Mainframe Rehosting: The Ugly Truth about Costs and Quality of Service!

Innes Read, IBM

Businesses trust their most mission-critical applications and data to the mainframe

70%
of top 500 System z customers run CICS

67%
of top 500 System z customers run CICS and DB2

21 of top **25**
insurance organizations use System z

23 of top **25**
retailers use System z

25 of top **25**
world's banks use System z

Today

IBM zEC12

1964

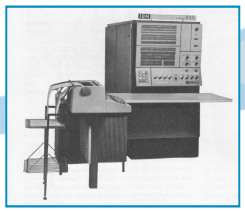

IBM S/360

Source: IBM

Yet, some mainframe clients are tempted to move workloads off the mainframe, allegedly to save money

z900
and
z/OS v1

… our hardware and software are old…

"…we're only running 87 MIPS…"

… our mainframe is pretty small…

$$$$

… accounting is telling me the mainframe is expensive…

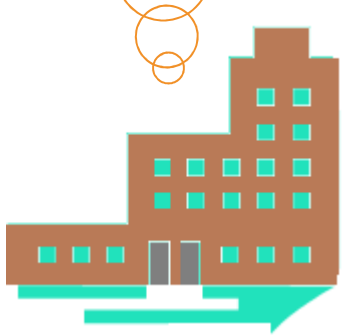

The IBM Eagle team can help customers understand mainframe costs and value

- **Worldwide** team of senior technical IT staff
- **Free of Charge** Total Cost of Ownership (TCO) studies
  - Help customers evaluate the lowest cost option among alternative approaches
  - Includes a one day on-site visit and is **specifically tailored to a customer's enterprise**
- Over 300 customer studies since formation in 2007
- Contact: eagletco@us.ibm.com

Eagle team data shows that in 96% of mainframe rehosting cases, clients ultimately end up spending *more* for an offload

In 38% of cases, the distributed platform was 2 or more times the cost of the System z platform

In only 4% of Eagle TCO studies was the cost of the distributed platform cheaper than the cost of the System z platform

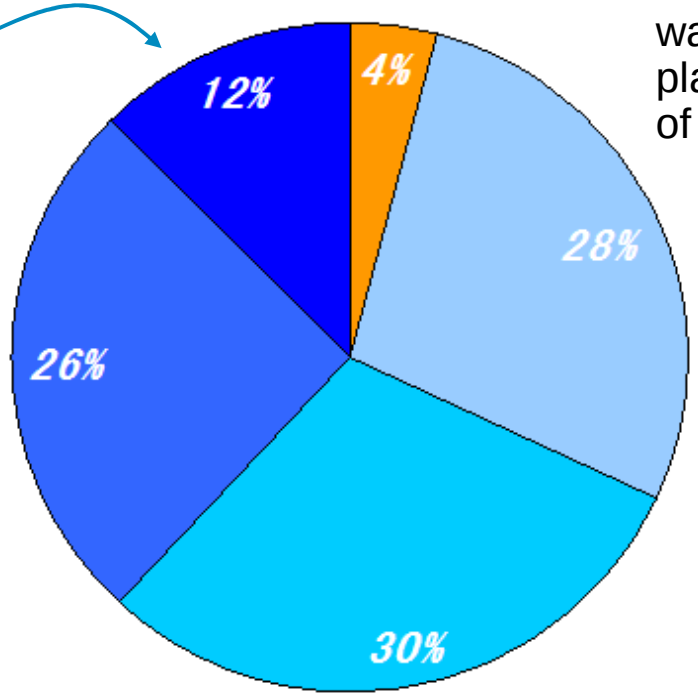

Cost ratio – Distributed to System z

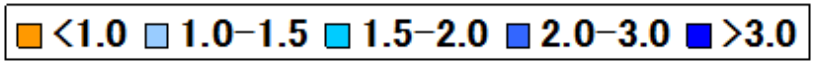

Sampling of 97 Eagle team TCO studies from 2007 - 2011

Example: Moving transaction processing off System z rarely reduces cost

Eagle TCO study for a financial services customer:

4 HP Proliant DL 980 G7 servers

Production

Development

256 cores total

Hardware	$1.6M
Software	$80.6M
Labor (additional)	$8.3M
Power and cooling	$0.04M
Space	$0.08M
Disaster Recovery	$4.2M
Migration Labor	$24M
Parallel Mainframe costs	$31.5M
Total (5yr TCO)	**$150M**

System z z/OS Sysplex

2,800 MIPS

Hardware	$1.4M
Software	$49.7M
Labor	Baseline
Power and cooling	$0.03M
Space	$0.08M
Disaster recovery	$1.3M
Total (5yr TCO)	**$52M**

65% *less cost!*

Rehosting costs are underestimated because of unrealistic equivalence data

From HP's "Mainframe Alternative Sizing" guide, published in 2012…

MIPS Level	z196 Models	Actual MIPS	z10 EC Models	z10 Actual MIPS	z10 BC Models	z10 BC Actual MIPS	z114 Models	z114 Actual MIPS	HP Cores Estimate	Total HP equivalent MIPS
1,000	2817-701	1,202	2097-701	889	2098-Z02	1250	2818-Z01	782	2	866
2,000	2817-702	2,272	2097-702	1,667	2098-Z03	1784	2818-Z03	2026	5	1,860
3,000	2817-703	3,311	2097-704	3,114	2098-Z05	2760	2818-Z05	3139	8	3,021

Can a 2-chip, quad-core x86-based Blade server really replace 3,000+ MIPS?

- Simple core comparisons are inherently inaccurate…
- Benchmarks can be deceiving…
- Real world use cases suggest this number is off by a factor of **10-20 times**

Eagle TCO study of a mid-sized workload demonstrates how HP's sizing guides are far from accurate

6x 8-way (x86) Production / Dev
2x 64-way (Unix) Production / Dev
Application/MQ/DB2/Dev partitions

$25.4M (5 yr. TCO)

176 processors

1,660 MIPS
(6 processors)

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

$17.9M (5 yr. TCO)

29x
more cores!

Eagle TCO Study shows a pure Intel offload was not effective…

3x HP DL580 (2ch/20co)
Production / Dev / Test
(2011 x86 technology)

z800 Production /
Dev / Test
(2002 mainframe technology)

499 MIPS
(2.1 processors)

60 processors

Despite a 9-year technology gap,
the Intel platform still required
29x more processors

Eagle TCO study shows this small workload was *not* cheaper on the distributed platform

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

$17.9M (4 yr. TCO)

36 processors

41x
more cores!

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

332 MIPS
(0.88 processors)

$4.9M (4 yr. TCO)

Eagle TCO study shows even this VERY small workload was not cheaper on the distributed platform

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

z890 Production / Test

88 MIPS
(0.24 processors)

8 processors

$8.1M (5 yr. TCO)

33x
more cores!

$4.7M (5 yr. TCO)

Slipped Schedules and Major Reduction in Scope – Customer Example

Schedule Objectives			
Met / Not Met	**Scheduled Completion Date**	**Actual Completion Date**	**Variance**
Not Met	Original 6/15/07 moved to 5/31/08 moved to 6/30/08 moved to 6/30/09	11/30/08	165% based on August 9, 2008 transfer of final application

Major Scope Changes
• **Extension of project timeline – September 2006**
• **Substitution of Micro Focus JCL Engine for original product ESPBatch – January 2007**
• **Removal of NDPERS from the migration – July 2007**
• **Removal of DHS TECS/Vision Application – September 2007**
• **Removal of DOT Drivers License Application – December 2007**
• **Removal of all of Phase IV (DHS and ITD Billing Applications) May 2008**
• **Removal of all DOT applications from the migration – July 2008**

- **Original Completion Date 6/2007 – Actual "completion" 11/2009**
- **Major Scope reductions throughout the project**

Slipped Schedules and Major Reduction in Scope

Schedule Objectives			
Met / Not Met	**Baseline Budget**	**Actual Expenditures**	**Variance**
Not Met	$8,271,274	$5,762,037	Applications Migrated: Planned 84, Actual 46 Percentage Completed = 55% CPU Reduction: Planned 77%, Actual 10% Budget Variance Based on Applications Migrated = 127% Budget Variance Based on CPU Reduction = 536%

***35 MIPS (of 350 MIPS)**

- **As usual this was presented in the press as a successful mainframe offload**
  - **Despite failing to remove 87% of planned workload**
  - **Despite the resulting pair of solutions costing more than the single solution did beforehand**
  - **Despite reducing future flexibility**
  - **Despite blowing the migration budget by 536%**
  - **Despite over running by years**

Better understanding of mainframe workloads and the platform can prevent embarking on a bad rehosting experience

The value and advantages of the System z platform

- Perfect workload management
- Multiple environments on one platform
- Disaster Recovery
- …

Why some workloads are best fit on System z

- I/O-intensive workloads
- CICS/COBOL workloads
- “Chatty” workloads
- …

Note that this is not intended to list *all* the advantages of the System z platform, nor is it intended to list *all* workloads that are best fit on System z.

System z is a highly efficient virtualized platform designed to benefit from statistical multiplexing of many workloads

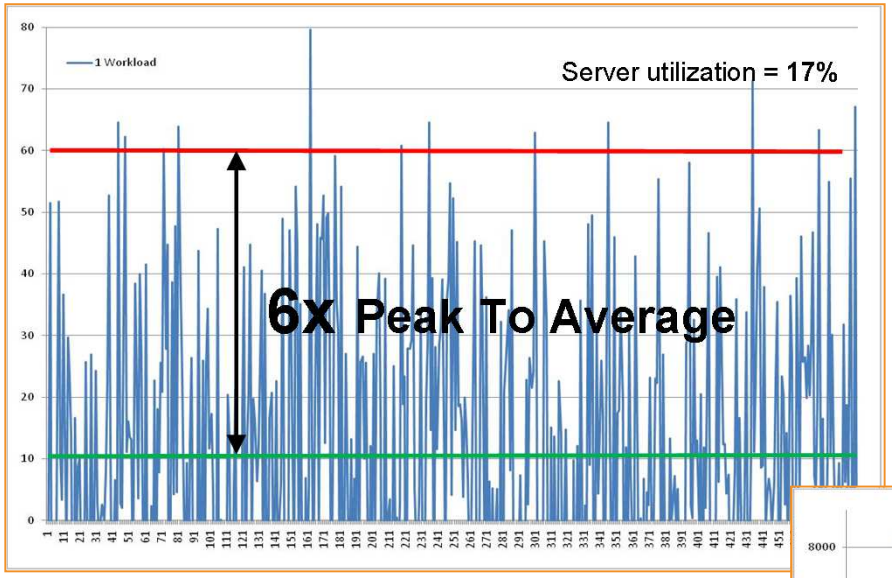

- Consolidating workloads with variance on a virtualized server reduces the overall variance (statistical multiplexing)
- Consequently, larger servers with capacity to run more workloads can be driven to higher average utilization levels without violating service level agreements

1 variable workload:
Machine capacity (red) =
6x average demand (green)

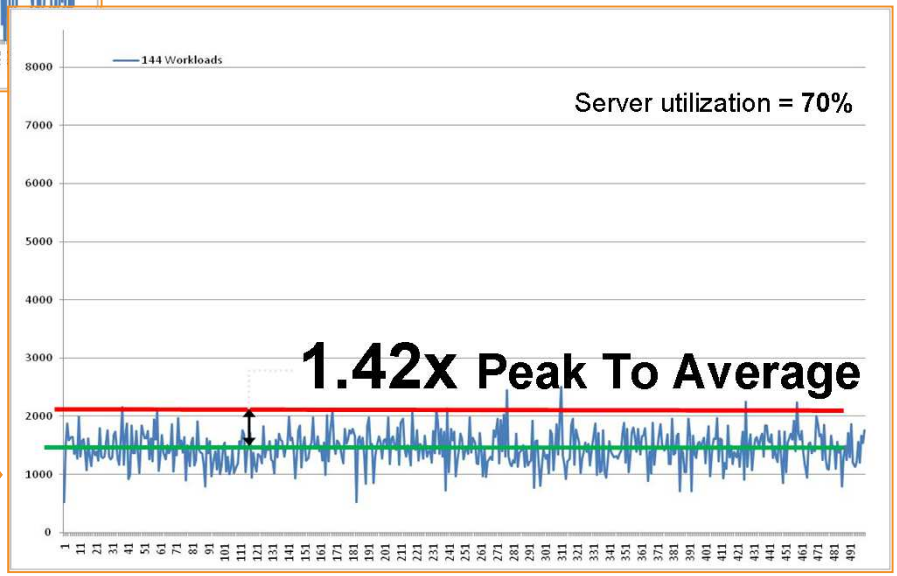

144 variable workloads:
Machine capacity (red) =
1.42x average demand (green)

System z demonstrates perfect workload management…

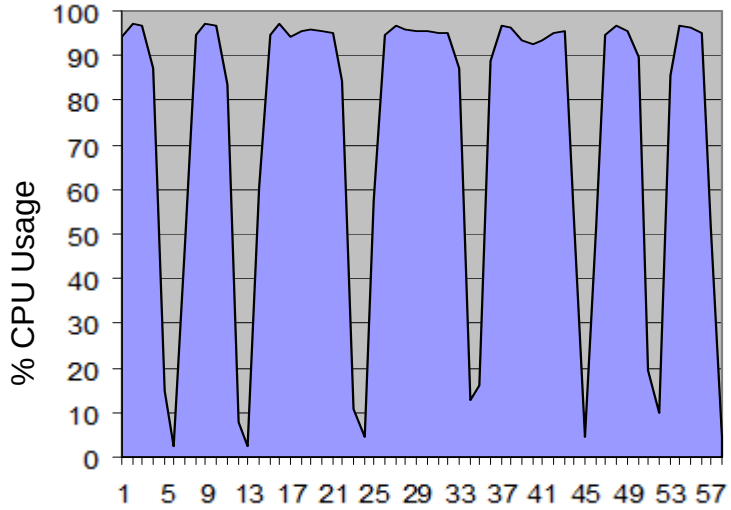

Demand curve for 10 high priority workloads running in 1 z/VM LPAR (PR/SM weight = 99)

- **Workloads consume 72% of available CPU resources**

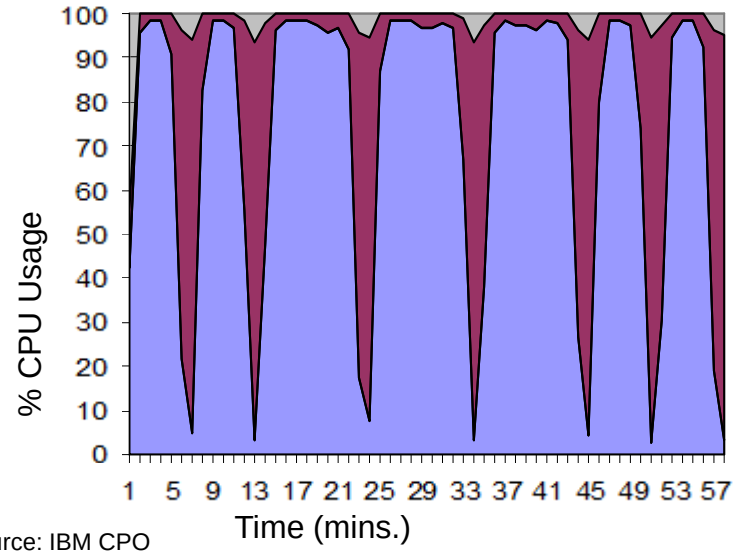

Demand curve when 14 low priority (PR/SM weight = 1) workloads are added in a second z/VM LPAR

- **High priority workload throughput is maintained**
- **No response time degradation**
- **All but 2% of available CPU resources is used**

Source: IBM CPO

…Unlike this common Intel hypervisor which demonstrates imperfect workload management

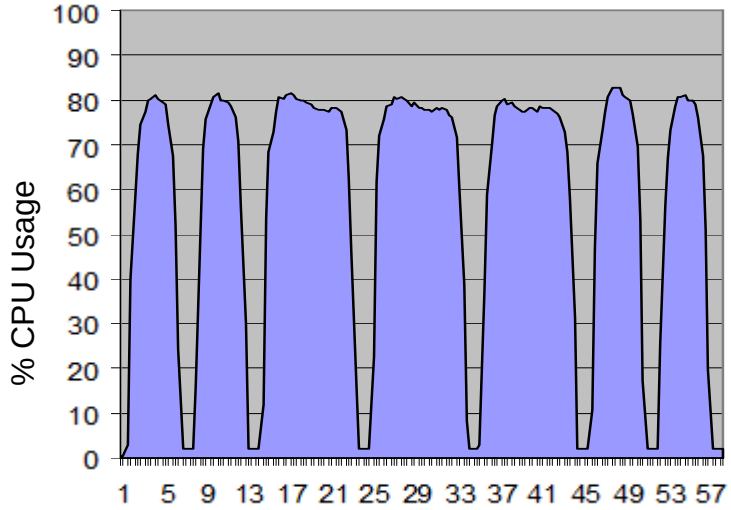

Demand curve for 10 high priority workloads running on a common Intel hypervisor (high share)

- **Workloads consume 58% of available CPU resources**

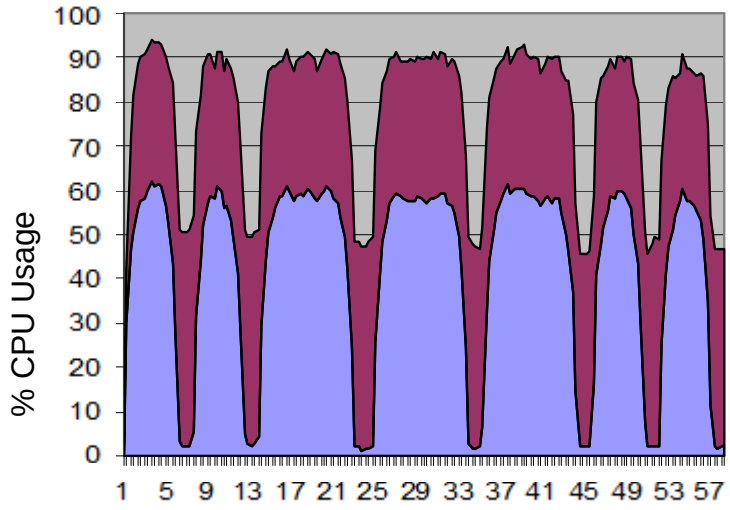

Demand curve when 14 low priority (low share) workloads are added

- **High priority workload throughput drops 31%**
- **Response time degrades 45%**
- **22% of available CPU resources is unused**

Source: IBM CPO

Non-production environments require fewer resources on the mainframe

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

Production *QA* *Dev/Test*

Mainframe Usage Profile

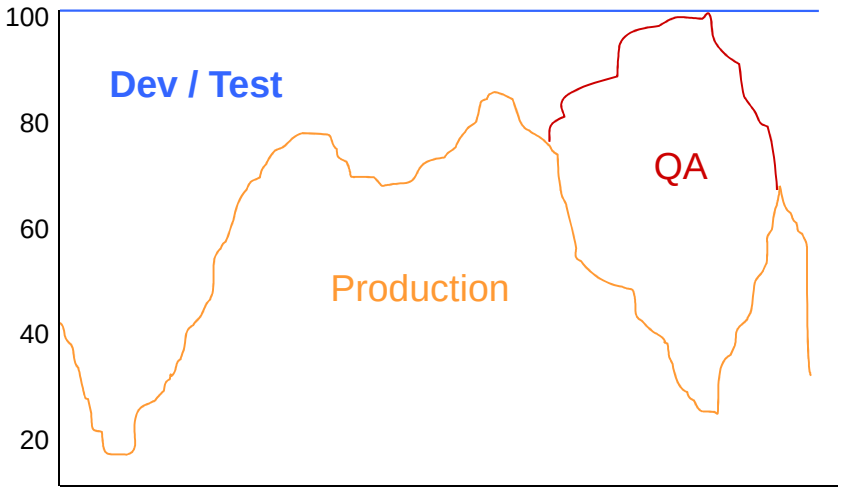

Source: IBM Eagle Team

Disaster Recovery on System z costs much less than on distributed servers

A large European insurance company with mixed distributed and System z environment:

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

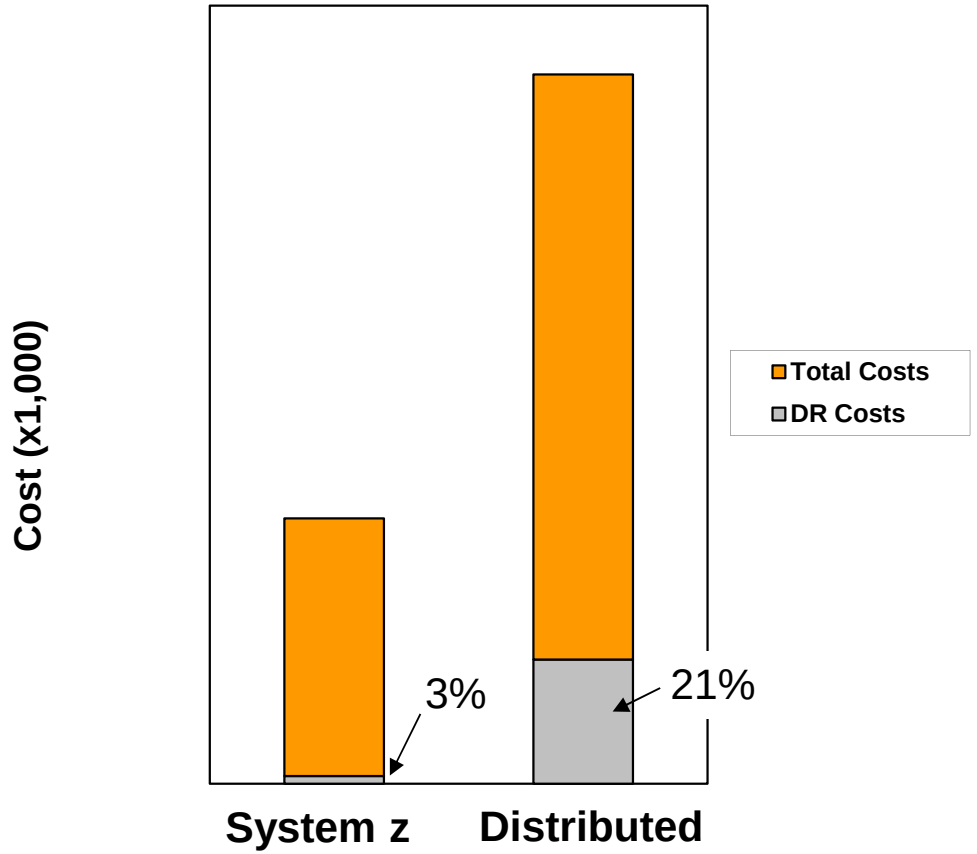

Source: IBM Eagle Team

Better understanding of mainframe workloads and the platform can prevent embarking on a bad rehosting experience

The value and advantages of the System z platform

- Perfect workload management
- Multiple environments on one platform
- Disaster Recovery
- …

Why some workloads are best fit on System z

- I/O-intensive workloads
- CICS/COBOL workloads
- “Chatty” workloads
- …

Note that this is not intended to list *all* the advantages of the System z platform, nor is it intended to list *all* workloads that are best fit on System z.

Dedicated I/O subsystem means System z is ideal for high bandwidth workloads

Capacity benchmark for Bank of China:

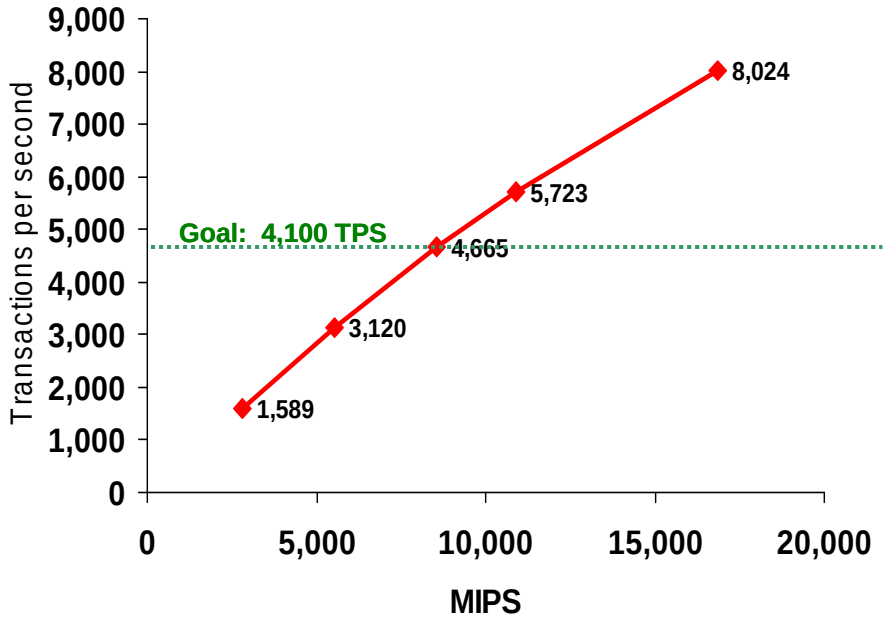

System z easily surpassed benchmark goal, and demonstrates near linear scalability

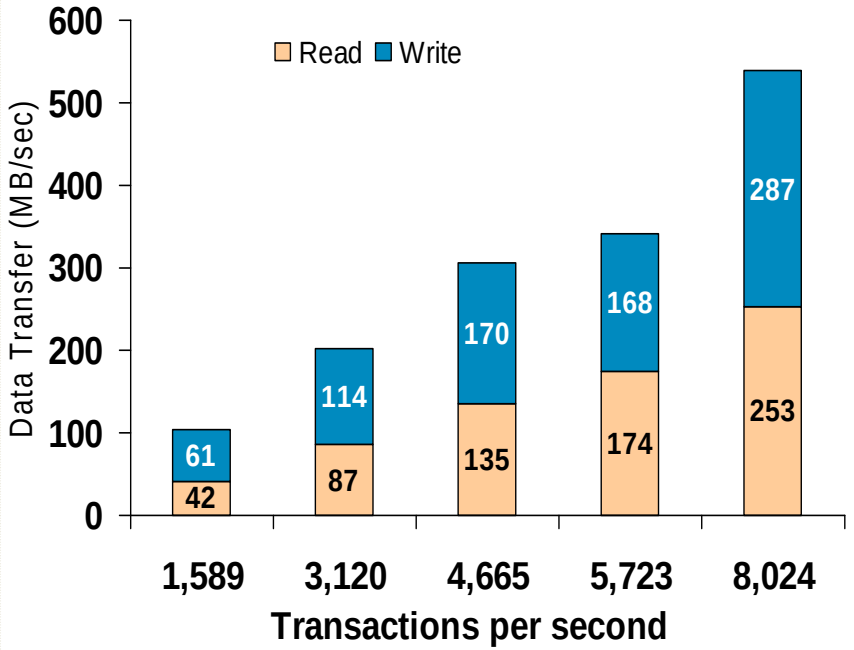

Reads and writes are well-balanced and scale linearly, demonstrating no constraints on I/O

Source: IBM CPO

Comparison test demonstrates System z supports significantly more high I/O bandwidth workloads

Comparison of consolidation platforms

Workloads

Online banking workloads, each driving **22** transactions per second, with **1 MB I/O per transaction**

1 workload per 16-core x86 blade

Virtualized on x86
16 core HX5 Blade

48 workloads per 32-way z/VM

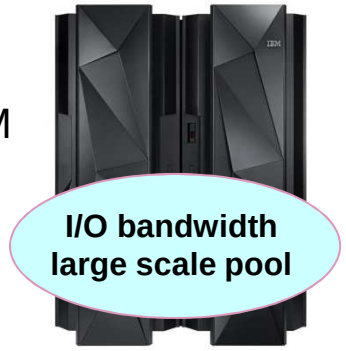

Virtualized on z/VM on zEC12
32 IFLs

***24x** more workload density*

Source: IBM CPO

Customer data demonstrates consolidated Oracle database workloads benefit from System z's high I/O bandwidth

Which platform provides the lowest TCA over 3 years?

Oracle DB workload

Customer Database Workloads each supporting 18K tps

Oracle Enterprise Edition

Oracle Real Application Cluster

3 Oracle RAC clusters
4 server nodes per cluster

12 total HP DL580 servers (192 cores)

$13.2M (3 yr. TCA)

3 Oracle RAC clusters
4 nodes per cluster

Each node is a Linux guest zEC12 with 27 IFLs

$5.7M (3 yr. TCA)

1/2 the cost!

TCA includes hardware, software, maintenance, support and subscription.
Workload Equivalence derived from a proof-of-concept study conducted at a large Cooperative Bank.

Distributed platforms don't often benefit from consolidation and therefore MUST run at low utilizations – even with virtualization!

- **Large insurance company considering moving applications to virtualized x86**
  - Believed this was a high utilization, low cost platform compared to other alternatives
  - Note costs are normally impacted largely by core count and software cost per core

- **Used readily available utilization data to demonstrate extremely low x86 utilization**
  - On average the provisioned systems were used at less than 15%, peak less than 20%
  - This despite many of the 75 hosts running up to 40 VMs each (unusually high)

- **Further investigation shows the various practical constraints that lead to this effect**
  - RAM shortages (normally no physical RAM overcommit allowed)
  - Limited virtual CPU overcommit (vCPU co-scheduling issues)
  - Enforced separation of production from non-production (isolation issues)
  - Limit to the number of VMs per host (to limit workload migration time requirement)
  - Presence of many idle workloads (wasting RAM and driving up the RAM/core ratio to impossible levels, thereby forcing idle cores)

- **System z does not suffer from these issues and normally runs at high utilizations**
  - Averages normally above 50%, often see 65% and above, unheard of on other platforms
  - Most System z machines run more workloads in a single LPAR than other platforms run on the whole physical server, even for large servers – and hence benefit from significant consolidation

Typical x86 hypervisor deployment rules of thumb

- As reported by a very large x86 hypervisor customer

  1. **Host memory virtualization <95%**
  2. **Number of VMs per host <40:1**
  3. **vcpu to physical <4:1**
  4. **PROD and Non-PROD separated**

- Note that none of these constraints are associated with CPU utilization!

- Basically the primary constraint is physical RAM – once virtual machines are defined that represent 95% of physical, no more guests are allowed
- The second constraint is the number of VMs on the server – this is likely due to the hypervisor only allowing a limited number of concurrent migrations, and they take a long time, so shutting down a physical box can require hours of manual work
- The third constraint suggests there may be vCore overheads and/or restrictions around co-scheduling of virtual to physical cores
  - ie. A VM will not run unless all of its virtuals can be backed by physicals at once
- The final constraint is not surprising given what we've seen in workload management testing – non production workloads can easily consume CPU intended for production workloads despite configuring the x86 hypervisor to explicitly avoid that ("bleed" between supposedly isolated workloads)

Consolidation capability is a much larger sizing factor than hardware age or differences in technology

- **Most customers use benchmark results to differentiate servers on some level**
  - e.g. RPE's from Gartner/Ideas, SAPS for SAP, tpcc results, etc.
  - These generally give a good estimate of change in technology over time
  - They may also be able to differentiate between specific technologies for specific workloads

- **Unfortunately these benchmarks do NOT capture the most significant factor in IT today**
  - The ability to effectively consolidate many mixed workloads into a shared pool of resource
    - Shared in this context means all unused resource is available to any workload immediately
    - x86 hypervisors try to define shared across "clusters" of servers but that is not truly shared
    - Oracle RAC tries the same thing with their "shared" cache – still not actually shared!

- **So let's examine the magnitude of the different factors for an example case**
  - To keep things simple we'll stay with "distributed" style workloads and compare x86 to large POWER
    - The result going to System z is dramatically more impressive, especially going to z/OS
  - Benchmarks suggest that the same generation x86 cores and POWER cores are similar on a core-to-core basis (of course some benchmarks say higher, others lower)
    - So the technology factor is basically 1 (and we set age factor to 1 by using same age hardware)
  - So what is a typical consolidation factor (assuming both platforms are virtualized)?

- **A typical real customer example for a single mid-sized application suggests a 3-4x factor**
  - So the ability to consolidate effectively is substantially more important in sizing than benchmarks suggest
  - Also note benchmarks tend to be single-image, incorrectly implying less capacity per additional core

Most customer cost tracking (or chargeback) creates adverse selection issues which require unintuitive steps to correct

- **Simple thought experiment involving a single shared infrastructure system**
  - We'll use a real customer example of an x86 based "private cloud" (production only)
    - 600+ virtual servers, mostly running Windows, some Linux
  - Calculating an average cost per VM for the whole system gives us a $/VM/yr

- **Now divide the same infrastructure into two pools**
  - Large pool of lightest VMs taking about 50% of the actual resource consumed
  - Small pool of heaviest VMs, also taking about 50% of the actual resource consumed
  - Calculate an average cost per VM for these two new pools
    - The light VM pool is now much cheaper than the heavy pool on a per VM basis
    - This is because both pools use about the same resources, but have very different VM counts

- **Follow the logic to understand that this creates a stunning case of adverse selection**
  - Line of Business users are now incented to move workload from the heavy to the light pool
  - Even though the total cost of the whole infrastructure is unchanged whatever happens!

- **What if you have a pool which is actually more efficient for most workloads?**
  - It tends to run the heaviest workloads, and therefore shows up as more expensive – sound familiar?
  - The incentives and cost data tell the business they should move off to the cheaper pool(s)
  - The real solution is actually the opposite – to move more workload onto the more efficient pool
    - This is trivial to prove mathematically, I'll leave that as an exercise for the reader!

Eagle studies show 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 resulted in CICS applications no longer meeting its SLA

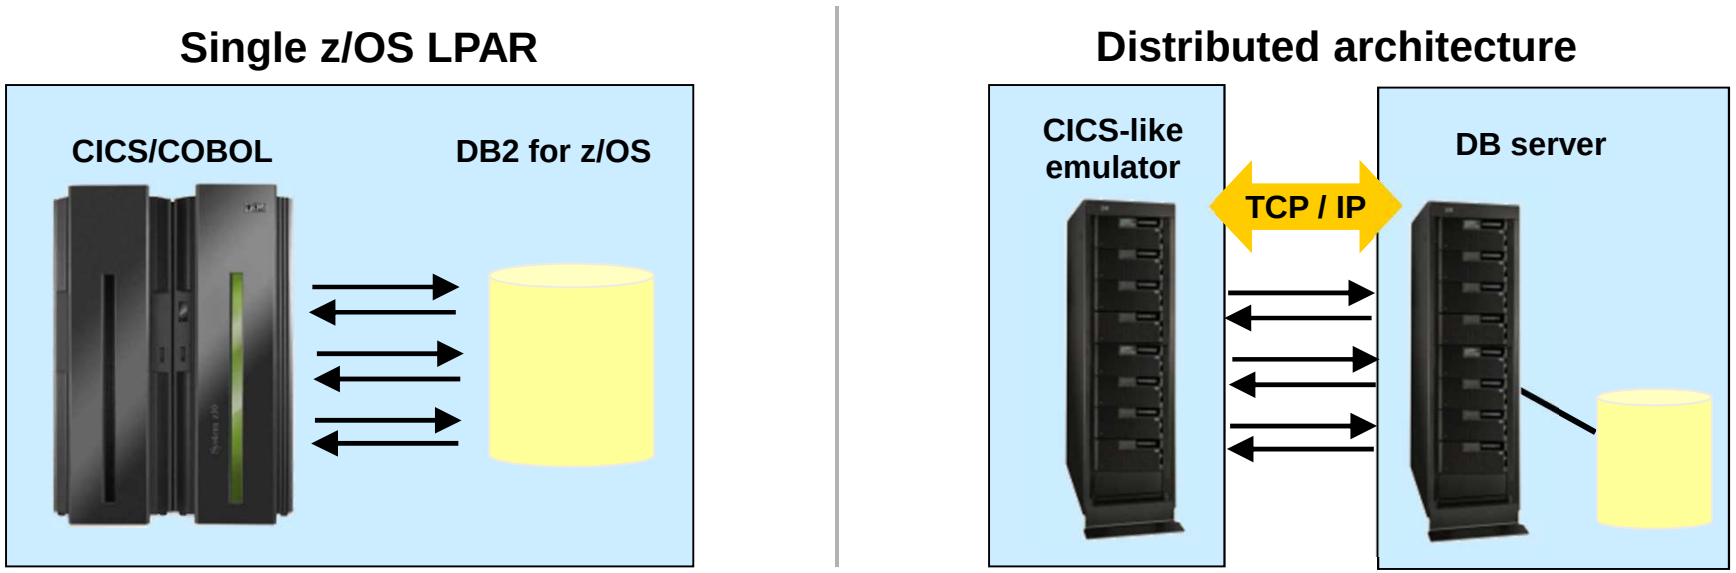

Source: IBM Eagle Team

Before you start a rehosting project, make sure you have evaluated *all* the risks

Look for hidden costs like:

- Missing functionality
- Sub-optimized performance
- Risks of failure

The IBM Eagle Team

Replacement technologies are not always available for many mainframe functions

Rehosted platform

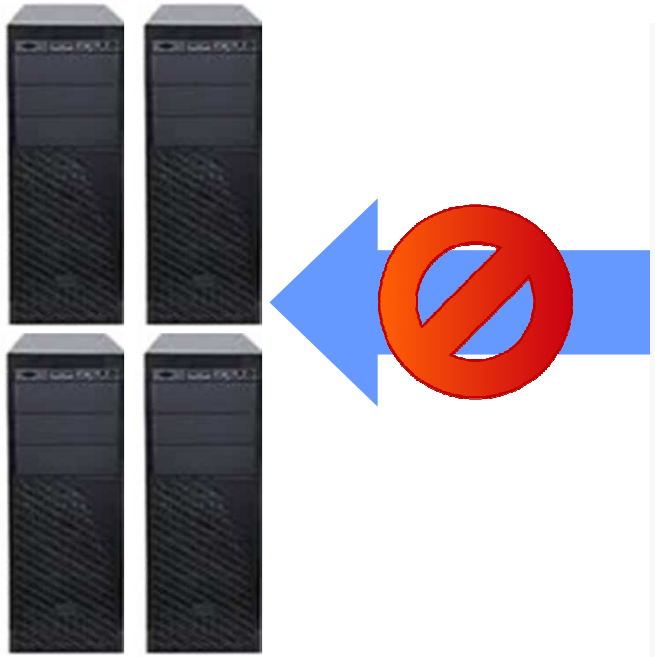

- Hierarchical databases – e.g., IMS DB and IMS DC
- Languages – e.g., PL/I, ASM …
- Batch environments including JCL with symbolic substitution, Batch pipes, Generation Data Group files for batch recovery
- System management and database tools
- 3270-style user interfaces, BMS maps, APIs…
- File structures – e.g., VSAM (alternate indexes not supported), QSAM and Partitioned Data Sets
- Print facilities including PSF, AFP, Info Print Server, JES2/3 spool
- Ability to read old backup tapes

Eagle studies for two US retailers highlight missing systems management functionality

Rehosted platform

- 15 replacement applications (7.5%)
- Cost = $8.4M OTC + $1.8M annually

- 53 replacement applications identifies (20%)

- 200 systems management products in total

- 261 systems management products in total

- Options?
  - Re-write applications to avoid usage
  - Write new code to perform the function
  - Add staff to manually perform the function

Source: IBM Eagle Team

Offloading CICS application results in suboptimal performance

- Offload project to move State of Montana Department of Motor Vehicles license registration system from CICS to Microsoft
  - Performed by Microsoft and Bearing Point
  - Cost of project $28.3M, 3 years late

	Response time
Before offload	Sub-second
After offload	30+ seconds

"Transferring titles is taking two to three hours instead 15 minutes."

One employee said she had never heard so many "four-letter words" from customers.

Source: http://spectrum.ieee.org/riskfactor/computing/it/montana-new-registration-and-licensing-system-still-having-hiccups

Customer tests show IBM Enterprise COBOL performs better than competition

- IBM Enterprise COBOL on z/OS performed best
- Micro Focus COBOL is a COBOL interpreter, and code is over 4.5 times less efficient
- ACUCOBOL, a compiler acquired by Micro Focus, was 12 times less efficient
- Micro Focus functional differences required additional debugging

Source: IBM CPO

Different compilers may potentially lead to different COBOL behavior

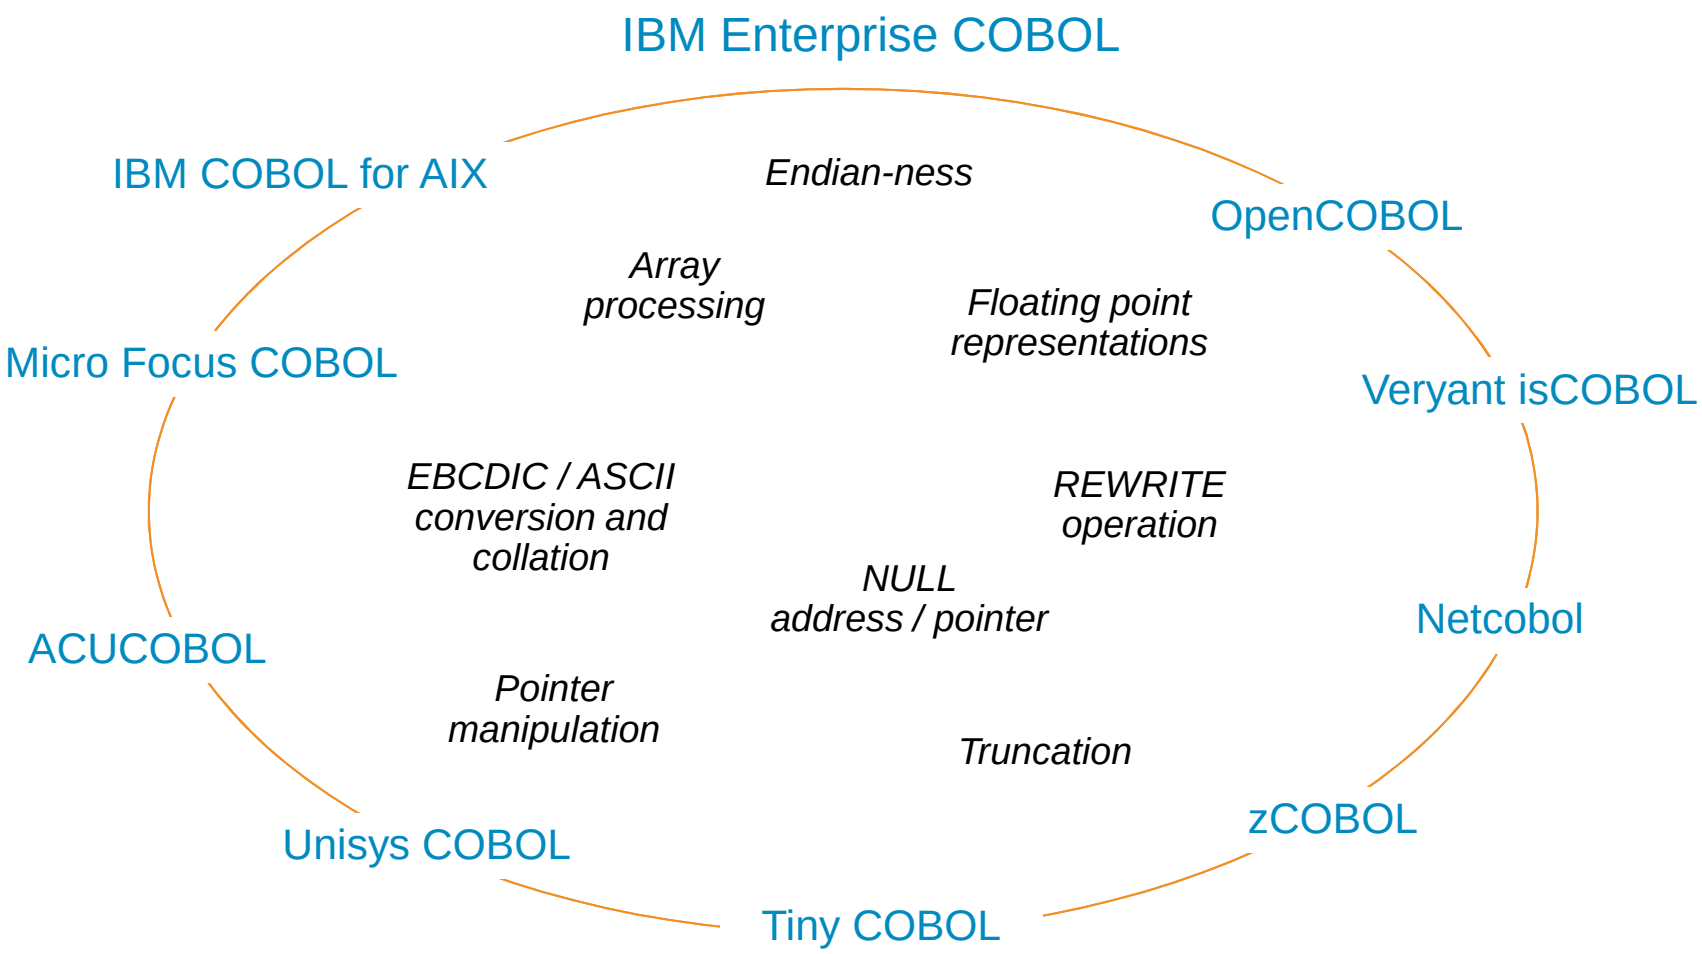

See http://download.oracle.com/docs/cd/E18050_01/artwb/docs11gr1/wbref/CobolConverter.html

Code stability is at risk on some distributed platforms

- Mature System z software is very stable
- Distributed software is typically less so…

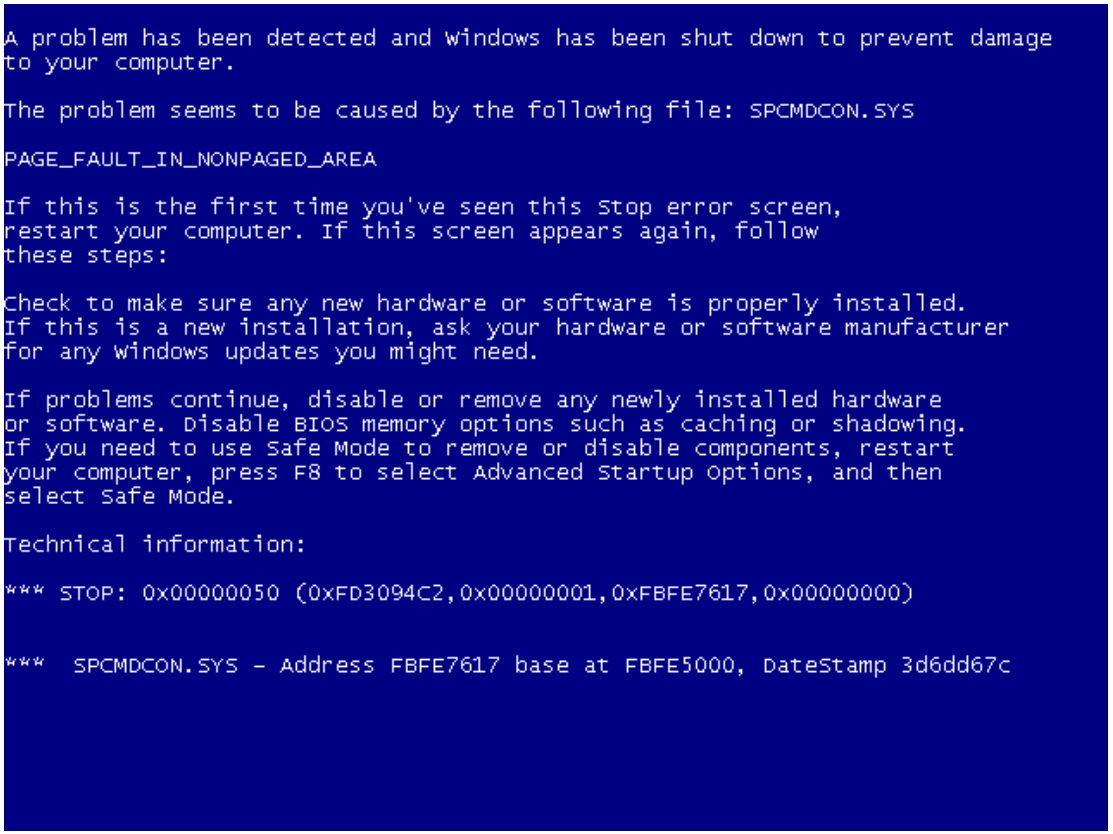

Familiar Microsoft “Blue Screen Of Death”

Oracle patches far outnumber those for DB2 on z/OS

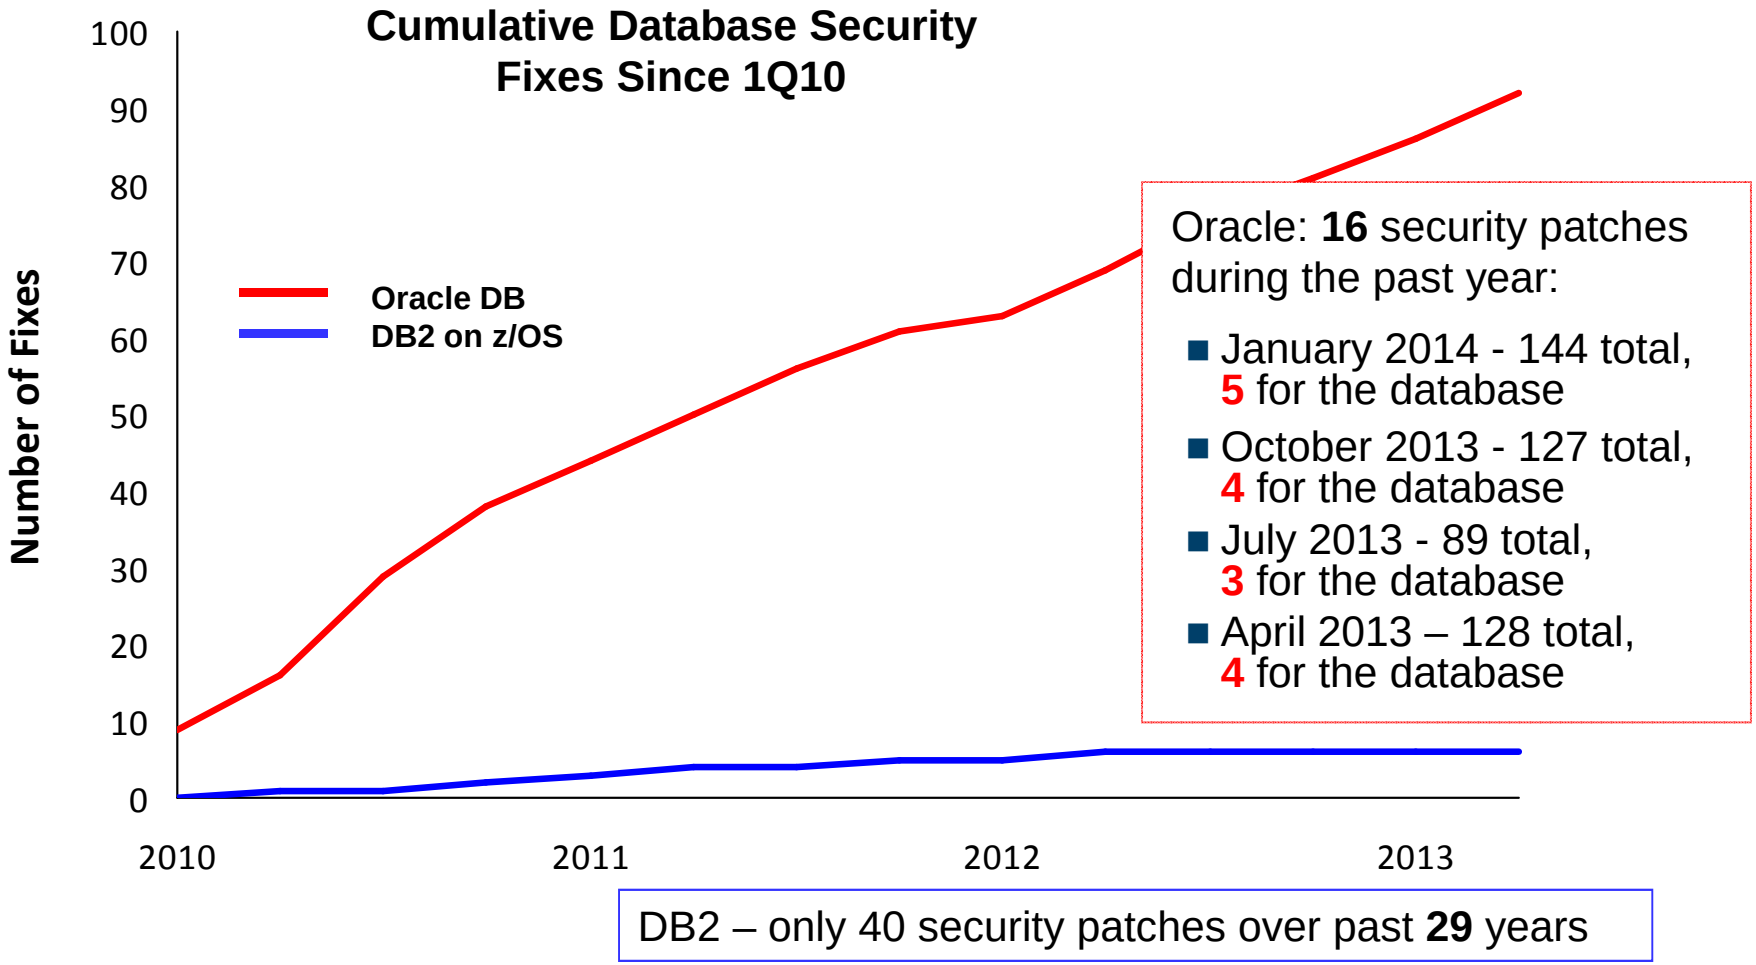

Source: http://www.oracle.com/technetwork/topics/security

In 2005, Canadian insurance company partnered with Micro Focus on a rehosting project…

Lombard Canada Ltd. wanted 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…"

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
- Millions of dollars spent 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

Asian bank project demonstrates another more recent example of failed rehosting

- 60 MIPS CICS/COBOL application plus additional 30 MIPS of Batch processing
  - 2.8M lines of COBOL code
  - 123K LOC in Assembler
  - 44K LOC of JCL
- IMS DB remained on System z
- **Two years later:**
  - Project abandoned after failing to complete development
  - $5.7M spent but unable to estimate eventual deployment costs
  - Team of 10 was disbanded and left the business – no one could describe the problems encountered
  - Management responsible was fired

Source: IBM Eagle Team

Ongoing rehosting project at US Retail company provides another example of the risks involved

Customer's stated objective:

- Offload 3,500 MIPS with Micro Focus…
- $10M budget…
- 1 year schedule…

- **18 months later:**
  - $60M spent, but only 350 MIPS offloaded
  - Increased staff to cover over-run
  - Required additional hardware over initial prediction
  - Implemented manual steps to replace mainframe automation
  - Extended the dual-running period of the rehost…
  - Executive sponsor no longer employed…

Eagle team had advised against this offload…

Source: IBM Eagle Team

Recent US government agency rehosting project also had to be abandoned

- 360 MIPS of CICS/COBOL for payroll and HR
  - 4M lines of COBOL code
  - Estimated 270K LOC needed to be changed
- Additional 30 MIPS of batch
- IMS DB to stay on System z
- Agency estimated a 5 year contract worth $80M to perform this offload

- **Project abandoned** and manager responsible for the decision left

Source: IBM Eagle Team

Can a rehosting vendor really meet your SLA requirements?

Distributed

- Insist the solution includes the same levels of backup, availability and disaster recovery.
- Can the same levels and complexity be reached? What is the *cost*? How much testing will be involved?
- Is this attainable? Can this be *guaranteed*?
- How many *years* have you spent fine-tuning? Are you prepared to spend that again – maybe more – to reach the same levels?

Mainframe

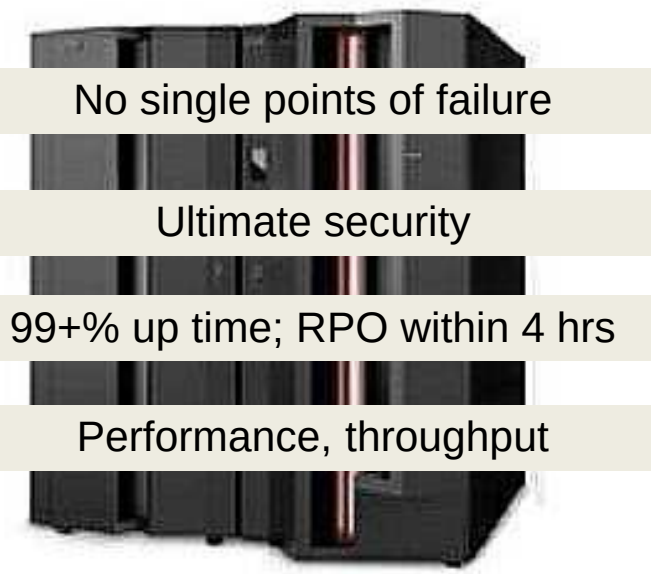

- No single points of failure
- Ultimate security
- 99+% up time; RPO within 4 hrs
- Performance, throughput

Know the risks! Know the costs!

What's next?

- **Re-examine your cost concerns; make sure chargebacks are accurate**
  - Read Forbes' (Dec 2013) It's 10 O'Clock -- Do You Know Where Your IT Costs Are?

    …*Applying IT financial management is like deploying enterprise architecture or setting up enterprise standards: a horribly messy discussion but one that will imbue structure into everything you do in the coming years. It's time to stop and think about budget, value, and costs…*

- **Examine the productivity of your mainframe compared to equivalent distributed platforms.**
  - Which generates more throughput in less time? Which gives you best cost per unit of work? Which has the lowest downtime and best security?
- **Ask IBM for an Eagle study… or a Portfolio Review and Analysis**

Remember:

Examine all costs and all risks; understand what the ROI will be

Consider upgrading the mainframe as a lower risk alternative

When talking to rehosting companies:

Ask to follow customers who are in the process of similar migrations rather than "completed" cases


thank
you!