

More/Slower vs. Fewer/Faster CPUs: Practical Considerations in 2024

Scott Chapman
Enterprise Performance Strategies, Inc.
Scott.Chapman@EPStrategies.com



SHARE Kansas City
August 2024

Contact, Copyright, and Trademarks



Questions?

Send email to performance.questions@EPStrategies.com, or visit our website at <https://www.epstrategies.com> or <http://www.pivotor.com>.

Copyright Notice:

© Enterprise Performance Strategies, Inc. All rights reserved. No part of this material may be reproduced, distributed, stored in a retrieval system, transmitted, displayed, published or broadcast in any form or by any means, electronic, mechanical, photocopy, recording, or otherwise, without the prior written permission of Enterprise Performance Strategies. To obtain written permission please contact Enterprise Performance Strategies, Inc. Contact information can be obtained by visiting <http://www.epstrategies.com>.

Trademarks:

Enterprise Performance Strategies, Inc. presentation materials contain trademarks and registered trademarks of several companies.

The following are trademarks of Enterprise Performance Strategies, Inc.: **Health Check[®], Reductions[®], Pivotor[®]**

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries: IBM[®], z/OS[®], zSeries[®], WebSphere[®], CICS[®], DB2[®], S390[®], WebSphere Application Server[®], and many others.

Other trademarks and registered trademarks may exist in this presentation

Abstract (why you're here!)



Mainframe customers have had a choice of processor speeds for years. As full-capacity engines have grown, more and more customers are finding that those full-capacity engines are not an ideal fit for them. For some customers the z16 A02 may even be preferable to an A01. What are the advantages to using sub-capacity engines? Does your LPAR configuration matter? What metrics should be examined to determine if a more/slower configuration is a good fit for your workload? Is it possible that this choice might impact your software bill?

Come to this session with Scott Chapman to learn why more/slower CPUs may be a better fit for many environments and how to determine what workloads might be at risk for moving to slower CPUs.

Agenda



- Mainframe processor nomenclatures
- Software cost issues
- Performance and processor caches
- Sharing CPs
- Capacity concerns
- Expected trade-offs
- Modeling changes with example scenarios
- Wrap-up

EPS: We do z/OS performance...



- Pivotor - Reporting and analysis software and services
 - Not just reporting, but analysis-based reporting based on our expertise
- Education and instruction
 - We have taught our z/OS performance workshops all over the world
- Consulting
 - Performance war rooms: concentrated, highly productive group discussions and analysis
- Information
 - We present around the world and participate in online forums
 - <https://www.pivotor.com/content.html>
 - <https://www.pivotor.com/webinar.html>



z/OS Performance workshops available



During these workshops you will be analyzing your own data!

- WLM Performance and Re-evaluating Goals
 - February 19-23, 2024
- Parallel Sysplex and z/OS Performance Tuning
 - August 20-21, 2024
- Essential z/OS Performance Tuning
 - October 7-11, 2024
- Also... please make sure you are signed up for our free monthly z/OS educational webinars! (email contact@epstrategies.com)

Like what you see?



- The z/OS Performance Graphs you see here come from Pivotor
- If you don't see them in your performance reporting tool, or you just want a free cursory performance review of your environment, let us know!
 - We're always happy to process a day's worth of data and show you the results
 - See also: <http://pivotor.com/cursoryReview.html>
- We also have a **free** Pivotor offering available as well
 - 1 System, SMF 70-72 only, 7 Day retention
 - That still encompasses over 100 reports!

All Charts (132 reports, 258 charts)

All charts in this reportset.

Charts Warranting Investigation Due to Exception Counts (2 reports, 6 charts, [more details](#))

Charts containing more than the threshold number of exceptions

All Charts with Exceptions (2 reports, 8 charts, [more details](#))

Charts containing any number of exceptions

Evaluating WLM Velocity Goals (4 reports, 35 charts, [more details](#))

This playlist walks through several reports that will be useful in while conducting a WLM velocity goal an.

EPS presentations this week



What	Who	When	Where
60 Years of Pushing Performance Boundaries with the Mainframe	Scott Chapman	Sun 17:00	Neptune D
Introduction to Parallel Sysplex and Data Sharing	Peter Enrico	Mon 13:15	Pomona
Macro to Micro: Understanding z/OS Performance Moment by Moment	Scott Chapman	Mon 15:45	Neptune D
WLM Turns 30! : A Retrospective and Lessons Learned	Peter Enrico	Tue 10:30	Neptune D
PSP: z/OS Performance Spotlight: Some Top Things You May Not Know	Peter Enrico Scott Chapman	Tue 13:00	Pomona
More/Slower vs. Fewer/Faster CPUs: Practical Considerations in 2024	Scott Chapman	Tue 14:15	Neptune D
z16 SMF 113s – Understanding Processor Cache Counters	Peter Enrico	Wed 13:15	Pomona

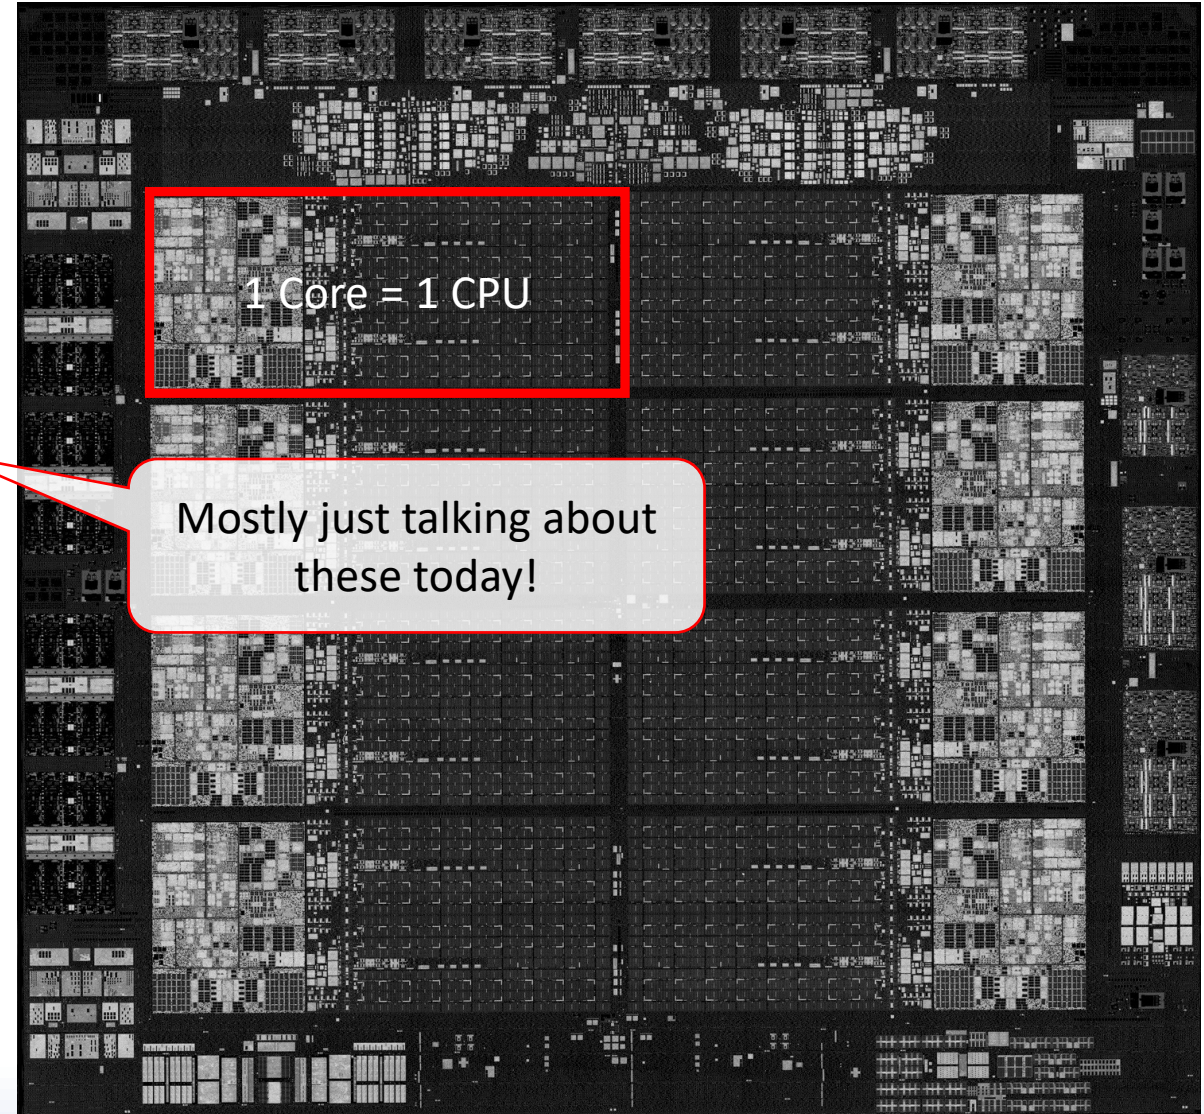


Processor Models and Nomenclature

CPU Characterizations



- Generically, a CPU is a core on a chip
- CPUs can be “characterized” for use with specific work:
 - GP (General Purpose, runs anything)
 - zIIP (Runs specific z/OS workloads)
 - IFL (Linux partitions)
 - ICF (Coupling Facility Partitions)
 - SAP (Service Assist Processor)
 - IFP (Integrated Firmware Processor)
 - Spare (not characterized)
- All are physically the same!



Measuring Capacity




- Capacity = How much of a certain type of work can be done per unit of time
- There are three terms we use to express the capacity:
 - MIPS (or PCI in IBM-speak)
 - MSUs
 - SU/Sec
- IBM publishes PCI, MSU, and SU/sec ratings for each machine
 - All are derived from the same IBM test workload results (LSPR)
 - All are pretty much the same, just different scales
 - I.E. there's an essentially constant relationship between PCI, MSU and SU/sec
- Other vendors publish MIPS ratings for machines
 - Which may vary from IBM PCI

GHz is not Capacity!



- Processor GHz is a measure of clock speed
 - Clock speed = timing function that controls on-chip signaling
 - Higher clock speed *may* imply the ability to execute instructions faster, but it's not that simple in terms how much useful work can be done per unit of time



What is a “slower” or “faster” CP?



- Sometimes (often) customers will specify a machine with “sub-capacity” (aka “sub-cap”, aka “knee-capped”) processors
 - These processors deliver less useful work per unit of time than the “full-cap” models
 - I.E. sub-cap CPs are “slower”
 - But... their clock speed is the same (although some reporting might estimate an “effective” clock speed)
 - **This is only for GP processors!** (zIIPs et al always run full speed)
- This counter-intuitive offering is used to:
 - Optimize (reduce) purchased capacity to limit software costs
 - Avoid performance & capacity problems caused by having too few CPs
 - But too slow CPs can be a problem too!

Mainframe Models



- Large, or EC (Enterprise Class), or x01, or multi-frame machines:
 - 4 “speed” settings: 4xx, 5xx, 6xx, 7xx
 - 7xx is “full” speed/capacity
 - The xx is replaced with the number of GPs
 - E.G. 410 = a machine with 10 of the slowest speed engines
 - E.G. 747 = a machine with 47 of the fastest speed engines
 - Sub-cap models limited to ~30 GPs
- Small, or BC (Business Class) or x02, or single-frame machines:
 - 26 “speed” settings: Axx (slowest) – Zxx (fastest)
 - Similarly, xx replaced with number of GPs (limited to 6 in recent machines)
- For a given generation, BC machines run at a lower clock speed than EC
 - Allows for manufacturing yield improvements (for one reason)

This session's question




- Different configurations of more/slower CPs and fewer/faster CPs may have very similar rated capacity: How do you decide?
- E.G. z15-710 and z15-620 both rated at 2037 MSUs
 - Will they actually deliver the same *effective* capacity though?
 - Your workloads and configuration is not the same as the IBM benchmarks
- Usually there will be multiple systems within a few MSUs that look like they might be plausibly delivering the same capacity
 - For many real-world configurations, there may be significant effective capacity differences between these “similar” machines



Software Costs

A very brief and somewhat simplified introduction

Why talk about software costs?



- Software costs are a significant portion of your mainframe TCO
 - Could easily be >50%
- Software costs generally driven by a capacity rating for the machine
- “Rated” capacity and may not reflect “effective” capacity
 - “Rated” = based on IBM test configurations and test workloads
 - Published: <https://www.ibm.com/support/pages/ibm-z-large-systems-performance-reference>
 - “Effective” = actual ability to get work done in the real world
- Your configuration and workloads will not be the same as the IBM test that determined the rated capacity
 - You may be able to get more work done at a lower rated capacity if you chose your engine speed configuration wisely

IBM Pricing Metrics



- Per-CEC
 - Flat charge per physical machine
- Full-capacity
 - Based on the full rated capacity of the machine
 - zIPLA (zOTC) products and some VUE products/agreements
- Sub-capacity
 - Based on how much of the capacity is actually used by the workloads
 - Primarily used for MLC products
 - Rolling 4-Hour Average (R4HA): price based on “monthly” peak rolling 4 hour average consumption
 - Tailored Fit Pricing (TFP): price based on total capacity consumed over the year

MSUs or MSUs?



- IBM doc refers to pricing for R4HA and TFP as MSUs

- But those are different MSU measurements!
- Example using a machine rated at 500 MSUs
- R4HA: Peak rolling 4HA utilization was 90% busy = 450 MSUs

- TFP:


Hour	% Util	MSUs
1	50%	250
2	75%	375
...
730	90%	450
Total		237,250

IBM may be encouraging you to move from R4HA to TFP. Should you?

- This may be a good deal if you're growing your workload.
- This could be a very bad deal if you are (or expect to be) shrinking your workload!

- Really with TFP, MSUs = MSU-hours (and we designate them as such on our reports)
- R4HA = pay for “peak” consumption, TFP = pay for all consumption in all hours

MSU Averages Comparisons




Example of R4HA where the cost for the month would be based on the highest peak blue dot.

Except in this case they're using a group cap to limit their cost exposure, at the expense of a capacity constraint when the R4HA exceeds the cap. (Which of course could impact performance.)

LPAR MSU-hour Cumulative Totals


2023-03-02 - 2024-03-01



Example of TFP, where the total consumption for the year is what drives the total cost for the year.



MSU-hour Totals by MLC Month



Comparing R4HA MSUs vs MSU-hours month by month.

Note that the two measurements don't always move in the same direction!

I.E. is a change in the peak period or in total across the month?

ISV Pricing




- Can be almost any model they dream up
- Often (alas) based on full installed capacity
- Many will convert to some form of a sub-cap pricing metric if pushed
 - Do this at contract renewal time!
 - Prepare to switch vendors if necessary
- Best ones tie pricing to something other than machine capacity
 - One print services vendor would base their cost on the number of printers you had
 - Pivotor: (basically) how much data we're processing for you

Outsourcers / MSPs



- If you outsource your mainframe operations, your pricing structure may be different
- In some cases, customer still pays IBM / ISVs for the software
- In some cases, outsourcer/MSP pays the vendors
 - Customer costs in this case could be based on any metric the MSP comes up with
 - For small customers in particular, there *may* be an opportunity to save money under this model

What are you optimizing for?



- Performance
 - A few customers need to maximize performance with limited regard to cost
- Balance
 - Many customers try to optimize for “acceptable” performance while limiting costs
- Cost
 - Some customers set limits to the costs (usually through capping) and try to make performance acceptable within the resulting capacity limits
- Variable costs are based on how much rated capacity you need/use
- Note: may have more opportunities for “software cost engineering” under RH4A where you only worry about your peak(s)
 - OTOH, TFP allows for easier, more direct relationship of consumption to cost

More effective capacity per rated capacity means better software cost optimization



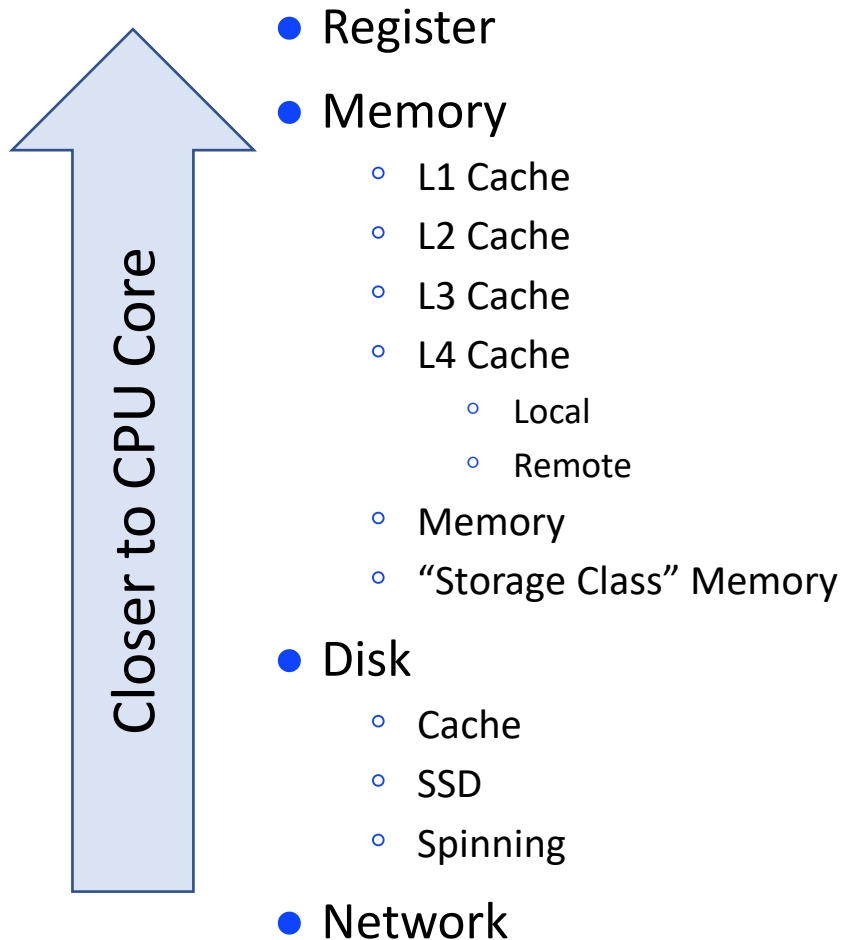
Performance & Processor Caches

Clock Speed and Cycles



- In one z16 clock cycle, light in a vacuum can only travel just over 2 inches!
 - Electrical signal in a circuit is much slower (40-70% of c)
 - 1 meter in fiber \sim 5 ns (>25 clock cycles!)
- Need to make a round trip
- Signal paths aren't as the mosquito flies
 - IBM's "Miles of wire in the chip" numbers:
 - zEC12 – 7.7 miles
 - z13 – Over 13 miles
 - z14 – 14 miles
 - z15 – 15.6 miles
 - z16 – 19 miles
- Physical distance matters!

Data Access Hierarchy




The farther the data is away from the processor, the more clock cycles will be spent accessing it.

Optimal performance & capacity utilization = keeping data as close to processor as possible!

CP CPU L1 Sourcing - Pct Breakdown by Cache Area

SMF 113

CP, SYSJ, z15, 8561, T01, 435B8



Note most L1 misses satisfied out of L2, then L3, very few out of L4 or memory.


This is common.



RNI - Breakdown by Memory Area for System

SMF 113

CP, SYSJ




The RNI (Relative Nest Intensity) is actually pretty good for this system. But note that a significant portion of the total RNI comes from the L4 and memory accesses.

RNI - Breakdown by Memory Area for System

SMF 113

CP, SYSJ



Changed to stack 100% graph just to show 40-50% of the RNI is due to those L4 and Memory references that make up a tiny portion of the total L1 misses.

This is all just to illustrate how much distance matters.



Sharing CPs

Physical & Logical Processors



- You've paid IBM for access to a certain number of physical processors
 - Actual processor cores
- You define a certain number of logical CPs to each LPAR
 - Logical CPs for an LPAR must be \leq physical CPs in the machine
 - Can have reserved CPs as well
- PR/SM assigns physical CP to the logical CPs on a time-sliced basis


Guaranteed Share as Processors



- Each LPAR's share can be translated into a number of processors
 - $LPAR\ Guaranteed\ Processors = LPAR\ Share * Shared\ Processor\ Count$
- In below example, there are 6 shared processors so:
 - $SYSB = 500/1000 * 6 = 3$ processors
 - $SYSC = 350/1000 * 6 = 2.1$ processors
 - $SYSD = 150/1000 * 6 = 0.9$ processors

z/OS LPARs with dedicated CPs are rare, but shown here to point out those dedicated CPs are separate

Dedicated to SYA



For ease of use, try to make weights add up to 1000 (like they do here).


Shared by SYSB, SYSC, SYSD

HiperDispatch




- HiperDispatch manages CPs “vertically”, meaning it endeavors to make the logical CPs a larger percentage of a physical
- Logical processors classified as:
 - High – The processor is quasi-dedicated to the LPAR (100% share) (VH)
 - Medium – Share between 0% and 100% (VM)
 - Low – Unneeded to satisfy LPAR’s weight (VL)
- This processor classification is sometimes referred to as “vertical” or “polarity” or “pool”
 - E.G. Vertical High = VH = High Polarity = High Pool = HP
- Parked / Unparked
 - Initially, VL processors are “parked”: work is not dispatched to them
 - VL processors may become unparked (eligible for work) if there is demand and available capacity

Physical to Logical: Vertical Mgt



With HiperDispatch, vertical high CPs are quasi-dedicated to an LPAR. Note that SYSB's VLs will only come into play when there's both demand from SYSB and the other LPARs aren't using the capacity.



Note that while reality may be a bit messier, vertical CPU management does greatly reduce the movement of logicals to different physicals. Also note VH CPs get longer dispatch intervals.

Important Note: 1 CPU = 1 Task



- A physical CPU can only be executing 1 task on 1 LPAR at any given moment in time
 - Time slicing just makes it seem like multiple tasks are running concurrently
 - SMT (for zIIPs, IFLs) allows 2 tasks within 1 LPAR to run simultaneously
 - But likely the 2 tasks will contend for core resources and so run somewhat slower
- Logical CPs can only execute work when PR/SM has assigned a physical CP to them
 - And PR/SM can/will take that physical away



Capacity Concerns


Some of which have become worse over time...

Sub-capacity Capacity Increases (or not)




- IBM sets the capacity of the sub-capacity models
- Sub-capacity models may not see the same per-processor capacity/performance increase that the full-speed machines see
 - Interesting that for some z16 A02 capacity settings, they dialed capacity *down* from the z15 T02 level for the same step
- Whether this is good or bad depends on your specific situation
 - **Always use zPCR to model your proposed upgrade!**

Dual-frame Uniprocessor Sizes (PCI)



Selected Single-frame Uniprocessor Sizes (PCI)



Big Processors = Big Steps




- The 7xx machines keep getting “faster” processors
 - This is obviously expected: you can’t release a new processor that’s slower!
- But larger “steps” makes “add an engine” upgrades more expensive
 - Also limits choices for selecting “right sized” machine
 - Choosing slower engines can mitigate those issues

	Generational Increase			
	z13	z14	z15	z16
401	3%	3%	3%	6%
501	18%	1%	4%	19%
601	13%	1%	6%	28%
701	12%	8%	11%	10%

	Uniprocessor capacity in MSUs				
	zEC12	z13	z14	z15	z16
401	30	31	32	33	35
501	80	94	95	99	118
601	119	134	136	144	185
701	188	210	227	253	278

	"Next step" from 10-way to 11-way			
	z13	z14	z15	z16
411	22	23	24	25
511	55	65	67	70
611	79	85	91	95
711	120	132	146	167

Big Difference on Big Machines



- Widening gap between sub-capacity models on the “big” machines increases risk for moving to slower speed engines
- 6xx in particular have not kept up with the 7xx series machines
 - Always has been a big step between 4xx and 5xx too


	Percent of full capacity engines			
	z13	z14	z15	z16
401	15%	14%	13%	13%
501	45%	42%	39%	42%
601	64%	60%	57%	67%
701	100%	100%	100%	100%

	Percent of next faster CP			
	z13	z14	z15	z16
401	38%	33%	34%	33%
501	67%	70%	70%	69%
601	63%	64%	60%	57%
701	N/A	N/A	N/A	N/A

Limitations on Small Machines



- Limitation of 6 GPs on the “small” machines makes it harder for customers on the “large” machines to go down to the “small” machines
 - Used to be actively discouraged, but I have seen customers do it
- Smallest engines have increased more than largest engines from z13-z15
 - z16 sized to apparently try to fix that
- Per-engine speed steps skewed for granularity of smaller environments
 - This may be addressing a need for the <200 MSUs market




z16 A01 and A02 Selected Models



- Under 200 MSUs probably want “small” machine
- 200-800 MSUs is potentially more of a toss-up

If you're considering a move, check IBM software pricing as the MLC pricing can be different for the A01 vs A02 machines.

z16 Machines <800 MSUs



Scott's Wishes/Hopes



- More processor “speed” options for the big machines
- Allow more GPs on the small machines (10 would be great!)


- Release “small” and “big” machines concurrently
 - Possibly collapse into a single model with 1xx-7xx and Axx-Zxx options

- The first two are easily doable, the third may be more “complicated” from a business / marketing / technical perspective



Expected trade-offs

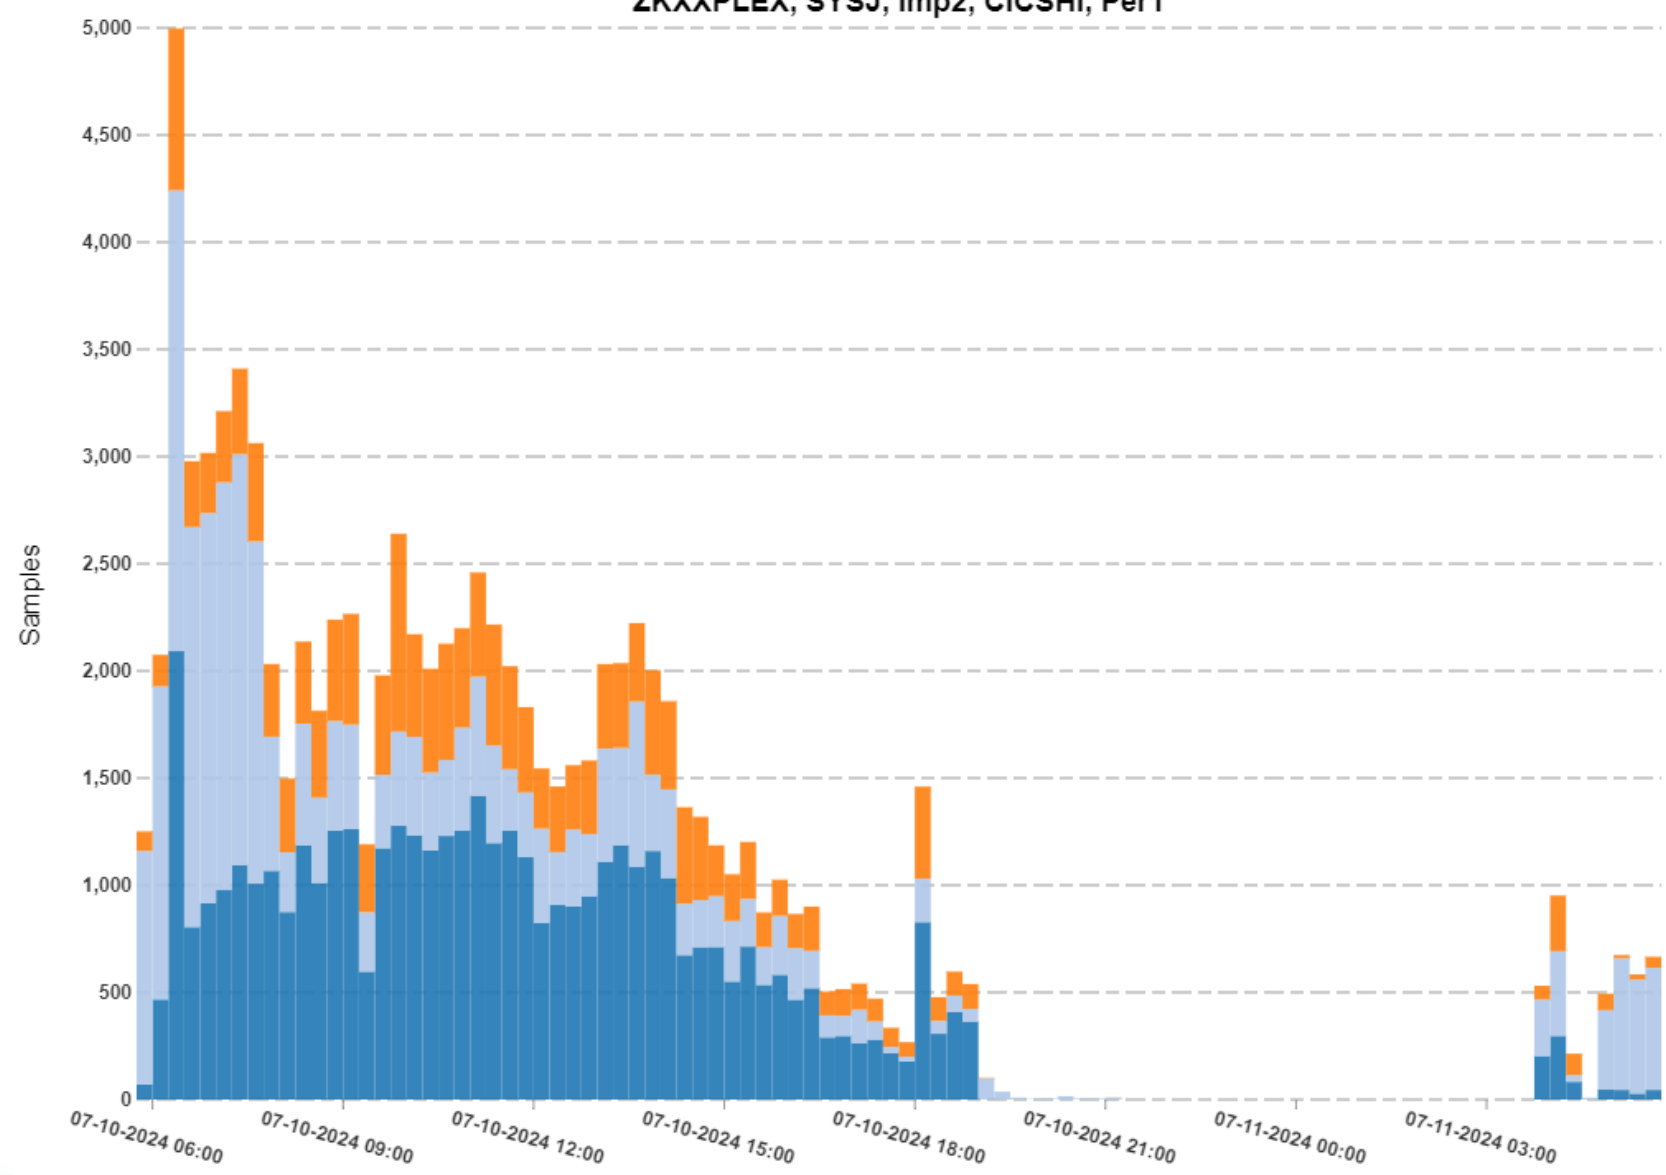
Response Time Components



- This is an exaggeration of course: in reality there's probably a lot more yellow and green!
 - And batch jobs probably won't have any black, per se
- The point is that CPU time and CPU wait may be a significant component, but it's certainly not the only thing that makes up application response time
 - And in some cases might be the minority of the time

Using & Delay Samples Breakdown

ZKXXPLEX, SYSJ, Imp2, CICS/SHI, Per1




- CPU Using
- CPU Delay
- I/O Using
- I/O Delay
- zIIP Using
- zIIP Delay
- AP Crypto Using
- AP Crypto Delay
- Contention Using
- Contention Delay
- Cap Delay
- MPL Delay
- Priv Pac
- Comm F
- XMEM
- VIO Del
- Queue
- Priv Pac
- VIO Srv
- MPL Srv

Just an example of WLM's view of a CICS workload.

In this case "CPU Using" is significant but so is CPU delay and I/O.

Using & Delay Samples Breakdown


ZKXXPLEX, SYSJ, Imp2, DDFOL, Per1



This DDF workload mostly dominated by I/O, not CPU.

Using & Delay Samples Breakdown

ZKXXPLEX, SYSJ, Imp3, BATMD, Per1



Similarly, CPU time is not the major consideration for this batch workload.

Here we even see some MPL, queue, and contention delays coming into play.

Fewer/Faster



- ± While CPU time measurements will be lower, CPU delays may be higher
- + Some workloads may perform better with faster processors
 - Primarily workloads whose elapsed time is dominated by CPU time
 - So not seeing a lot of CPU delay or time spent doing I/O
 - Likely need to be running at a high dispatching priority and/or running during a less busy time
- Some workloads may not be able to move to slower processors
 - Classic case: single-CICS region applications that are dependent on the QR TCB and are consuming a significant portion of a CP with that QR TCB
- Less important tasks may be delayed more with fewer CPs to dispatch on
- Tasks dominating a CP will dominate a larger portion of the total capacity
 - Including “spin” time for sync CF requests or zHyperLink I/O

More/Slower



- ± While CPU time measurements will be higher, CPU delays may be lower
- + Processor efficiency may improve, i.e. more net effective capacity
 - More L1 / L2 cache from having more physical CPs
- + Responsiveness of less important work likely better
 - More CPs to dispatch on = more opportunities to be dispatched
- + Workloads dominating a CP will dominate less of the total capacity
 - Especially good for “spinning” requests
- Workloads dominating a faster CP could be negatively impacted on slower
 - Bad if that workload is performance-sensitive and important to the business
 - E.G. a busy QR TCB in that important CICS region
 - Maybe not bad if the workload isn't the above
 - E.G. a batch job that runs overnight and is finishing hours before the business really needs it done

Biggest Potential Problems



- For fewer/faster CPs:
 - Are there tasks that are dominating a CP that will consume a greater percentage of your total capacity?
 - Do you need more CPs to dispatch on?
- For more/slower CPs:
 - Do you have tasks that are dominating a CP that need that faster CP?
- For both:
 - How will the effective capacity change relative to the rated capacity?



Modelling Changes




- zPCR is your tool to analyze the relative capacity difference in the machines based on your specific LPAR configuration
 - Free download from IBM
 - Easy to use
- Lets you explore the relative capacity impacts of various changes based on your specific LPAR configuration and your workload characteristics
 - Using the relative capacity change predicted by zPCR is more accurate than just looking at the “rated” capacity the machines
 - May find configurations that are more efficient per rated MSU

<https://www.ibm.com/support/pages/getting-started-zpcr-ibms-processor-capacity-reference>

zPCR Tips



- Change the name of your configurations to include the rated MSUs to make it easier to remember/compare later
- Table of MSU ratings as well as the ability to find machines of a similar rated capacity included right in the app
- Single-CP comparison can be used to predict CPU time changes on new machine
- Applying actual zIIP loadings can help refine estimates
- Provides HiperDispatch details for each configs
 - May want to use even for non-upgrade config changes



See also my SHARE (summer 2022) presentation “Planning Your Next Mainframe Upgrade”

Also available on our website: https://www.pivotor.com/library/content/Chapman_MainframeUpgradePlanning_SHARE_202208.pdf

Example scenario



- z14 605 with a couple of production LPARs, test LPAR and Sysprog test LPAR
- Want to upgrade to z16 and stay near same rated capacity due to software cost considerations (ISV restrictions can be problematic!)

Define General Purpose Partitions
Based on LSPR Data for IBM Z Processors
Study ID: Not specified
#1 z14-605 (591 MSUs)
z14 Host = 3906-M01/600 with 10 CPs: GP=5 zIIP=3 ICF=2
10 Active Partitions: GP=4 zIIP=4 ICF=2

Include	Partition Identification					Partition Configuration				Capping	
	No.	Type	Name	SCP	Workload	Mode	LCPs	Weight	Weight %	INIT	ABS
<input checked="" type="checkbox"/>	1	GP	PROD1	z/OS-2.5*	Average	SHR	5	500	50.00%	<input type="checkbox"/>	
<input checked="" type="checkbox"/>		zIIP	PROD1	z/OS-2.5*	Average	SHR	2	500	50.00%	<input type="checkbox"/>	
<input checked="" type="checkbox"/>	2	GP	PROD2	z/OS-2.5*	Average	SHR	4	300	30.00%	<input type="checkbox"/>	
<input checked="" type="checkbox"/>		zIIP	PROD2	z/OS-2.5*	Average	SHR	2	200	20.00%	<input type="checkbox"/>	
<input checked="" type="checkbox"/>	3	GP	TEST	z/OS-2.5*	Average	SHR	2	150	15.00%	<input type="checkbox"/>	
<input checked="" type="checkbox"/>		zIIP	TEST	z/OS-2.5*	Average	SHR	2	200	20.00%	<input type="checkbox"/>	
<input checked="" type="checkbox"/>	4	GP	IPO	z/OS-3.1*	Average	SHR	2	50	5.00%	<input type="checkbox"/>	
<input checked="" type="checkbox"/>		zIIP	IPO	z/OS-3.1*	Average	SHR	1	100	10.00%	<input type="checkbox"/>	

Finding "similar" options



z/OS-2.4 LSPR Data (04/04/2023)

IBM zPCR 9.6.5

LSPR Multi-Image Capacity Ratios

IBM Z General Purpose CPs (Similar CPCs)

21 CPCs in the range of 7.30 to 8.92 ITRR for "Average" are displayed

Capacity basis: 2094-701 @ 0.9440 ITRR for a typical multi-partition configuration

Capacity for z/OS on z10 and later processors is represented with HiperDispatch turned On

IBM Z Processor	Features	Flag	MSU	LSPR Workload Category				
				Low	Low-Avg	Average	Avg-High	High
3906-423	23W	=	547	8.560	7.998	7.505	7.114	6.762
3931-421	21W	=	548	8.396	7.931	7.516	7.196	6.902
3906-507	7W	=	562	8.385	8.029	7.702	7.232	6.816
3932-U05	5W	=	562	7.993	7.814	7.643	7.304	6.994
3906-424	24W	=	568	8.891	8.301	7.785	7.380	7.015
3931-422	22W	=	570	8.753	8.258	7.816	7.485	7.180
3932-T06	6W	=	582	8.321	8.114	7.917	7.552	7.218
3906-425	25W	=	588	9.220	8.604	8.064	7.645	7.266
3906-605	5W	=	591	8.727	8.405	8.107	7.620	7.187
3931-423	23W	=	592	9.110	8.585	8.117	7.773	7.457
3932-W04	4W	=	594	8.417	8.247	8.083	7.748	7.440
3932-Y03	3W	=	600	8.478	8.319	8.165	7.865	7.586
3906-426	26W	=	608	9.547	8.904	8.342	7.908	7.516
3931-424	24W	=	614	9.464	8.910	8.416	8.060	7.733
3931-506	6W	=	616	8.843	8.637	8.440	8.079	7.748
3906-703	3W	=	620	9.114	8.805	8.515	8.048	7.629
3906-427	27W	=	628	9.872	9.203	8.619	8.169	7.764
3906-508	8W	=	633	9.501	9.074	8.683	8.150	7.679
3931-425	25W	=	635	9.817	9.234	8.716	8.347	8.008
3932-V05	5W	=	641	9.107	8.903	8.708	8.323	7.970
3906-428	28W	=	647	10.195	9.501	8.895	8.429	8.010

Note:

- There's no z16 6xx within a rated +/- 10% of our z14 6xx
- z16 506 is rated at about +4%
- z16 3932 may be good?

Comparing Full Capacity



Host Capacity Comparison Summary

IBM zPCR 9.6.5

LPAR Host Configuration Capacity Comparison Report

Capacity basis: 2094-701 @ 1.000 ITRR for a shared single-partition configuration
 Capacity for z/OS on z10 and later processors is represented with HiperDispatch turned ON

LPAR Configuration			Full CPC Capacity (based on usable RCP count)					
Identity	Hardware	SMT	GP	ZAAP	zIIP	IFL	ICF	Total
#1	z14-605 (591 MSUs)	3906-M01/600: GP=5 zIIP=3 ICF=2	8.09	n/s	8.07		4.24	20.40
#2	z16-506 616 MSUs	3931-A01(Max39)/500: GP=6 zIIP=3 ICF=2	8.62	n/s	10.18		5.26	24.05
Percent Delta from "z14-605 (591 MSUs)"			+6.4%		+26.1%		+24.2%	+17.9%
#3	z16-423 (592 MSUs)	3931-A01(Max39)/400: GP=23 zIIP=3 ICF=2	8.97	n/s	9.31		4.70	22.97
Percent Delta from "z14-605 (591 MSUs)"			+10.9%		+15.3%		+10.8%	+12.6%
#4	z16-T06 (582 MSUs)	3932-A02(Max16)/T00: GP=6 zIIP=3 ICF=2	8.02	n/s	8.91		5.13	22.06
Percent Delta from "z14-605 (591 MSUs)"			-0.8%		+10.4%		+21.0%	+8.2%
#5	z16-W04 (594 MSUs)	3932-A02(Max16)/W00: GP=4 zIIP=3 ICF=2	8.02	n/s	9.00		5.19	22.21
Percent Delta from "z14-605 (591 MSUs)"			-0.9%		+11.5%		+22.5%	+8.9%

Content Control: Show Capacity Deltas

Show capacity as: Full CPC, Single-CP

Based on "z14-605 (591 MSUs)" Incremental

For significant configuration changes, such as upgrading the processor family, all capacity comparisons are subject to a +/-5% margin-of-error.

IBM does not guarantee the results from this tool. This information is provided "as is", without warranty, expressed or implied. You are responsible for the results obtained from your use of this tool.

Interested in these blue percentages

Comparing Per-CPU Capacity



Host Capacity Comparison Summary

IBM zPCR 9.6.5

LPAR Host Configuration Capacity Comparison Report

Capacity basis: 2094-701 @ 1.000 ITRR for a shared single-partition configuration
 Capacity for z/OS on z10 and later processors is represented with HiperDispatch turned ON

LPAR Configuration			Single-CP Capacity (based on usable RCP count)						
Identity	Hardware	SMT	GP	zAAP	zIIP	IFL	ICF	Total	
#1	z14-605 (591 MSUs)	3906-M01/600: GP=5 zIIP=3 ICF=2	1.62	n/s	2.69		2.12	2.04	
#2	z16-506 616 MSUs	3931-A01(Max39)/500: GP=6 zIIP=3 ICF=2	1.44	n/s	3.39		2.63	2.19	
	Percent Delta from "z14-605 (591 MSUs)"		-11.3%		+26.1%		+24.2%	+7.2%	
#3	z16-423 (592 MSUs)	3931-A01(Max39)/400: GP=23 zIIP=3 ICF=2	0.39	n/s	3.10		2.35	0.82	
	Percent Delta from "z14-605 (591 MSUs)"		-75.9%		+15.3%		+10.8%	-59.8%	
#4	z16-T06 (582 MSUs)	3932-A02(Max16)/T00: GP=6 zIIP=3 ICF=2	1.34	n/s	2.97		2.56	2.01	
	Percent Delta from "z14-605 (591 MSUs)"		-17.3%		+10.4%		+21.0%	-1.7%	
#5	z16-W04 (594 MSUs)	3932-A02(Max16)/W00: GP=4 zIIP=3 ICF=2	2.00	n/s	3.00		2.60	2.47	
	Percent Delta from "z14-605 (591 MSUs)"		+23.9%		+11.5%		+22.5%	+21.0%	

Note that 3 of our 4 options involve "slower" CPUs

Content Control

Show Capacity Deltas

Based on "z14-605 (591 MSUs)"

Incremental

Show capacity as

Full CPC

Single-CP

For significant configuration changes, such as upgrading the processor family, all capacity comparisons are subject to a +/-5% margin-of-error.

IBM does not guarantee the results from this tool. This information is provided "as is", without warranty, expressed or implied. You are responsible for the results obtained from your use of this tool.

Rated MSUs vs Effective Capacity



- The rated MSU change vs the effective capacity change is a means of understanding the software cost efficiency of the change
 - I.E. more effective capacity for a smaller MSU increase implies more work done for less software cost

Option	MSUs	MSUs Change	Effective Capacity Change
z14-605	591	n/a	n/a
z16-506	616	+4.2%	+6.4%
z16-423	592	+0.1%	+10.9%
z16-T06	582	-2.5%	-0.8%
z16-W04	594	+0.5%	-0.9%

This looks particularly intriguing!
(And a 422 might appear even better?)

Which would I choose?



- Impossible to say without understanding many more details!
 - Hardware costs of the different options
 - Software contract limits and upgrade costs (in particular ISVs)
 - IBM MLC software pricing of A01 vs A02 in this environment
 - Are there other CECs in a sysplex with this one?
 - Are there single-TCB tasks sensitive to decreased per-engine capacity?
 - How sensitive?
 - What are the current performance pain points for the organization?
 - What are the mainframe budgetary concerns?
 - What is the capacity requirements forecast over the life of the new CEC?
- Important point: **There are multiple options to consider!**
 - If your Business Partner is providing only one option, dig deeper!

Larger CEC scenario



- Z14 710 (1793 MSUs) upgrade to z16 and allow for about 10% growth

IBM zPCR 9.6.5

Partition Detail Report

Based on LSPR Data for IBM Z Processors
Study ID: Not specified

#1 z14 710 (1793 MSUs)

z14 Host = 3906-M01/700 with 19 CPs: GP=10 zIIP=6 ICF=3
12 Active Partitions: GP=5 zIIP=4 ICF=3

Capacity basis: 2094-701 @ 1.000 ITRR for a shared single-partition configuration
Capacity for z/OS on z10 and later processors is represented with HiperDispatch turned ON

Include ✓	Partition Identification					Partition Configuration									
	No.	Type	Name	SCP	Assigned Workload	Mode	Logical CPs	Weight	Weight Percent	Capping		SMT		Capacity	
										INIT	ABS	✓	Benefit	Minimum	Maximum
<input checked="" type="checkbox"/>	1	GP	SYA1	z/OS-2.5*	Average	SHR	6	350	35.00%	<input type="checkbox"/>				8.85	15.17
<input checked="" type="checkbox"/>		zIIP	SYA1	z/OS-2.5*	Average	SHR	3	100	25.00%	<input type="checkbox"/>		<input type="checkbox"/>		3.82	7.64
<input checked="" type="checkbox"/>	2	GP	SYA2	z/OS-2.5*	Average	SHR	4	200	20.00%	<input type="checkbox"/>				5.21	10.42
<input checked="" type="checkbox"/>		zIIP	SYA2	z/OS-2.5*	Average	SHR	2	100	25.00%	<input type="checkbox"/>		<input type="checkbox"/>		3.90	5.20
<input checked="" type="checkbox"/>	3	GP	SYB1	z/OS-2.5*	Average	SHR	6	300	30.00%	<input type="checkbox"/>				7.59	15.17
<input checked="" type="checkbox"/>		zIIP	SYB1	z/OS-2.5*	Average	SHR	3	100	25.00%	<input type="checkbox"/>		<input type="checkbox"/>		3.82	7.64
<input checked="" type="checkbox"/>	4	GP	SYB2	z/OS-2.5*	Average	SHR	2	100	10.00%	<input type="checkbox"/>				2.54	5.08
<input checked="" type="checkbox"/>		zIIP	SYB2	z/OS-2.5*	Average	SHR	2	100	25.00%	<input type="checkbox"/>		<input type="checkbox"/>		3.98	5.30
<input checked="" type="checkbox"/>	5	GP	SYT1	z/OS-2.5*	Average	SHR	2	50	5.00%	<input type="checkbox"/>				1.36	5.43
<input checked="" type="checkbox"/>	6	ICF	ICF-01	CFCC	CFCC	DED	1	n/a						2.26	2.26
<input checked="" type="checkbox"/>	7	ICF	ICF-02	CFCC	CFCC	DED	1	n/a						2.26	2.26
<input checked="" type="checkbox"/>	8	ICF	ICF-03	CFCC	CFCC	DED	1	n/a						2.26	2.26

Purple is warning us about the LCP count vs. the weight. NBD here.

Full Capacity Comparison



Host Capacity Comparison Summary

LPAR Host Configuration Capacity Comparison Report
 Capacity basis: 2094-701 @ 1.000 ITRR for a shared single-partition configuration
 Capacity for z/OS on z10 and later processors is represented with HiperDispatch turned ON

LPAR Configuration			Full Capacity (based on usable RCP count)					
Identity	Hardware	SMT	GP	zAAP	zIIP	IFL	ICF	Total
#1	z14 710 (1793 MSUs)	3906-M01/700: GP=10 zIIP=6 ICF=3	25.55	n/s	15.51		6.78	47.84
#2	z16 708 (1866 MSUs)	3931-A01(Max39)/700: GP=8 zIIP=6 ICF=3	26.11	n/s	19.74		8.45	54.30
	Percent Delta from "z14 710 (1793 MSUs)"		+2.2%		+27.2%		+24.6%	+13.5%
#3	z16 709 (2061 MSUs)	3931-A01(Max39)/700: GP=9 zIIP=6 ICF=3	29.14	n/s	19.67		8.40	57.20
	Percent Delta from "z14 710 (1793 MSUs)"		+14.0%		+26.8%		+23.8%	+19.6%
#4	z16-613 (1894 MSUs)	3931-A01(Max39)/600: GP=13 zIIP=6 ICF=3	27.22	n/s	19.25		8.17	54.65
	Percent Delta from "z14 710 (1793 MSUs)"		+6.5%		+24.1%		+20.5%	+14.2%
#5	z16-614 (2009 MSUs)	3931-A01(Max39)/600: GP=14 zIIP=6 ICF=3	29.20	n/s	19.25		8.12	56.57
	Percent Delta from "z14 710 (1793 MSUs)"		+14.3%		+24.1%		+19.7%	+18.2%
#6	z16-523 (1904 MSUs)	3931-A01(Max39)/500: GP=23 zIIP=6 ICF=3	27.92	n/s	17.92		7.62	53.45
	Percent Delta from "z14 710 (1793 MSUs)"		+9.3%		+15.5%		+12.4%	+11.7%

Content Control: Show Capacity Deltas

Show capacity as: Full CPC, Single-CP

For significant configuration changes, such as upgrading the processor family, all capacity comparisons are subject to a +/-5% margin-of-error.

IBM does not guarantee the results from this tool. This information is provided "as is", without warranty, expressed or implied. You are responsible for the results obtained from your use of this tool.

Capacity target between 708 and 709

523 seems to hit the capacity target almost exactly.

Per-CPU Capacity



Host Capacity Comparison Summary IBM zPCR 9.6.5

LPAR Host Configuration Capacity Comparison Report
 Capacity basis: 2094-701 @ 1.000 ITRR for a shared single-partition configuration
 Capacity for z/OS on z10 and later processors is represented with HiperDispatch turned ON

LPAR Configuration				Single-CP Capacity (based on usable RCP count)					
Identity	Hardware	SMT	GP	zAAP	zIIP	IFL	ICF	Total	
#1	z14 710 (1793 MSUs)	3906-M01/700: GP=10 zIIP=6 ICF=3	2.55	n/s	2.59		2.26	2.52	
#2	z16 708 (1866 MSUs)	3931-A01(Max39)/700: GP=8 zIIP=6 ICF=3	3.26	n/s	3.29		2.82	3.19	
	Percent Delta from "z14 710 (1793 MSUs)"		+27.8%		+27.2%		+24.6%	+26.8%	
#3	z16 709 (2061 MSUs)	3931-A01(Max39)/700: GP=9 zIIP=6 ICF=3	3.24	n/s	3.28		2.80	3.18	
	Percent Delta from "z14 710 (1793 MSUs)"		+26.7%		+26.8%		+23.8%	+26.2%	
#4	z16-613 (1894 MSUs)	3931-A01(Max39)/600: GP=13 zIIP=6 ICF=3	2.09	n/s	3.21		2.72	2.48	
	Percent Delta from "z14 710 (1793 MSUs)"		-18.0%		+24.1%		+20.5%	-1.4%	
#5	z16-614 (2009 MSUs)	3931-A01(Max39)/600: GP=14 zIIP=6 ICF=3	2.09	n/s	3.21		2.71	2.46	
	Percent Delta from "z14 710 (1793 MSUs)"		-18.4%		+24.1%		+19.7%	-2.3%	
#6	z16-523 (1904 MSUs)	3931-A01(Max39)/500: GP=23 zIIP=6 ICF=3	1.21	n/s	2.99		2.54	1.67	
	Percent Delta from "z14 710 (1793 MSUs)"		-52.5%		+15.5%		+12.4%	-33.7%	

Content Control

Show Capacity Deltas

Based on "z14 710 (1793 MSUs)"

Incremental

Show capacity as

Full CPC


Single-CP

For significant configuration changes, such as upgrading the processor family, all capacity comparisons are subject to a +/-5% margin-of-error.

IBM does not guarantee the results from this tool. This information is provided "as is", without warranty, expressed or implied. You are responsible for the results obtained from your use of this tool.

523 does have about half-speed engines relative to current.

Rated MSUs vs Effective Capacity



- Again, slower engines seem to be price-performance leaders
- Decision here should take into consideration timing & confidence of the capacity forecast
 - Also consider the next step if the forecast is wrong

Option	MSUs	MSUs Change	Effective Capacity Change
z14-710	1793	n/a	n/a
z16-708	1866	+4.1%	+2.2%
z16-709	2061	+14.9%	+14.0%
z16-613	1894	+5.6%	+6.5%
z16-614	2009	+12.0%	+14.3%
z16-523	1904	+6.2%	+9.3%



Wrap-up

Know your limits



- What's your budgetary limit for hardware?
- How do your IBM software costs change on a new machine?
- Do you have ISV contracts with big upgrade charges if you pass a limit?
- Are there workloads that are dependent on the per-CPU speed?
- How much leeway do you have in your batch window(s)?

Think outside the box



- More/Slower is often better for:
 - Overall system efficiency
 - Software cost efficiency
 - More/slower is more better with more LPARs
- Highly transactional workloads may not suffer noticeably from slower CPs
 - But beware the QR TCB, especially for old CICS applications
- Consider the “small” machines if your machine is in that < 500 MSUs range



Don't forget
Your
Session
Evaluations!!

Questions?



Don't forget
Your
Session
Evaluations!!



Don't forget
Your
Session
Evaluations!!

Your feedback is important!

Submit a session evaluation for each session you attend:

www.share.org/evaluation

