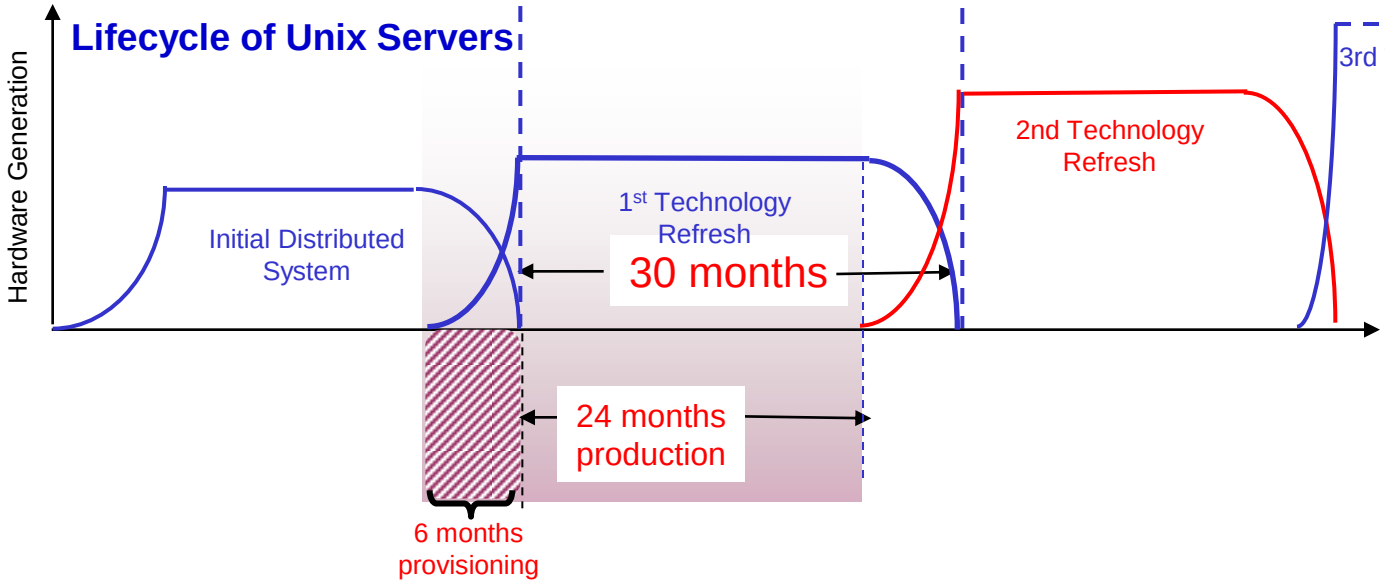

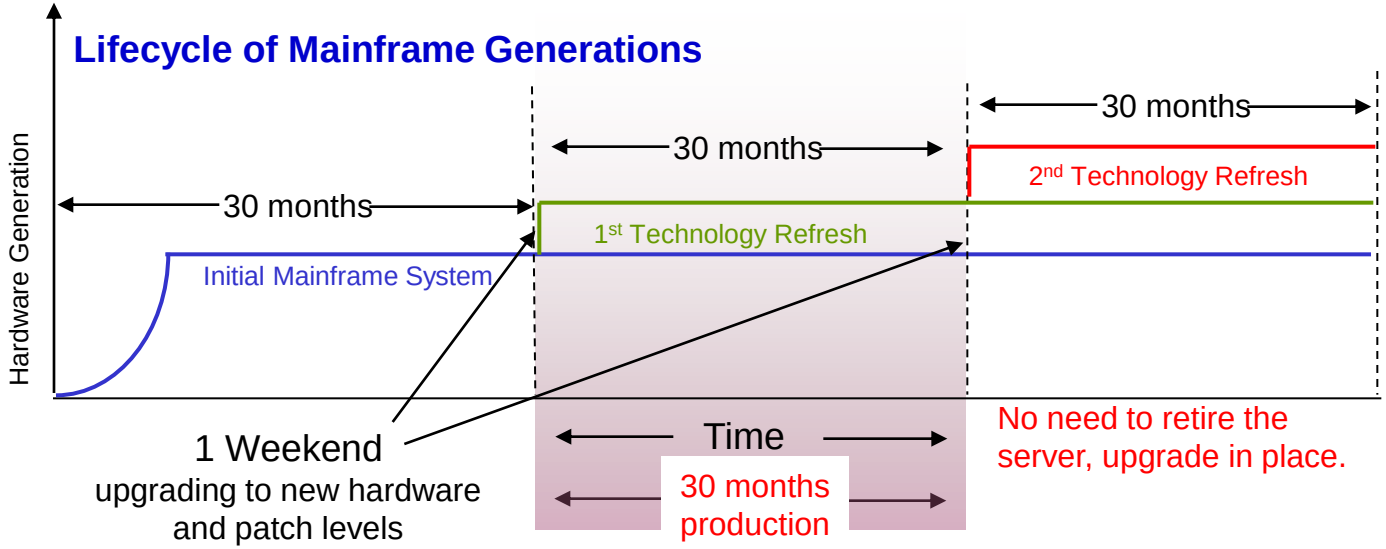

Refresh is normally even worse than just re-purchasing existing capacity as this real customer demonstrates:

Non-mainframe systems must co-exist for months at a time while being refreshed, requiring space, power, licenses etc. In this case only 24 months of productive work is realized for each 30 month lease period and the leases overlap up to 6 months

The mainframe by contrast is upgraded over a weekend and is fully productive at all times

Disaster Recovery On System z Costs Much Less Than On Distributed Servers

Mixed distributed and System z environment at a large European insurance company:

Disaster Recovery Cost as a percentage of Total Direct Costs:

System z – **3%**

Distributed – **21%**

Two mission-critical workloads on distributed servers had DR cost > 40% of total costs

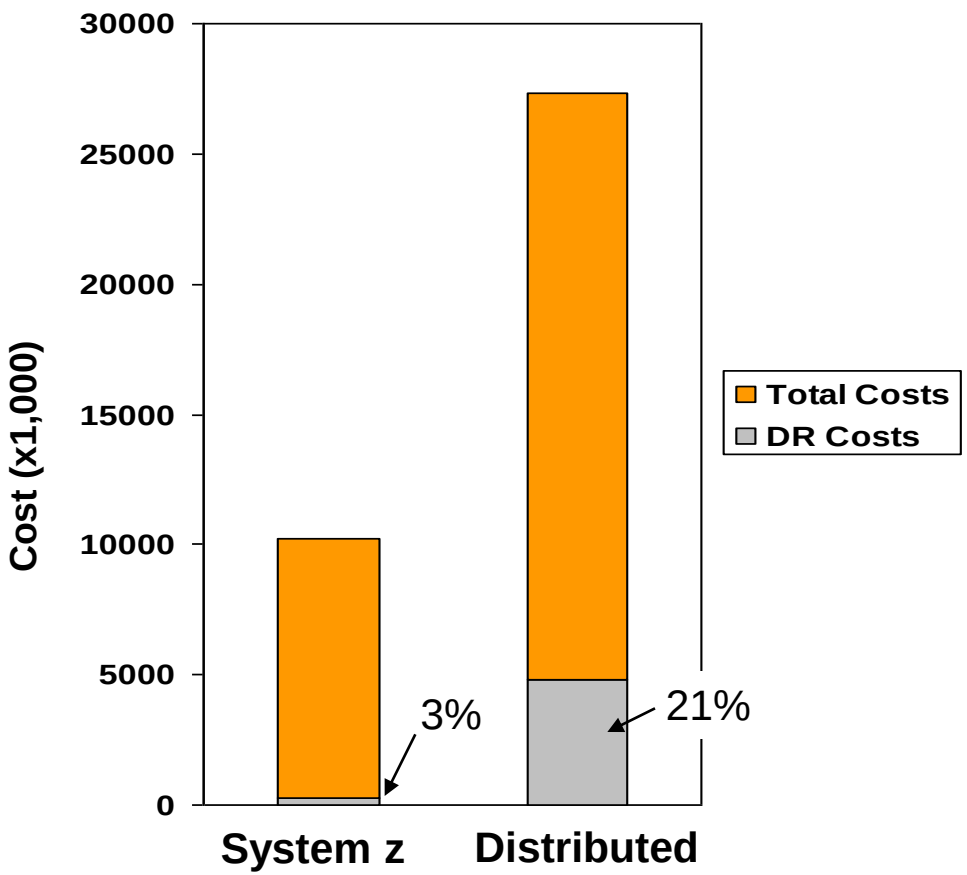

Disaster Recovery Testing Is Typically More Expensive On Distributed Platforms Too

- A major US hotel chain
  - ~ 200 Distributed Servers (LinTel, Wintel, AIX, and HP-UX)

	Person-hours	*Elapsed days*	*Labor Cost*
Infrastructure Test (7 times)	1,144	7	$89,539
Full Test (4 times)	2,880	13	$225,416
Annual Total – Distributed	14,952*	73	$1,170,281
Mainframe Estimate	2,051*	10	$160,000

* Does not include DR planning and post-test debriefing

- Customer Recovery Time Objective (RTO) estimates:
  - Distributed ~ 48 hours to 60 hours
  - Mainframe ~ 20 minutes

- Conclusion: Mainframe both simplifies and improves DR testing

Customers Often Overlook Significant Tool Replacement Cost

- Customers often struggle to identify all the replacement tools and middleware they will need for an offload
- Straight-line extrapolation of cost from the easily identified subset is often accurate enough

- Customer example: 261 total software products on z/OS
  - 37 product replacements identified in vendor proposal and IBM identified an additional 16 for a total of 53 products of 261 (20%)
- 208 products missing – how to estimate their likely cost, especially given that not all products will end up with one-for-one replacements:
  - Applications may be re-written to not need missing products
  - New code could be written to perform the function from scratch
  - Adding operations labor to manually do the function could be an option

- We extrapolated from the known products cost and a few years later were proven to be very close to the mark

Not All Mainframe Management Software Could Be Replaced

Total **Distributed** Software Costs

$53.8M (5 yrs.)

Distributed Software Identified	Initial OTC	Maint. (per yr)
DB2	$4.50M	$0.99M
DB2 Recovery Expert	$1.58M	$0.35M
DB2 Optim Perf. Manager	$1.31M	$0.29M
DB2 Adv. Access Control	$1.23M	$0.27M
DB2 PureScale	$2.18M	$0.48M
IBM Optim DB Admin.	$0.66M	$0.15M
MQ	$0.82M	$0.18M
System Automation	$3.56M	$0.78M
Workload Scheduler	$0.78M	$0.17M
Access Manager	$0.51M	$0.11M
Micro Focus	$8.89M	$1.60M
Micro Focus Studio Ed.	$0.84M	$0.11M
Additional Products	$2.61M	$0.57M

Total **System z** Software Costs

$30.0M (5 yrs.)

44% less

- Only 12 of 26 mainframe system management products available on distributed platform
- Of those, functionality not equivalent
  - Operations automation not as robust
  - Tape solution missing
  - Database tools missing
  - No RTM1 and RTM2
  - Lack of SMF and RMF
- Development costs for repair of missing functionality not included

(3) Consider Additional Platform Differences That Affect Cost

- System z responds flexibly to unforeseen business events
- Disaster Recovery is better and cheaper on the mainframe
- Cost of adding incremental workloads to System z is less than linear
- Offloading chatty applications introduces latency
- Non-production environments require fewer resources on System z
- Batch challenges non-mainframes
- Cost of administrative labor is lower on System z
- System z cost per unit of work is much lower than distributed

Mainframe Can Respond Flexibly To Unforeseen Business Events

- Transportation company experienced a natural disaster
  - Required them to re-run a whole weeks business while continuing to operate normally
  - Able to turn on double capacity immediately to achieve this
- Customer informed IT department that they would be running a Super Bowl advertisement with very short notice
  - Massive capacity spike
  - Temporarily turned on additional capacity
  - Stress tested their systems prior to the event despite short notice

Cost Of Adding Incremental Workloads To System z Is Less Than Linear

- Mainframes are priced to deliver a substantial economy of scale as they grow
- Doubling of capacity results in as little as a 30% cost growth for software on z/OS
- Average Cost is significantly more than incremental cost

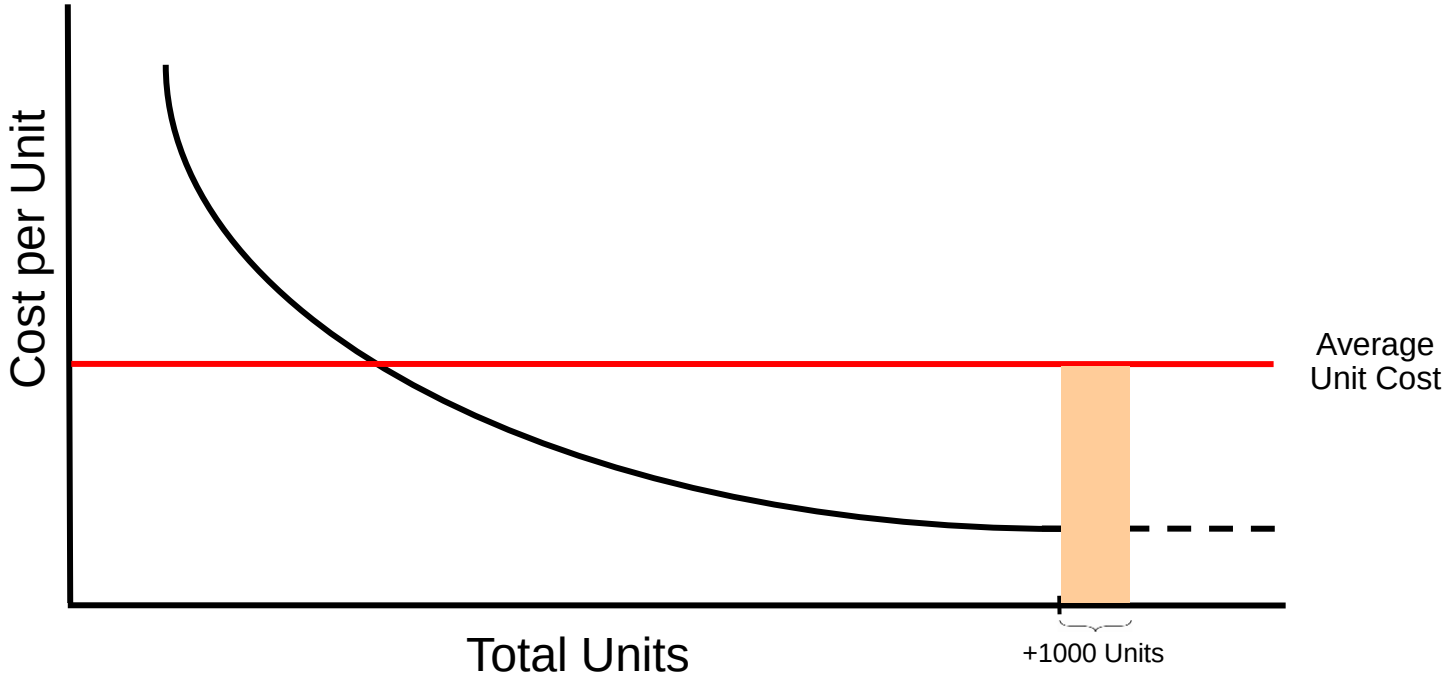

Cost Of Adding Incremental Workloads To System z Is Less Than Linear (Example)

- Customer determined that their current deployment of WAS applications would cost a similar amount each year on either the mainframe, or on a distributed platform
- Then they examined the incremental cost of adding one more large WAS application to each platform
- The mainframe demonstrated a clear advantage

	Incremental Mainframe	Incremental Distributed
5 Year TCO	**1.29M** **(657K OTC, 42K Y1, 147K Y2-5)**	**1.56M** **(378K OTC, 192K Y1, 249K Y2-5)**

- Although moving existing WAS applications between platforms is unlikely, future WAS deployments will therefore be targeted to the mainframe

Some Applications Originally Designed With Co-located Data Are Not Good Offload Candidates

- Large insurance company rehosted portion of application as POC
  - Found TCP/IP stack consumed considerable CPU resource, and introduced security compromises and network latency
- European bank tried rehosting CICS workload to Linux while maintaining VSAM and DB2 data on System z
  - Induced latency meant CICS applications no longer met SLAs

Single z/OS LPAR

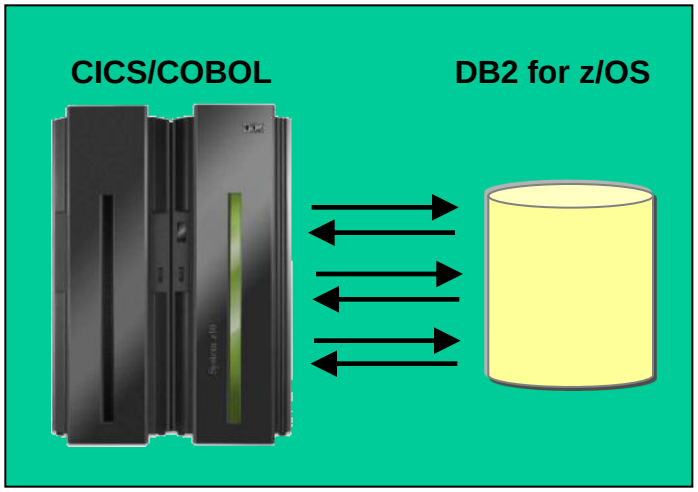

Distributed architecture

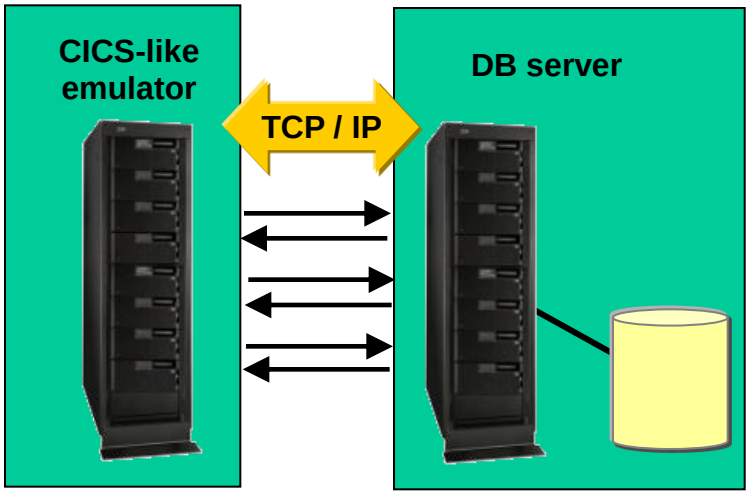

Co-locating In the Same Address Space Is More Efficient

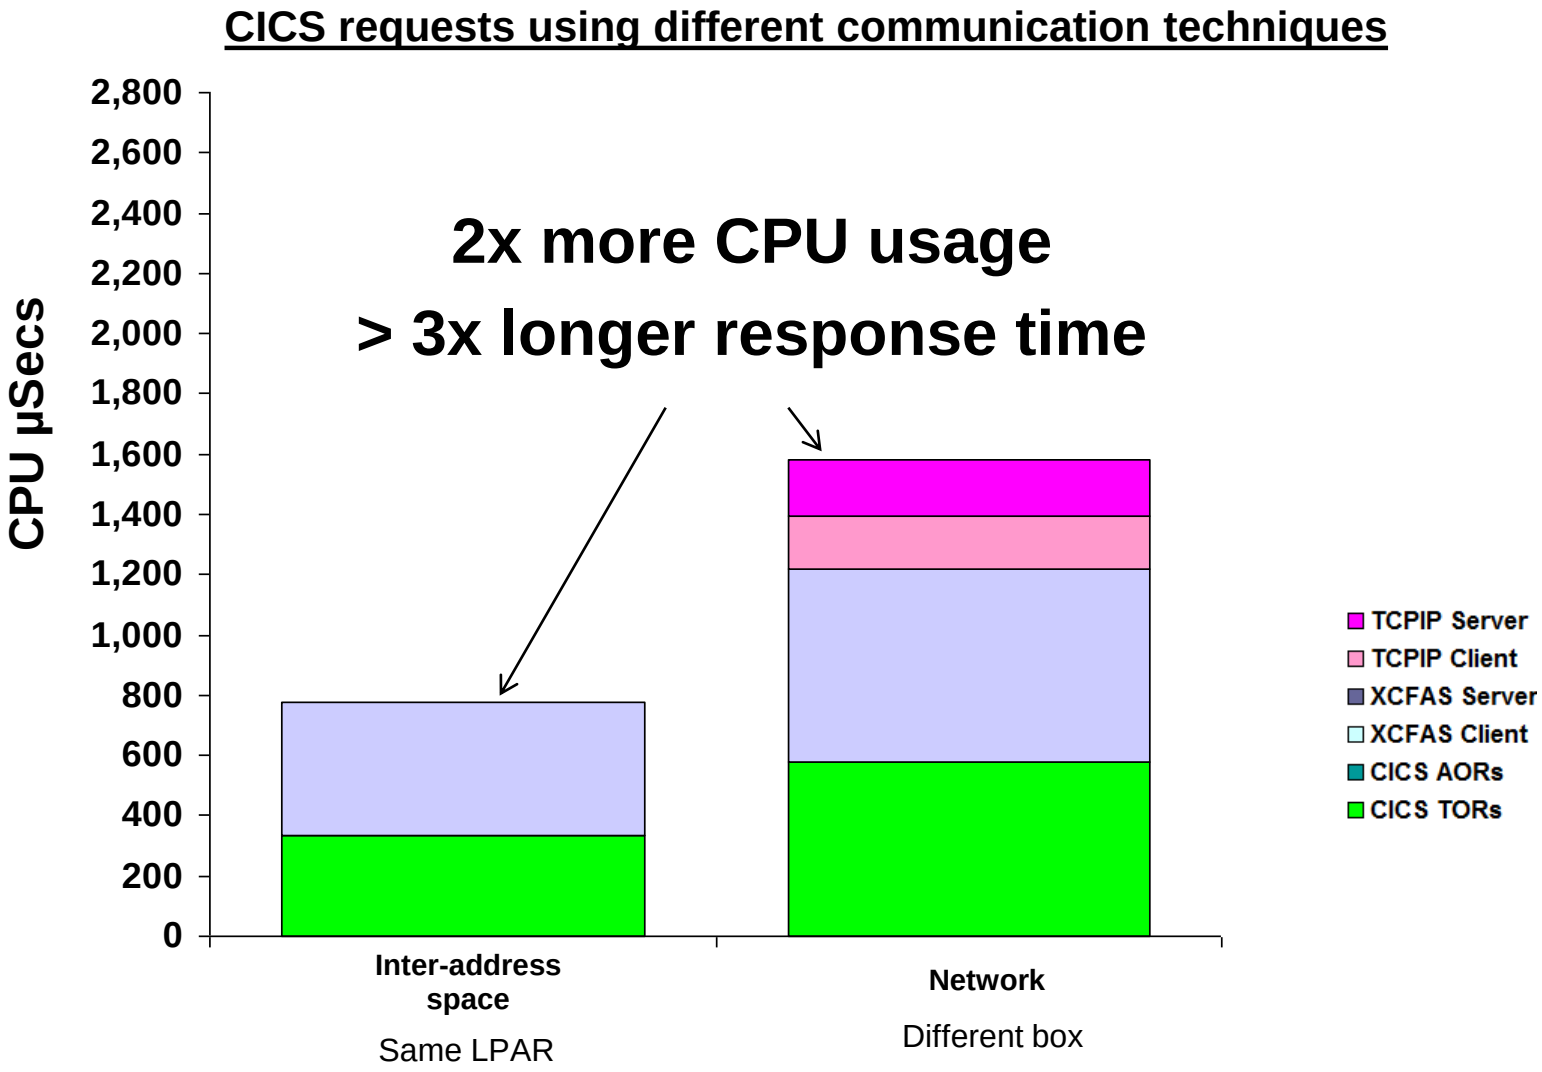

Source: http://hurgsa.ibm.com/projects/t/tp_performance/public_html/OS390CICS/reports/CICS%20TS%20V4.2%20Performance.ppt and email with z/OS Communications Server development team

Non-Production Environments Require Fewer Resources On The Mainframe

- High Availability mechanisms for Production
  - Dedicated failover (Prod x 2.5)
  - N+1 clustered (Prod x 2 worst case)
  - Mainframe (usually Prod x 1, sometimes less!)

- Development and Test Capacity
  - Mainframe – Prod +20%
  - Distributed – a range, often Prod +200%

Production *QA* *Dev/Test/DR*

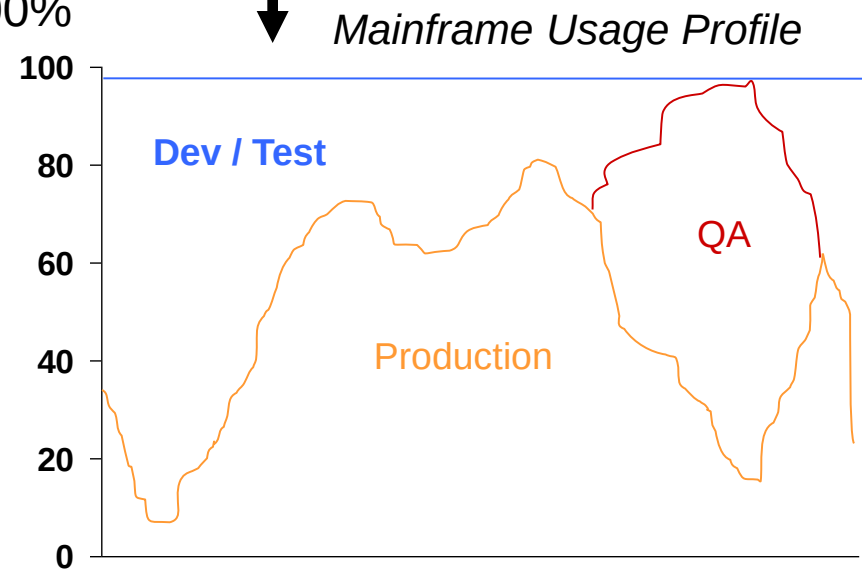

Batch Challenges Non-Mainframe Systems

- Rehosting vendor committed to a quick partial migration to avoid mainframe growth
  - Customer was an outsourced account with linear pricing (incents shrinkage of MIPS)
  - MIPS increase would drive a multi-million dollar one-time charge in 3rd party software
  - Growth would also incur another $1M per year in recurring outsourcing charges
- Current mainframe configuration very simple

Current Mainframe

- CPU = 1 unit
- Elapsed time = 1 unit

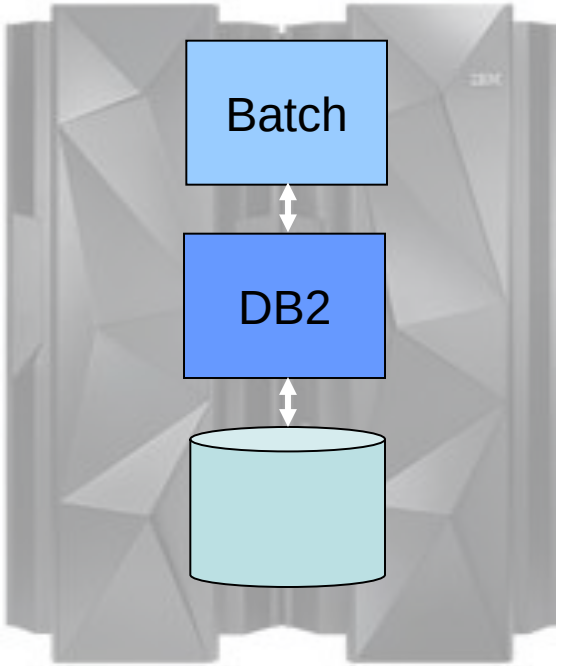

System z

Batch Challenges Non-Mainframe Systems

- Additional network latency dramatically increased elapsed job time (between 10x and 25x)
- Additional DRDA processing doubled mainframe CPU usage even though the application was now running on Intel

Resulting Mainframe

- CPU = 2 units
- Elapsed time = 10-25 units

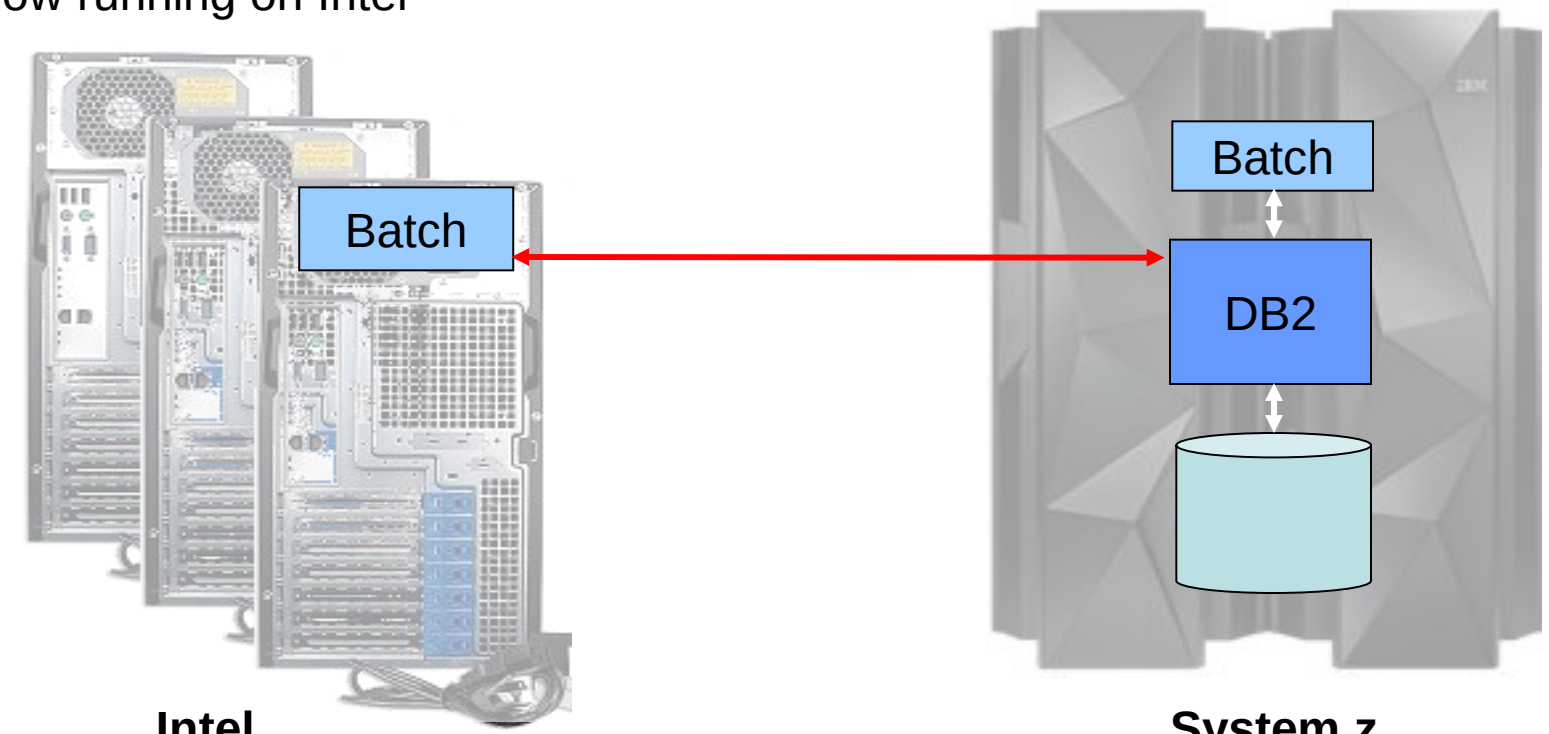

Large Systems With Centralized Management Deliver Better Labor Productivity

Large US Insurance Company

HP Servers + ISV

Production Servers
HP 9000 Superdome RP4440
HP Integrity RX6600

Dev/Test Servers
HP 9000 Superdome RP5470
HP Integrity RX6600

Claims per year **327,652**

$0.12 per claim

$0.79 per claim

Mainframe support staff has 6.6x better productivity

IBM System z CICS/DB2

Total MIPS **11,302**

MIPS used for commercial claims processing prod/dev/test **2,418**

Claims per year **4,056,000**

Mainframe Cost Per Unit Of Work Much Lower Than Distributed

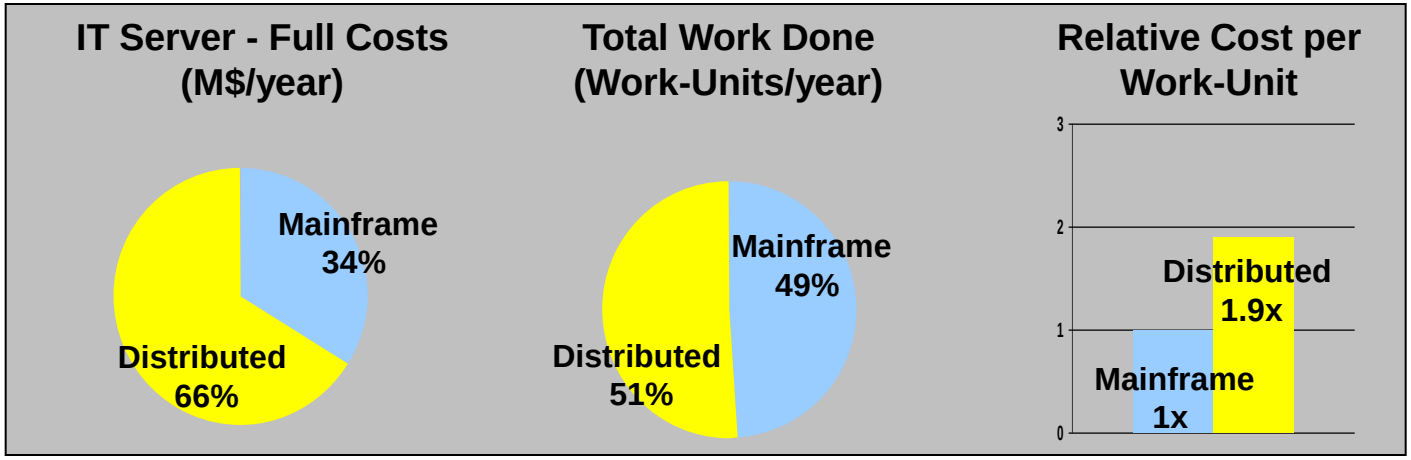

Case 1:
Very large retail bank
total of service delivery plus application development

Mainly UNIX distributed (>5000 servers)

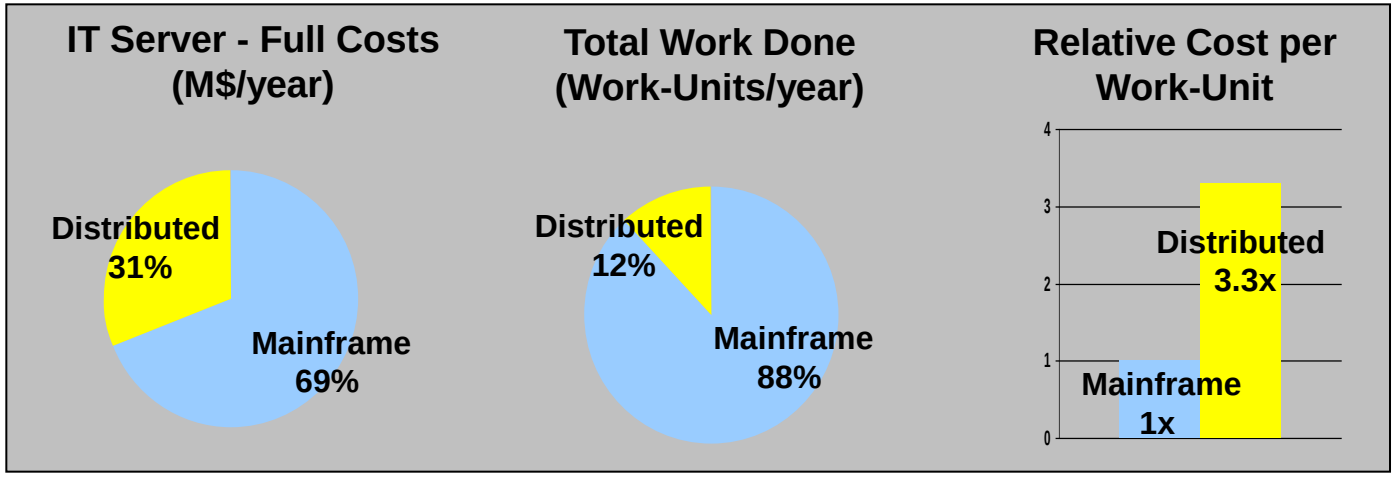

Case 2
Medium retail bank
service delivery only
Mainly mainframe
2500 mips, 13 M txns/day
Limited distributed servers
Windows + some UNIX
(~350 servers, 12% util'n)

Data from 3Q06 Scorpion studies

The distributed IT Total Cost/Work-Unit is approx. 2-3x Mainframe Cost/Work-Unit
The Mainframe typically does more work, Distributed has a lot of supporting infrastructure

Numerous TCO Studies Prove These Learned Lessons

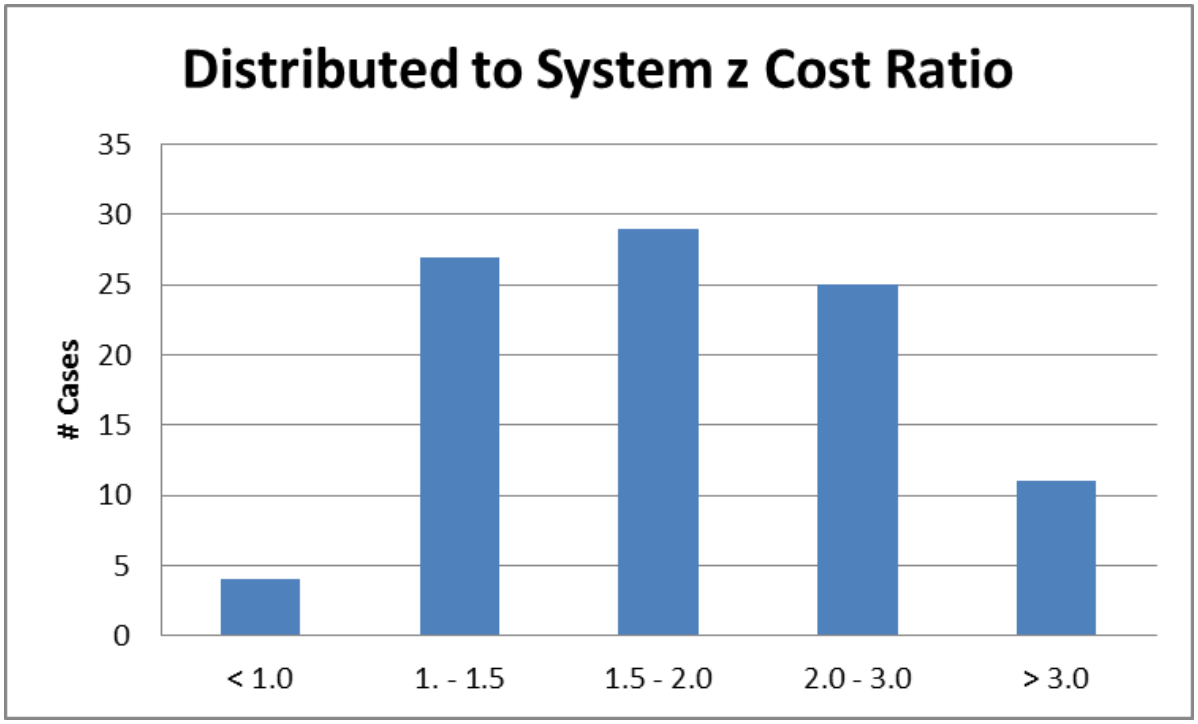

- 97 total customer studies
- Average cost of distributed alternative is 2.2 times greater than System z
- Only 4 out of 97 studies showed lower costs on distributed

Notice Regarding Specialty Engines (e.g., zIIPs, zAAPs And IFLs):

Any information contained in this document regarding Specialty Engines ("SEs") and SE eligible workloads provides only general descriptions of the types and portions of workloads that are eligible for execution on Specialty Engines (e.g., zIIPs, zAAPs, and IFLs). IBM authorizes customers to use IBM SE only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the “Authorized Use Table for IBM Machines” provided at www.ibm.com/systems/support/machine_warranties/machine_code/aut.html (“AUT”).

No other workload processing is authorized for execution on an SE.

IBM offers SEs at a lower price than General Processors/Central Processors because customers are authorized to use SEs only to process certain types and/or amounts of workloads as specified by IBM in the AUT.