Planning Your Next Mainframe Processor Upgrade in 2026

Scott Chapman
Enterprise Performance Strategies
Scott.Chapman@epstrategies.com

SHARE Orlando
February 2026
Session TECH_135

Session 44850

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:

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

Many customers only replace their mainframe every 3-5 years, so properly planning for what machine to go to next is not something they have a lot of practice at. Many simply rely on their vendor to make a good suggestion. But the wise customer will consider multiple factors before making a final decision.

In this session Scott Chapman will share the insight he has garnered from years of planning processor upgrades. He'll explain the basics of running zPCR, why you want to do so, and how you can use the output. Preparing for processor speed changes will be discussed. Finally, some thoughts will be provided about understanding the performance changes after the upgrade.

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

z/OS Performance workshops available

During these workshops you will be analyzing your own data!

- Essential z/OS Performance Tuning
  - March 30 – April 3, 2026 (4 days, excl Wednesday the 1st)
- WLM Performance and Re-evaluating Goals
  - June 22 – 26, 2026 (4 days, excl Wednesday the 24th)
- Parallel Sysplex and z/OS Performance Tuning
  - May 12-13 2026
- 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!

EPS presentations this week

What	Who	When	Where
z/OS Performance Management If You Only Have 20 Minutes A Day	Scott Chapman	Mon 9:45	Salon 14
PSP: z/OS Performance Tuning - Some Top Things You May Not Know	Peter Enrico Scott Chapman		
Planning Your Next Mainframe Processor Upgrade in 2026	Scott Chapman	Tue 15:45	Salon 15
Processor MSU Consumption Analysis	Peter Enrico	Wed 13:15	Salon 14
Can We All Share Fairly? Detection and Remediation of inter-LPAR Performance Impacts	Scott Chapman	Wed 14:30	Salon 14
Standard z/OS Measurements When Monitoring Transactions	Peter Enrico	Thu 13:15	Salon 19

Agenda

- Why I want to talk about this
- Finding your options
- Evaluating your options
- Confirming your success

Why should we talk about this?

This is not a "common" process

- Most customers only upgrade every 3-5 years so they don't have a lot of practice at this
  - Larger shops and shops with more than 1 machine may do this more
- These are not small transactions: it's worth spending some effort on them
  - The first offered solution may be the best one, but understand any trade-offs
- Your business partner will helpfully suggest an upgrade path for you
  - But their recommendations may not be entirely optimal for you
  - I'm not saying anybody is specifically and intentionally leading people astray, but…

We've seen multiple customers in unfortunate situations because they didn't fully vet the business partner's recommended configuration

Three Cost Concerns

- Hardware
  - “Easy” and obvious one
- IBM Software
  - Upgrade usually means new software discount tiers, but understand the specific impact
  - If you’re considering switching between “EC” and “BC” class machines that can also impact the costs involved
- ISV Software
  - If you’re adding capacity, this will likely be the most important one to look at
  - I’ve seen plenty of configurations changed because of ISV software costs
  - Get a good estimate of this before agreeing to anything

Understanding the software costs is a critical point in the evaluation!

Three Configuration Concerns

- Processor Configuration
  - More/Slower vs. Fewer/Faster is generally the big question (much more shortly…)
- I/O Configuration
  - Carrying forward I/O channels “just because” may not always be smart
    - Make sure you have a need for the I/O channels that you’re ordering
    - But do order what you will plausibly need during the life of the machine!
- Plan-ahead Details
  - What is your next capacity upgrade step?
  - How many drawers do you really need?

For z17: pre-plan for Spyre cards!

Many reconfigurations can be done non-disruptively, but in some cases some reconfigurations may require an outage, so talk about future plans now!

Finding your options

How much do you need?

- Over the life of the new machine do you expect your processing requirements to:
  - Increase?
  - Stay about the same?
  - Decrease?
- For example:
  - We expect to process 20% more accounts over the next 3 years
  - We don’t expect any significant changes in processing requirements
  - We’ve been told 50% of the work will leave the mainframe over the next 5 years

What are your financial constraints?

- Do you have a requirement to not increase your spend?
- Do you have a need to reduce your overall spend?
- Do you have software contract limitations which will limit how much capacity you can add?
  - E.G. Software upgrade payments that are more than the hardware cost
  - I've frequently seen software cost issues result in non-optimal capacity decisions

Figure out about how much you need

- At this point you're only interested in rough sizing based on your currently installed machine
- So if you need 20% growth and you have a 300 MSUs now: 360 MSUs
  - This is just a high-level rough estimate at this point!
  - Your predictions about the future probably have significant margins of error
- "Stay the same" targets are have more inherent risk
  - If you get it wrong you'll feel it now instead of some time in the future
  - If you do everything perfectly, IBM says there's still a +/-5% margin of error when comparing across machine generations
  - If missing by 5% is going to cause issues, you probably should be thinking about adding capacity!

CEC Physical Machine CP Busy% by CEC Serial Number

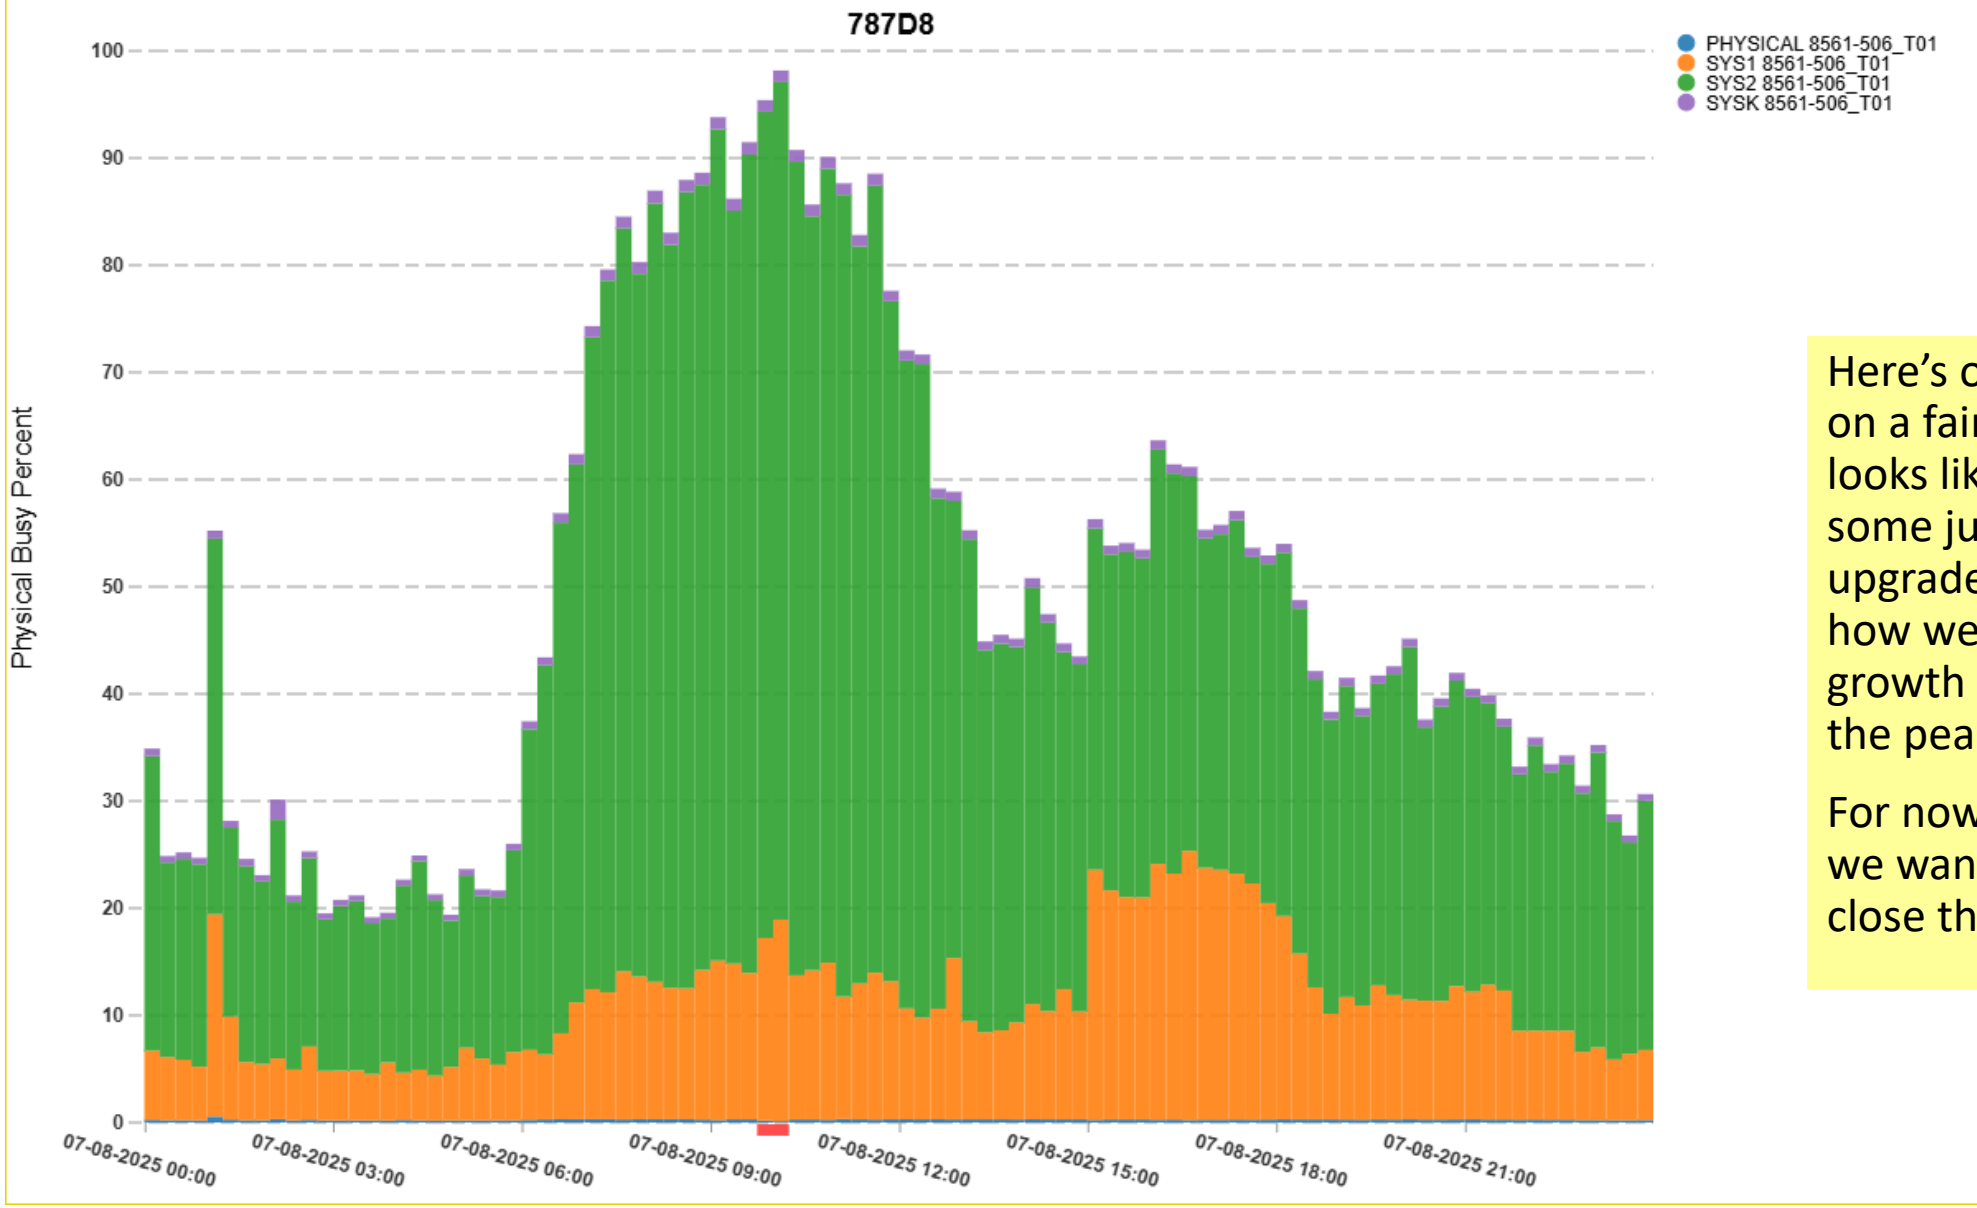

Here's our example z15 on a fairly typical day. It looks like we might have some justification for an upgrade, depending on how we feel about the growth (or not) within the peak.

For now, we'll assume we want/need to keep it close the existing size.

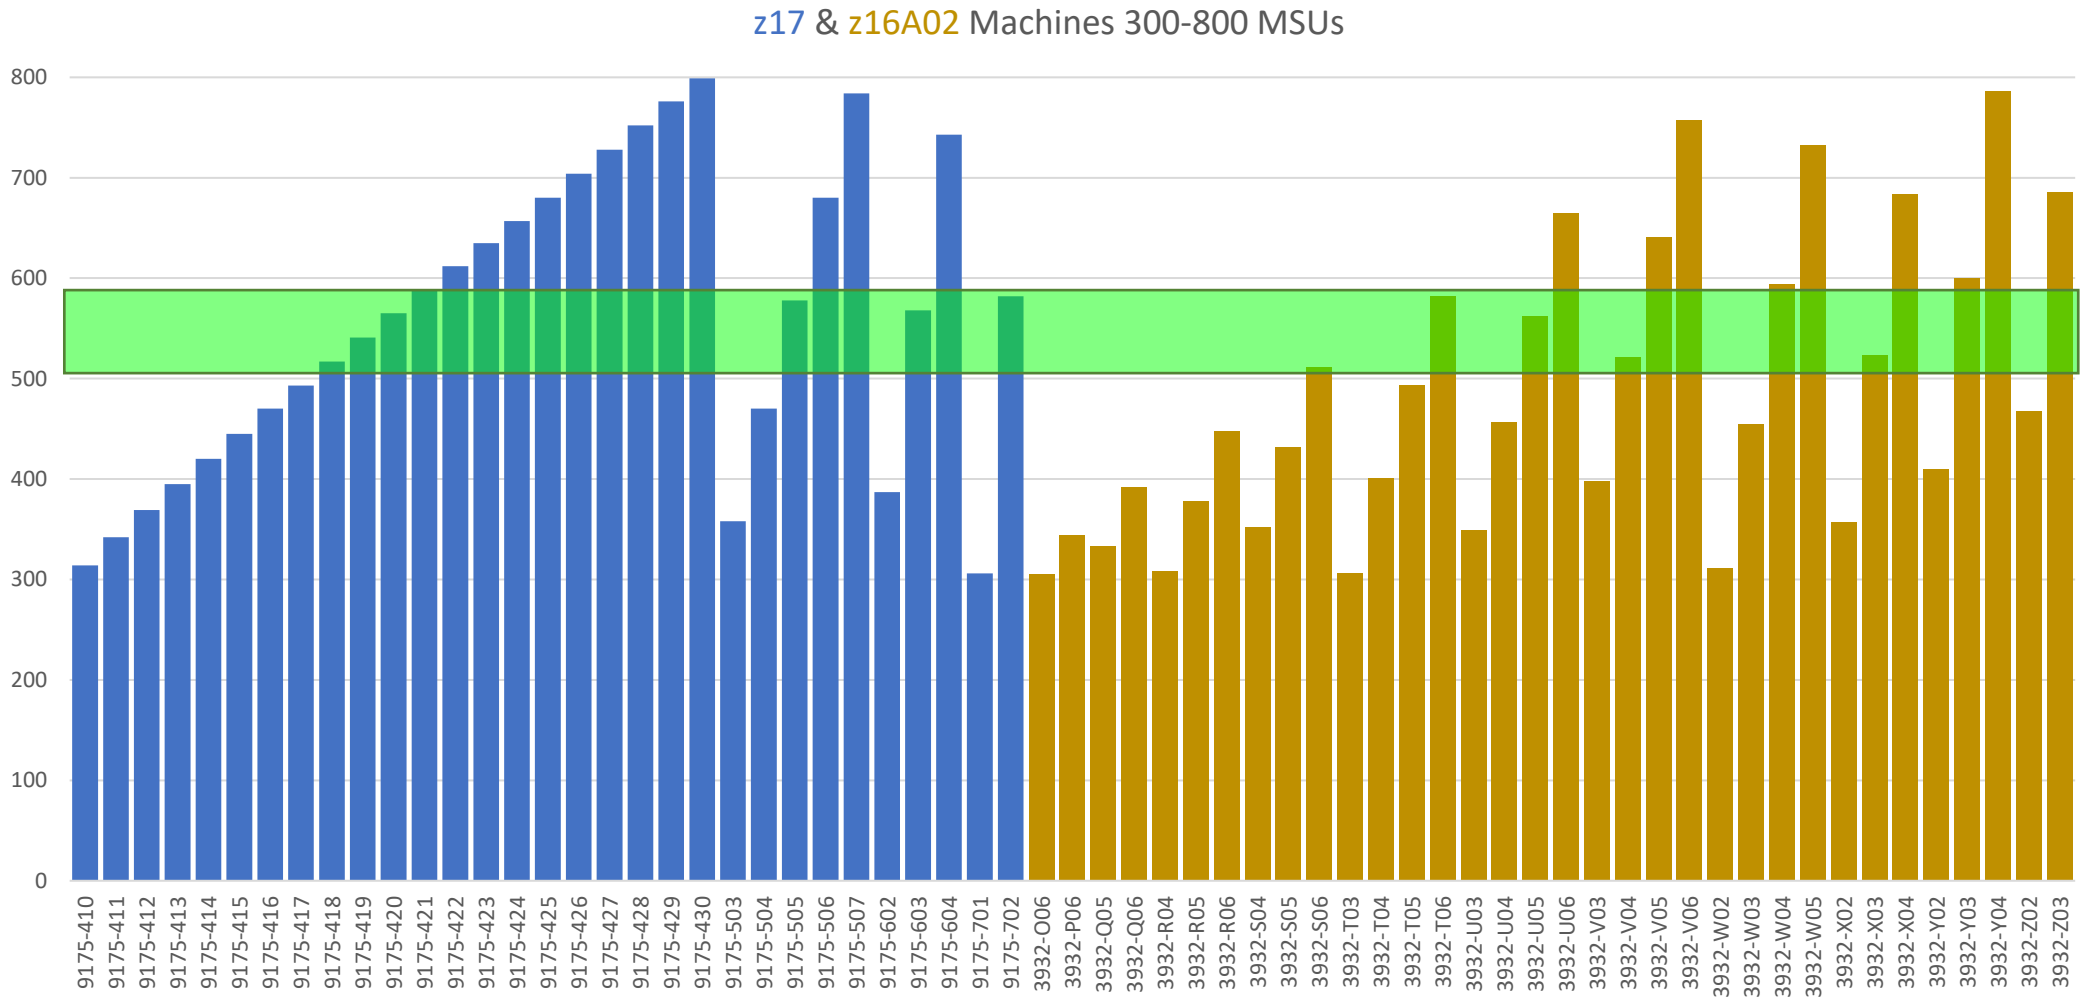

There may (or may not) be several potential options. Let's pretend we're coming from a z15-506 and want to stay close to that 512 MSUs on z17 or z16A02 machine.

Detailed Options

- The gaps between the capacity levels on the large machines is forcing some customers into uncomfortable
- Here, the closest z17 ME1 is a 418
  - People will be concerned about going to a slower CP
- A 505 adds 66 MSUs
  - Could be a concern for full-capacity SW contracts
- The 3932-S06 is almost an exact match
  - Going down to a business class z16 may have different concerns, especially if you have a multiple-CEC sysplex running under R4HA

Processor	#CP	PCI	MSU
9175-417	17	4004	493
9175-418	18	4199	517
9175-419	19	4393	541
9175-504	4	3815	470
9175-505	5	4698	578
9175-506	6	5557	680
9175-603	3	4614	568
9175-604	4	6065	743
9175-702	2	4731	582
3932-S06	6	4118	511
3932-T05	5	3976	493
3932-T06	6	4695	582
3932-U05	5	4532	562
3932-U06	6	5351	665
3932-V04	4	4202	521
3932-V05	5	5164	641
3932-X03	3	4212	523

Broad Generalities

- More/slower engines generally better for system efficiency
  - This becomes more true with more LPARs and more concurrent work
    - CPU time will go up, but CPU wait time will go down
  - Do have to be careful about important workloads that need a faster CP though
    - E.G. Older CICS applications that are dependent on the QR TCB
- More/slower often results in better effective capacity utilization per MSU
  - E.G. spend less on software to get the same or more relative capacity
- Fewer than 3 engines troubling unless the machine is dominated by 1 LPAR
  - You can run on a 1- or 2-way but you will likely have more sporadic performance
- "Business class" machines may impact your software pricing
  - May be an advantage, or could be a disadvantage depending on the details
  - Do have to be concerned about the limited number of available CPs though

What if you have to add capacity?

- Business changes: what happens if you have to add capacity before your next machine replacement? What's the next step?
- Easy, low-risk in-machine upgrades:
  - Increasing the CPU count
  - Increasing the CPU speed
- Riskier in-machine upgrades:
  - Go to fewer, but faster CPs
  - Go to more, slower CPs
  - These are doable, but you need to evaluate more closely (like we're about to do for the purchase itself)

Processor	#CP	PCI	MSU
9175-417	17	4004	493
9175-418	18	4199	517
9175-419	19	4393	541
9175-504	4	3815	470
9175-505	5	4698	578
9175-506	6	5557	680
9175-603	3	4614	568
9175-604	4	6065	743
9175-702	2	4731	582
3932-S06	6	4118	511
3932-T05	5	3976	493
3932-T06	6	4695	582
3932-U05	5	4532	562
3932-U06	6	5351	665
3932-V04	4	4202	521
3932-V05	5	5164	641
3932-X03	3	4212	523

Don’t forget about CPU count limits

- z16-A02 can only have 6 GPs
- Z17-ME1 can have up to 43 sub-capacity CPs
  - Up from 39 in the z16 and 34 in the z15
- But specific limits in your configuration may be lower due to
  - Other characterized processors (zIIPs, ICFs, IFLs)
  - Specific “MaxN” machine ordered
- My pleas to IBM:
  - Allow more than 6 GPs on the “BC” z17 (ME2?)
  - Add intermediate speed steps on the next “EC” machine

Evaluating your options

So how do we evaluate those 4 options?

- MSU/MIPS ratings are only gross generalizations of capacity
  - Based on specific tested configurations with specific test workloads
  - Your configuration and workload are different!
- zPCR is your tool to analyze the relative capacity difference in the machines based on your specific LPAR configuration
  - Free download from IBM
  - Relatively easy to use
  - Let's you explore the relative capacity impacts of various changes

https://www.ibm.com/support/pages/getting-started-ibm-z-processor-capacity-reference

Importance of SMF 113 Data

- For the most accurate planning purposes, zPCR needs to understand how the LPARs are utilizing the processor "nest"
  - Basically: how effective are the processor caches for the workload
- This data is recorded by HIS in the SMF 113 data
- If you have not already done so, enable HIS and the SMF 113 data!
  - Note that you don't normally need HIS profiler (aka sampling) output
  - https://www.ibm.com/docs/en/zos/2.5.0?topic=operations-setting-up-hardware-event-data-collection
- Workload intensity is set from combination of:
  - L1MP – Level 1 Misses Per 100 instructions
  - RNI – Relative Nest Intensity
  - Examine these values from key time periods

L1MP	RNI	Workload Hint
<3	>= 0.75 < 0.75	AVERAGE LOW
3 to 6	>1.0 0.6 to 1.0 < 0.6	HIGH AVERAGE LOW
>6	>= 0.75 < 0.75	HIGH AVERAGE

Processor Caches - CP CPU Key Measurements

SMF 113

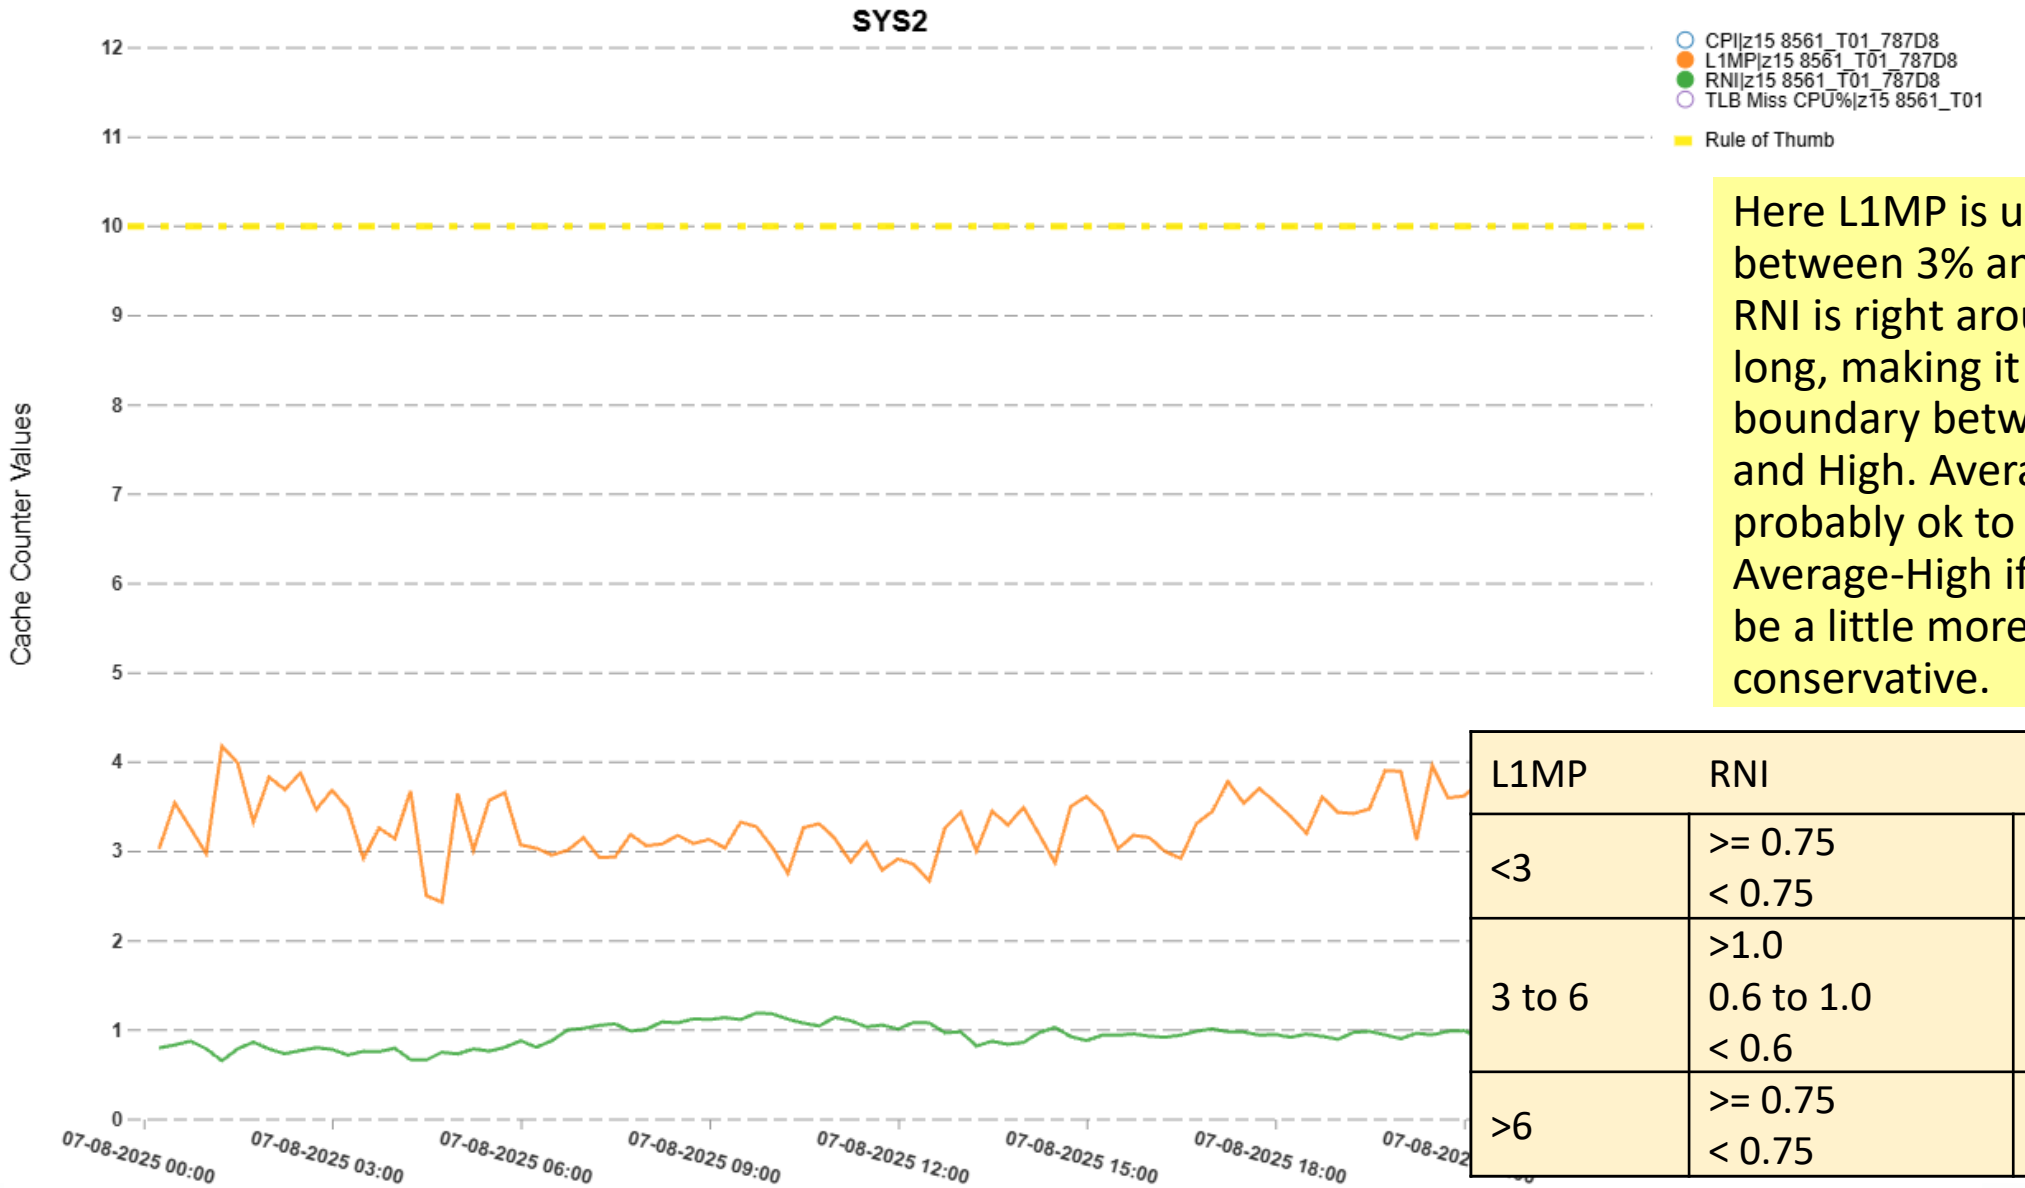

Here L1MP is usually between 3% and 4% and RNI is right around 1 all day long, making it right on the boundary between Average and High. Average is probably ok to use here. Average-High if you want to be a little more conservative.

L1MP	RNI	Workload Hint
<3	>= 0.75 < 0.75	AVERAGE LOW
3 to 6	>1.0 0.6 to 1.0 < 0.6	HIGH AVERAGE LOW
>6	>= 0.75 < 0.75	HIGH AVERAGE

Processor Caches - CP CPU Key Measurements

SMF 113

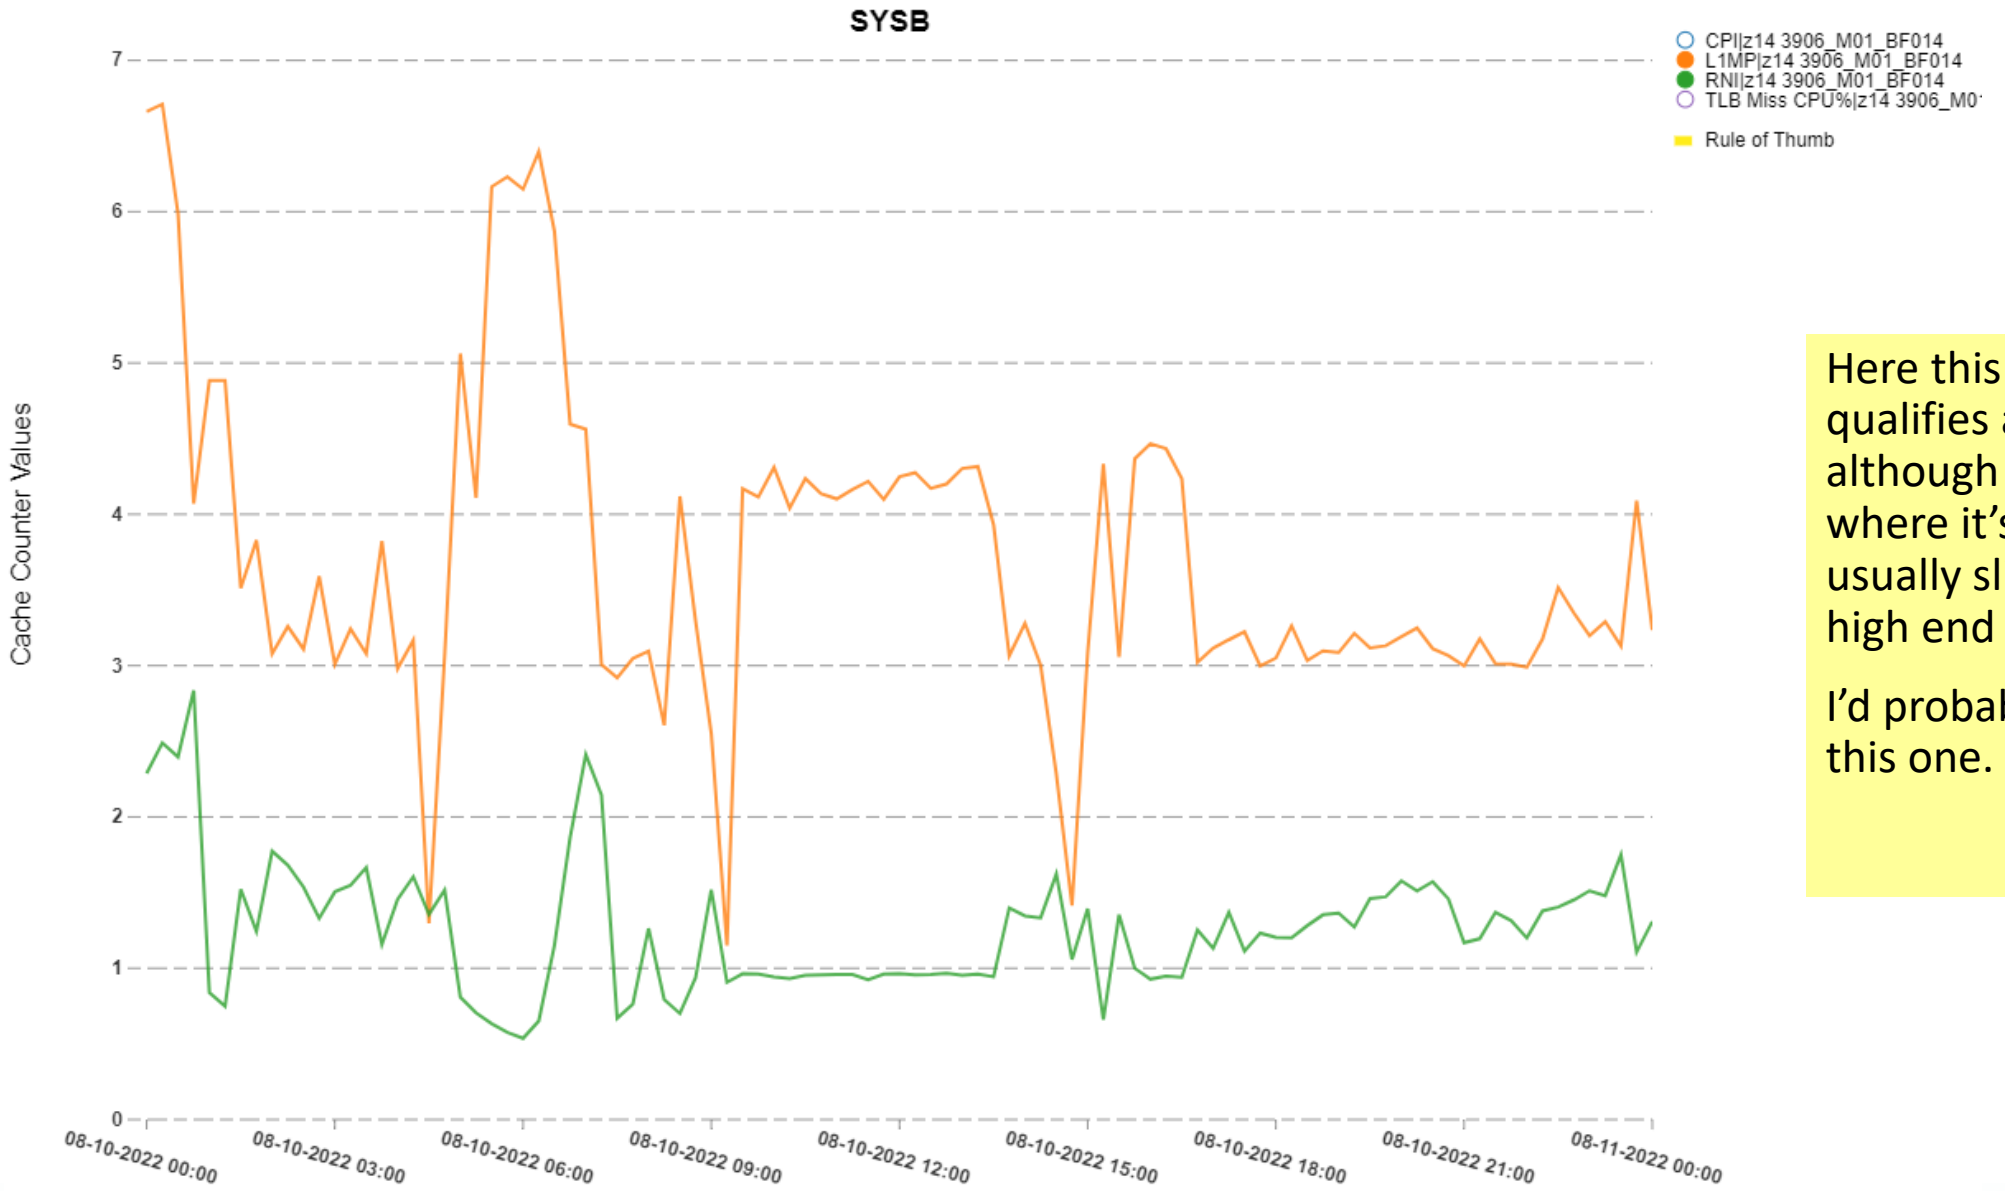

Here this system usually qualifies as “high”, although during periods where it’s busiest it usually slips back to the high end of average.

I’d probably use high for this one.

Configure your baseline in zPCR

- You can import data out of your SMF data by using CP3KEXTR to generate an EDF file, or you can manually define them
  - https://www.ibm.com/support/pages/zos-data-extraction-program-cp3kextr-zpcr-and-zbna
  - EDF input does populate some detailed numbers automatically
  - Choose your interval wisely
- Frankly, doing it manually is pretty quick, but you will need:
  - Number of GPs and zIIPs online to each LPAR
  - LPAR weights for both GPs and zIIPs
  - Workload assignment (previous two slides) for each LPAR
  - Ideally would also be nice to have an estimate of:
    - SMT benefit if using
    - zIIP loading (if they're not as busy as your GPs)
      - Default 100% loading is probably overly pessimistic for most configurations, but "safe"

Partition Detail Report

Edit Graph Documentation

IBM zPCR 9.7.5

Partition Detail Report

Based on LSPR Data for IBM Z Processors

Study ID: Not specified

#1 z15-506 512 MSUs

z15 Host = 8561-T01(Max34)/500 with 13 CPs: GP=6 zIIP=6 ICF=1

6 Active Partitions: GP=3 zIIP=2 ICF=1

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									
									Capping		SMT		Capacity	
	No.	Type	Name	SCP	Assigned Workload	Mode	Logical CPs	Weight	Weight Percent	INIT	ABS	√	Benefit	Minimum	Maximum
☑	1	GP	SYS1	z/OS-3.1	Average	SHR	2	200	20.00%	☐				1.45	2.42
☑		zIIP	SYS1	z/OS-3.1	Average	SHR	2	200	20.00%	☐		☐		3.73	6.22
☑	2	GP	SYS2	z/OS-3.1	Average	SHR	6	750	75.00%	☐				5.25	7.00
☑		zIIP	SYS2	z/OS-3.1	Average	SHR	6	800	80.00%	☐		☐		13.99	17.49
☑	3	GP	SYSK	z/OS-3.1	Average	SHR	2	50	5.00%	☐				0.38	2.50
☑	4	ICF	CF01	CFCC	CFCC	DED	1	n/a						2.81	2.81

Table View Controls

Display zAAP/zIIP/IFL/ICF Associated Partitions

(●) With Parent GP ( ) Separate by Pool

Show	GP Pool	Specialty Pools	
(●) All Partitions	☑ GP	☐ zAAP	☑ zIIP
( ) Includes Only		☐ IFL	☑ ICF

Capacity Summary by Pool

CP Pool	Real CPs	LPs	DED LCPs	SHR LCPs	SHR LCP:RCP	Sum of Weights	SMT Benefit	Capacity Totals
GP	6	3		10	1.667	1,000		7.08
zIIP	6	2		8	1.333	1,000		17.72
IFL								
ICF	1	1	1					2.81
Totals	13	6	1	18				27.61

Host Summary | SMT Benefit | LCP Alternatives | zAAP/zIIP Loading | HiperDispatch

For significant configuration changes, such as upgrading the processor family, all capacity comparisons are subject to a +/-5% margin-of-error.
For configurations estimating SMT benefits (10-25%), the margin-of-error is +/-10% for
the LPAR workloads using SMT. For higher estimates, the margin-of-error would be greater.

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.

Input fields have white background; Single-click a "selection field" for drop-down list; Double click a "key-in field" to open.

In the end should have your LPARs defined for the "from" machine.

Either get zIIP loading from EDF, change it manually, or just let it default to the pessimistic 100%.

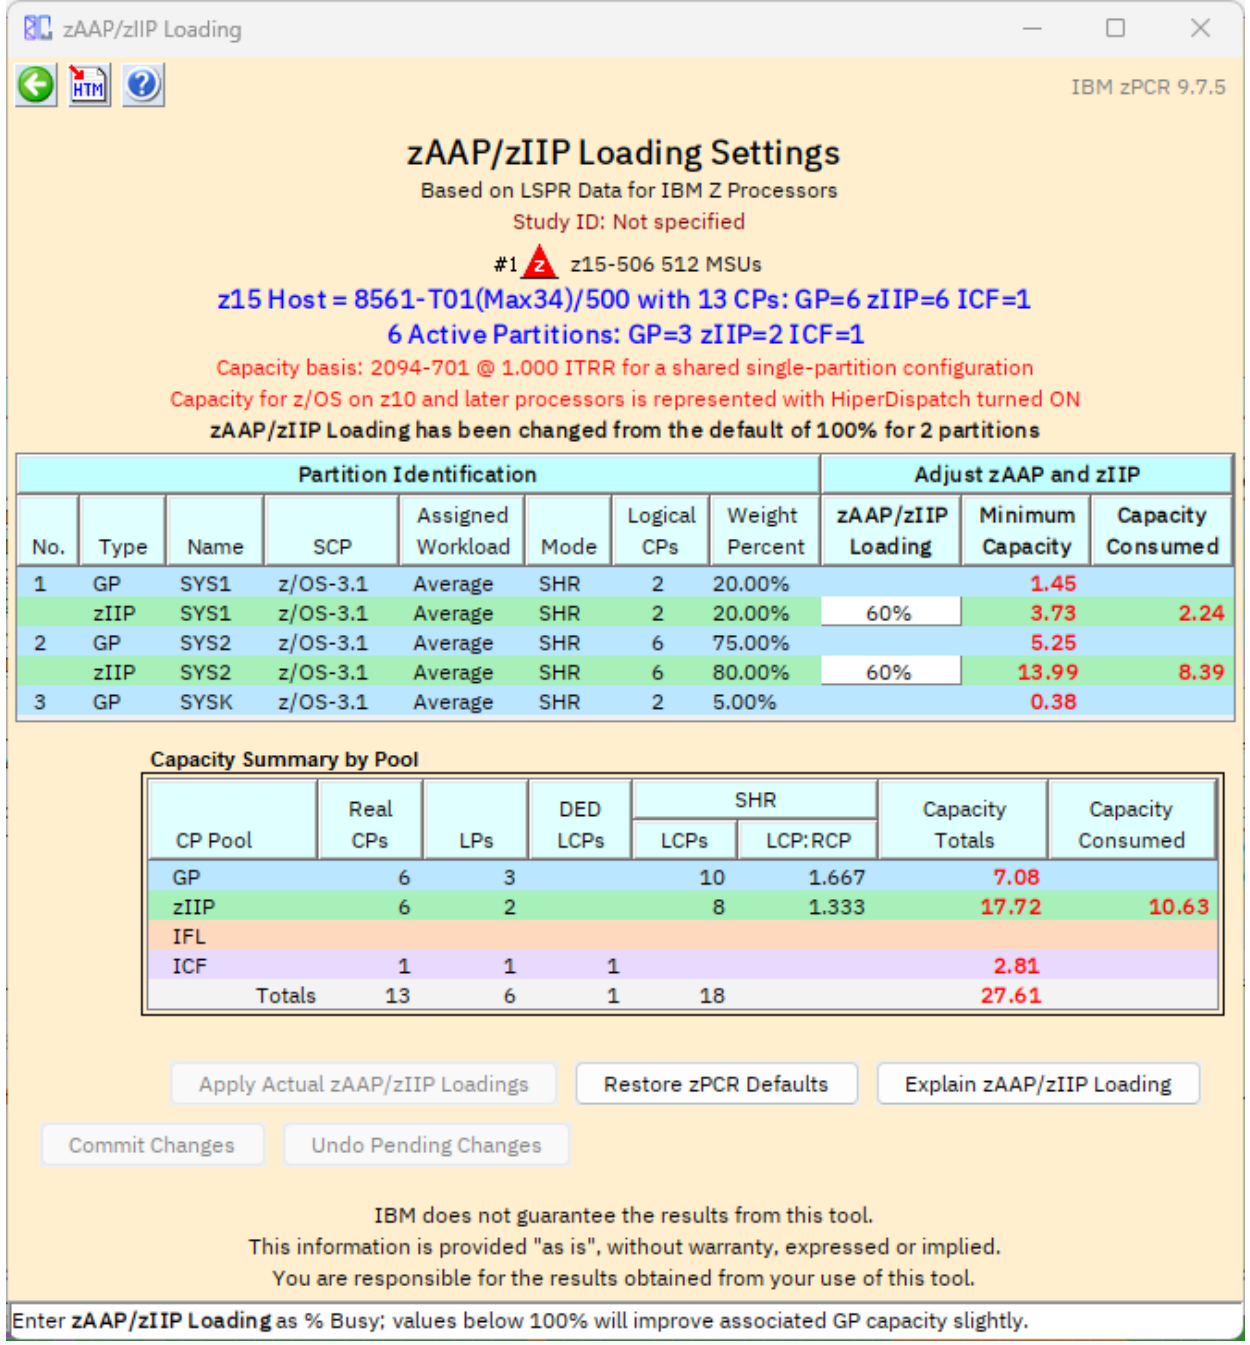

I did manually adjust it down here. This should be based on the relative logical utilization of the zIIPs and GPs. Frankly, the impact is small enough that swagging a relatively high approximation should be fine.

Control Panel [C:\...z15-z17Upgrade2.zpcr]

File Edit CPcalculator Registration Documentation Help

IBM zPCR 9.7.5

Capacity Planning Control Panel

Study ID:

Double click on a tree branch below to access the relevant windows

Reference-CPU

2094-701 @ 1.000 ITRR (SI); 0.9440 ITRR (MI)

LSPR Multi-Image Processor Table

IBM Z General Purpose CPs

IBM Z IFL CPs

IBM LinuxONE CPs

LPAR Configurations

#1 z15-506 512 MSUs
z15-T01(Max34) 8561-506 I=6 C=1

#2 z17-505 578 MSUs
z17-ME1(Max43) 9175-505 I=6 C=1

#3 z17-603 568 MSUs
z17-ME1(Max43) 9175-603 I=6 C=1

#4 z17 418 517 MSUs
z17-ME1(Max43) 9175-418 I=6 C=1

IBM z17

Manage Compare Copy & Move Partitions

QuickStart Guide

#1	z15-506 512 MSUs z15/500 LPAR Host: 8561-T01(Max34)/500					
Pool CP Type	#1 GP	#2 zAAP	#3 zIIP	#4 IFL	#5 ICF	CPC Total
RCPs	6	0	6	0	1	13
Partitions	3	0	2	0	1	6
LCPs	10	0	8	0	1	19
Capacity	7.08	n/a	17.72	n/a	2.81	27.61

LPAR Host [C:\...z15-z17Upgrade2.zpcr]

IBM zPCR 9.7.5

LPAR Host Processor

Study ID: Not specified

#2 z17-505 578 MSUs

Description: Cloned from z15-506 512 MSUs

Select Brand

IBM Z LinuxONE

Family: z17 (9175)

Speed Class: z17/500

Model: 9175-ME1(Max43)/500

Configure Real CP Types

GP	zIIP	IFL	ICF
5	6	0	1

Of 43 available CPs, 12 have been configured

define a z17 (9175) config

After the first machine, you can use the copy button to clone the configuration and just change the host machine to whatever target you want. Remember to adjust the number of processors online to the LPARs if the new machine has a different number of processors.

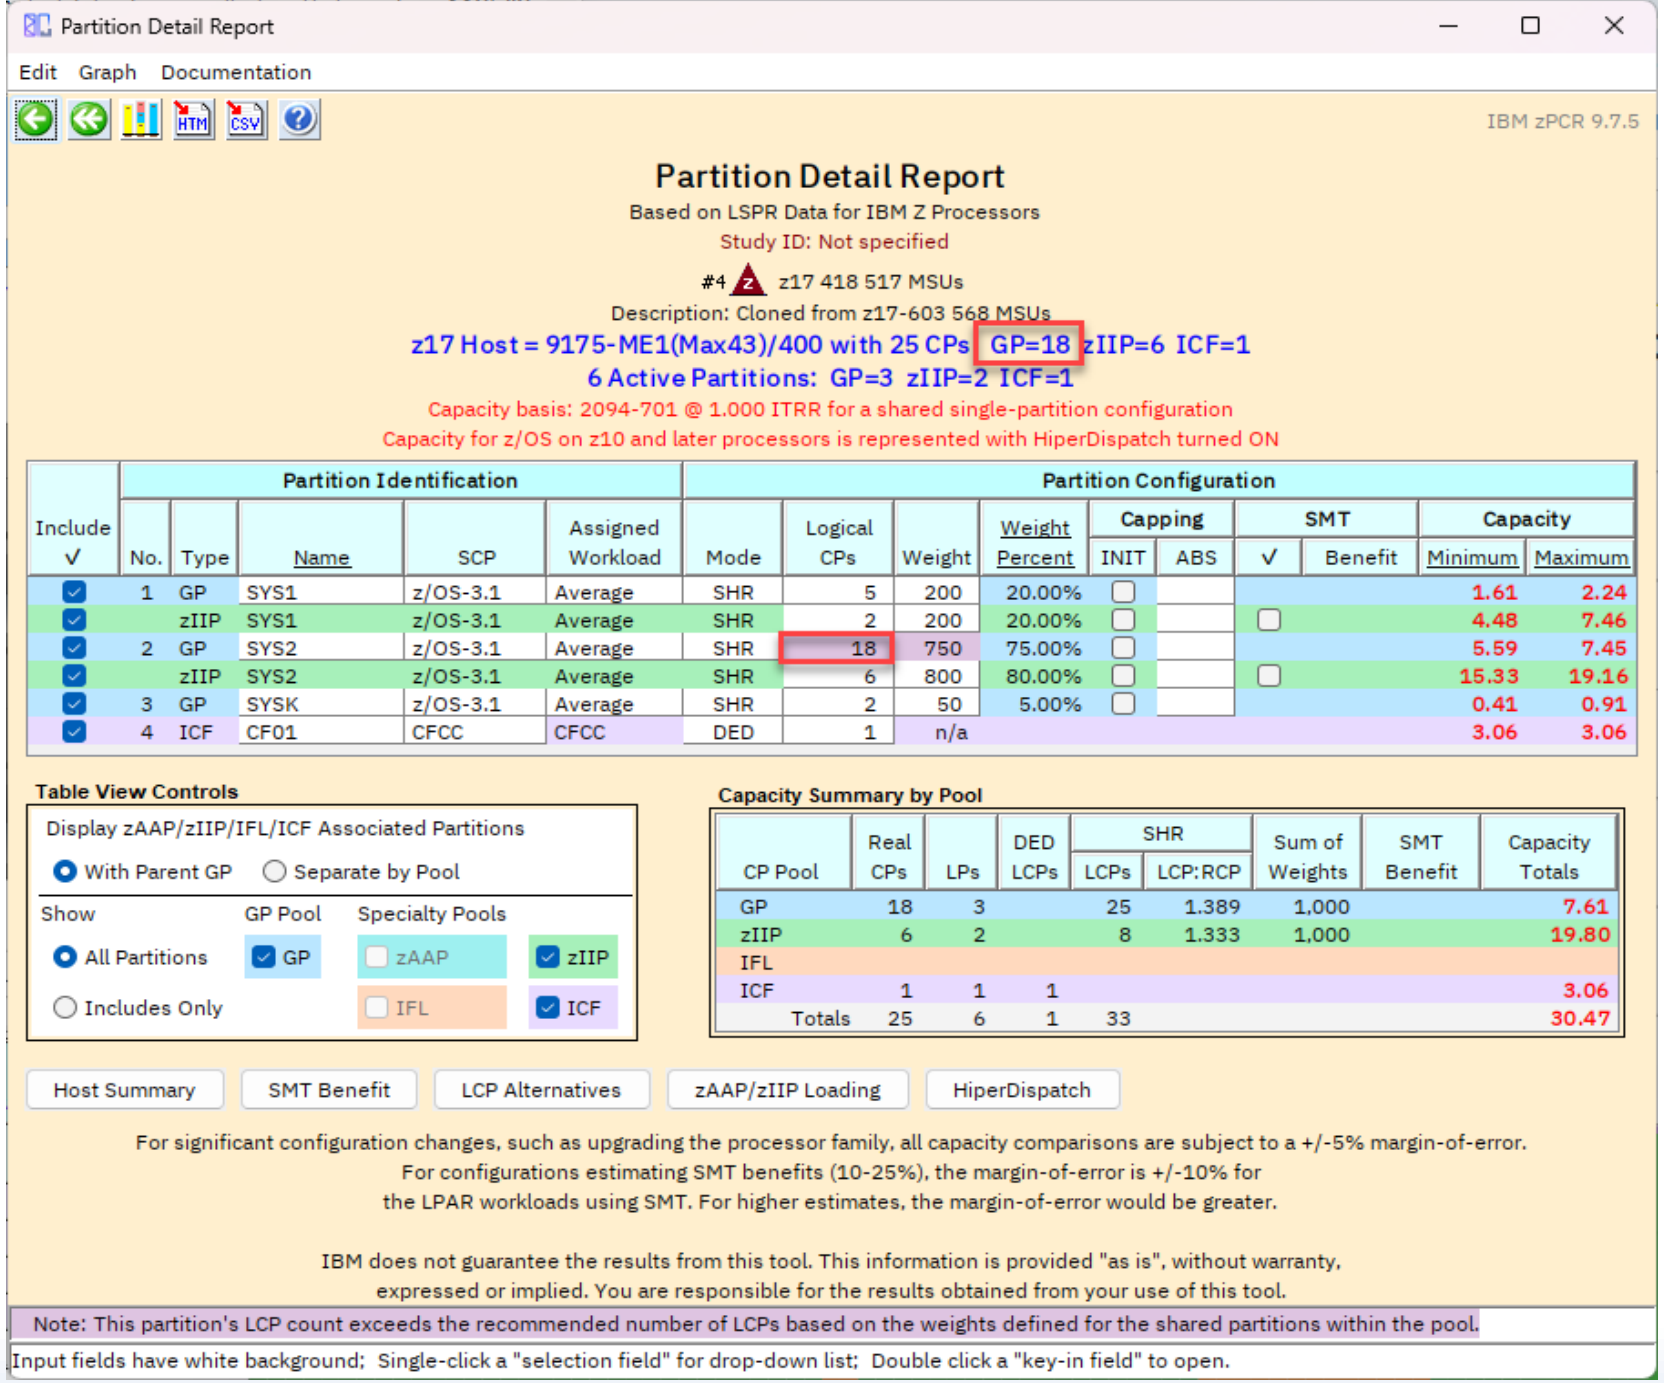

Note the change to the number of logical CPs to the larger LPARs when we increased the number of CPs on the machine.

Have to reduce when you go down in CPs.

Also note that it's warning us about having too many low-pool processors. A potentially valid point, but there also may be good reasons for doing this.

In this case, left it like this because on the "from" machine the LPAR had access to 100% of the capacity too.

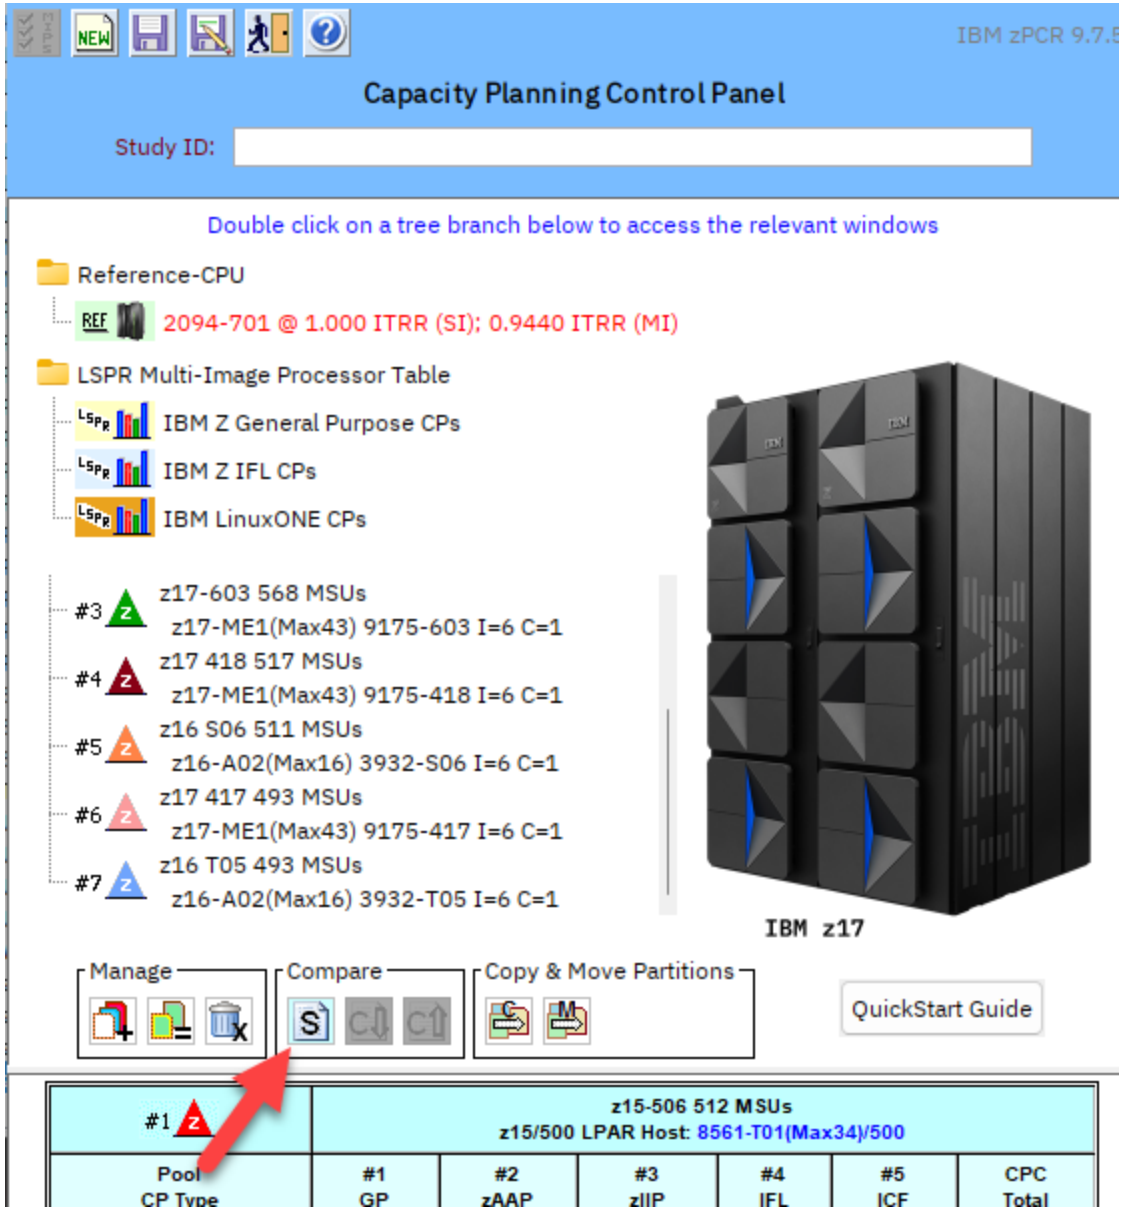

Note that I renamed the configurations to include the machine type and capacity that each scenario represents. This will make it easier to remember what scenario is what when you hit the compare button…

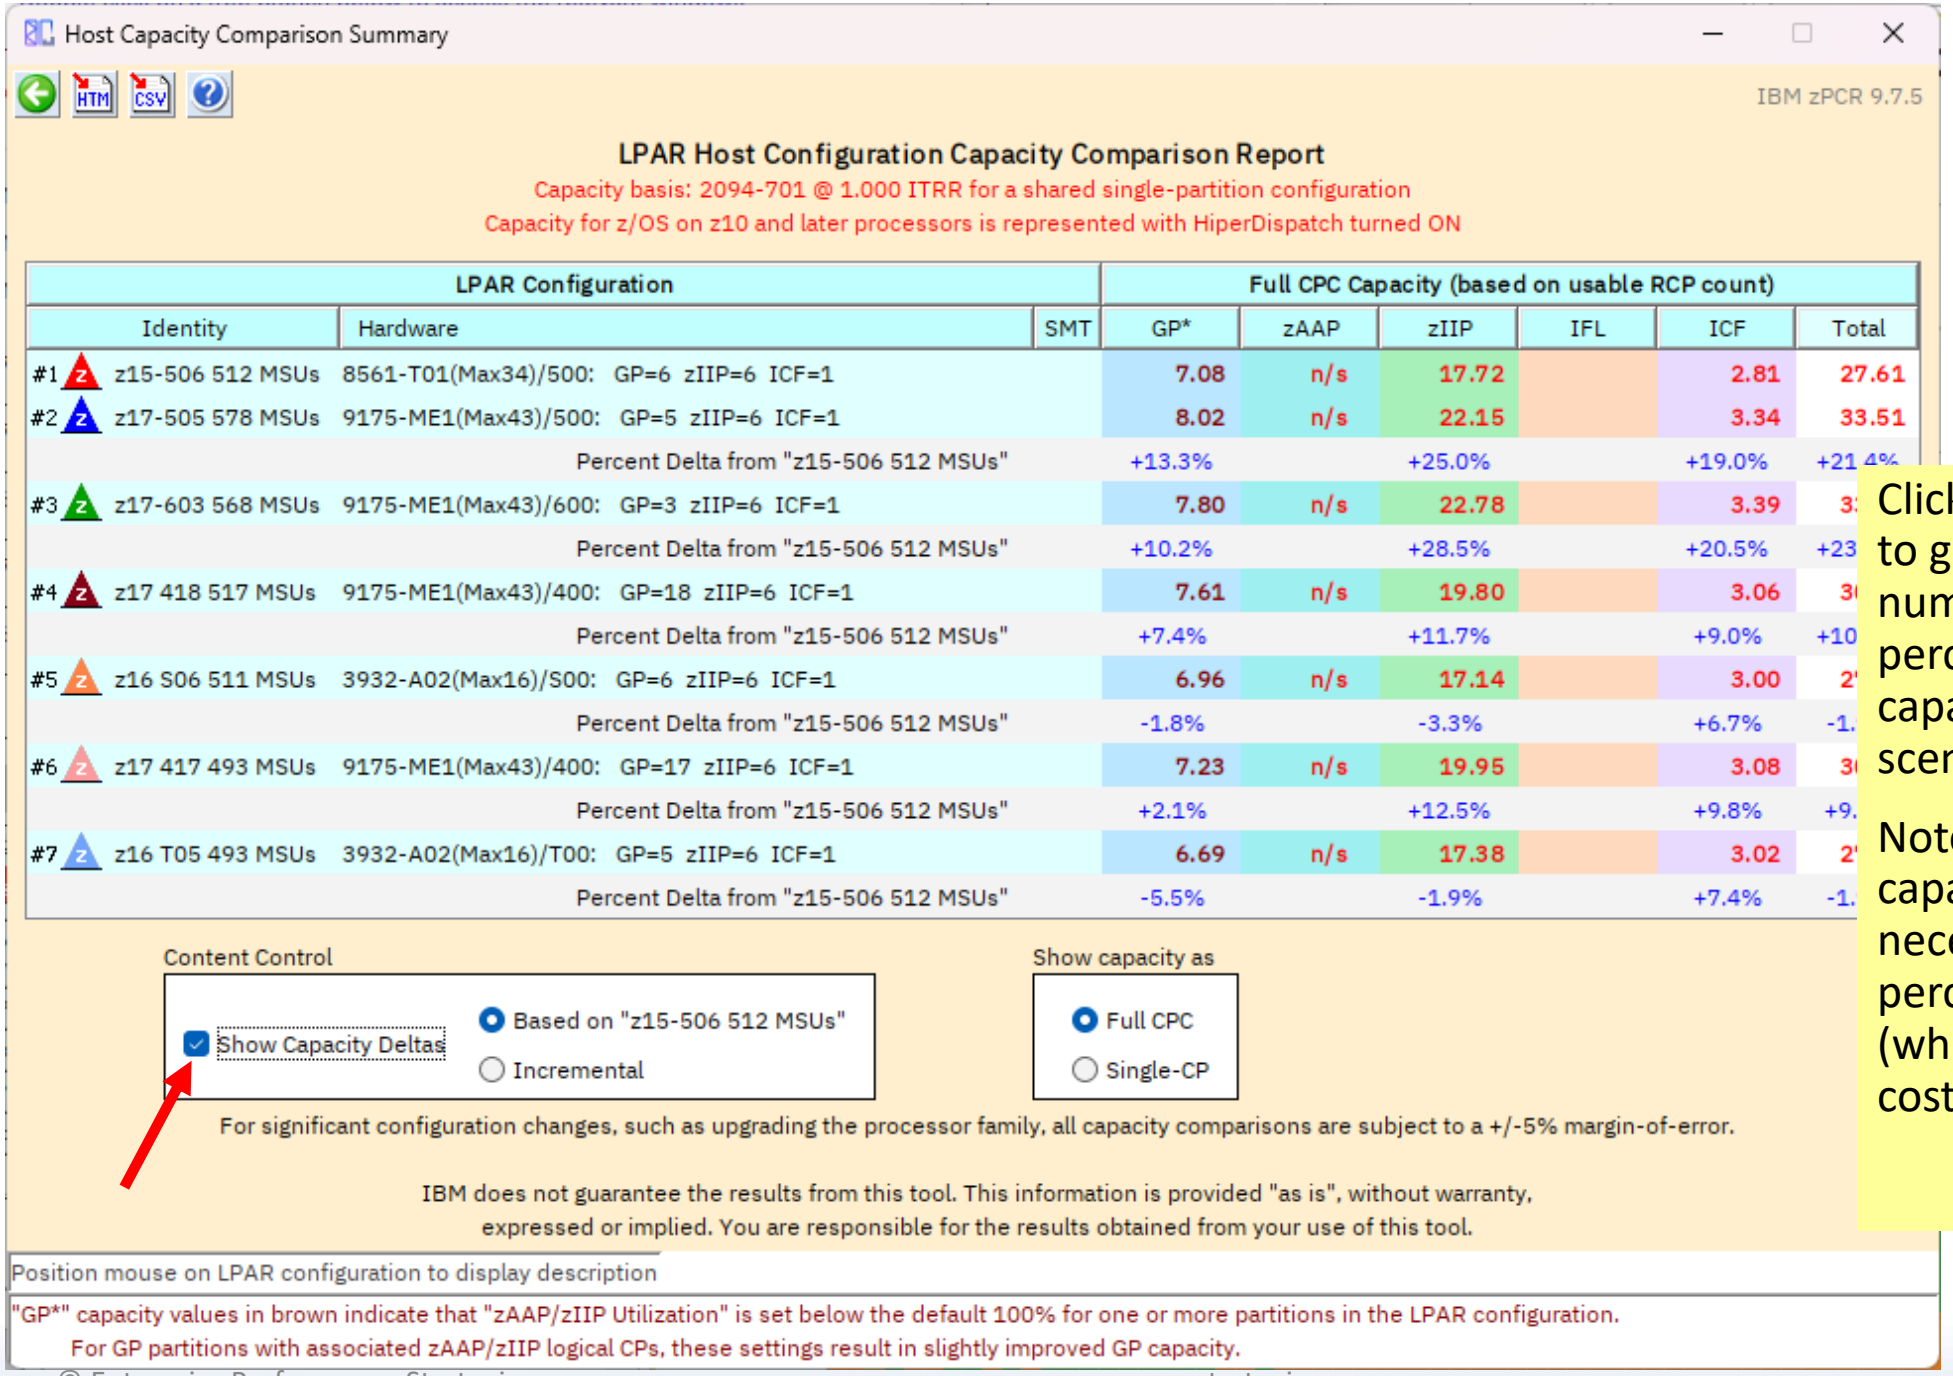

Click “Show Capacity Deltas” to get the important blue numbers which shows the percentage change in overall capacity from our baseline scenario.

Note that the predicted capacity changes are not necessarily aligned with the percent change in MSUs (which drive your software costs).

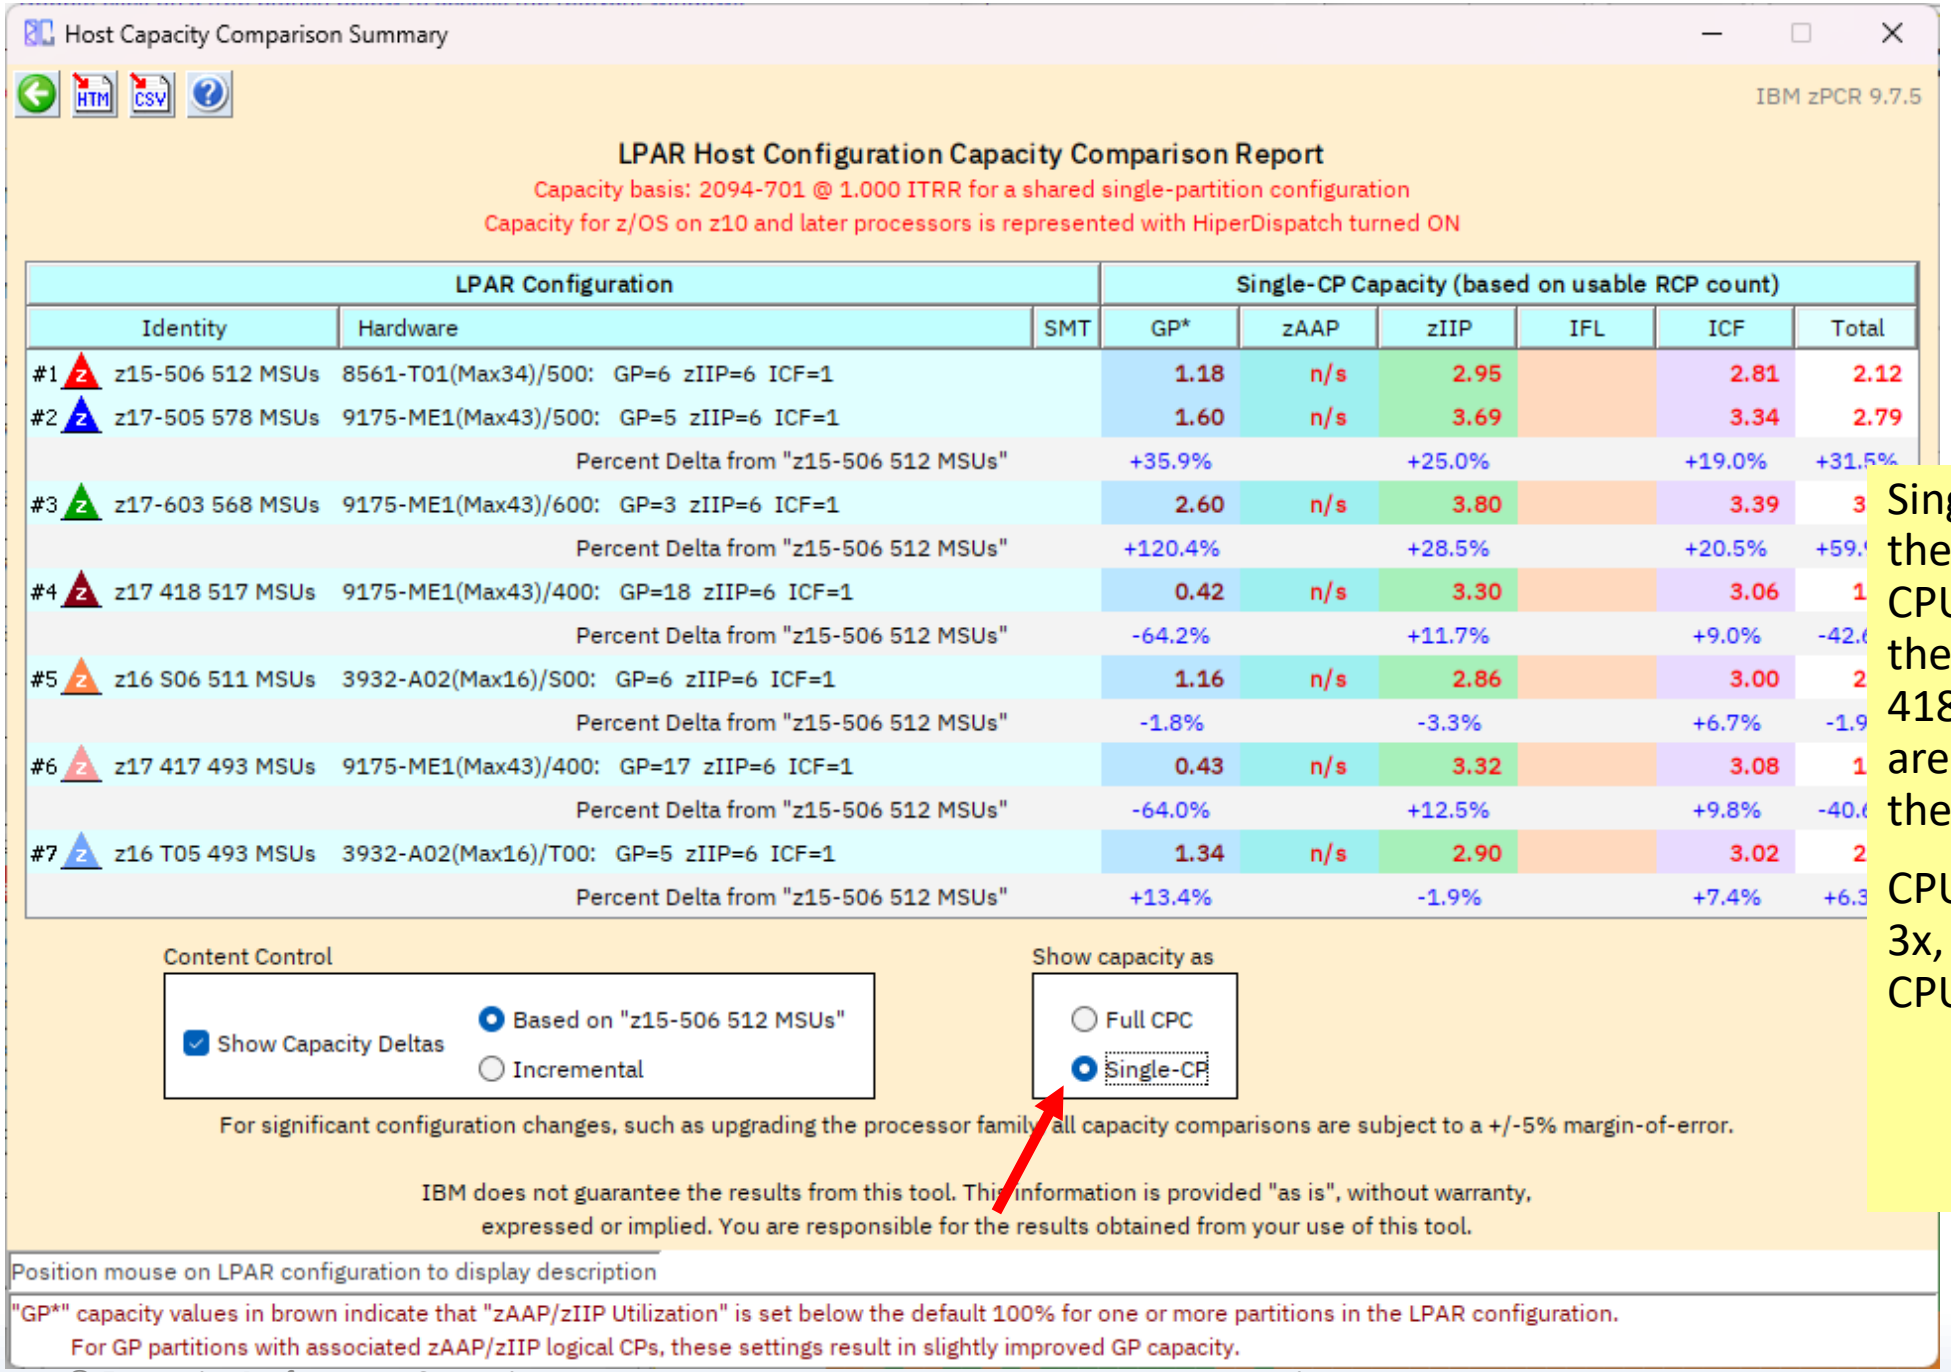

Single CP option shows you the impact on the individual CPU speeds. Here we see the potential issue with the 418 option: those engines are only about 1/3rd as fast the current ones.

CPU time will go up by about 3x, but there will be 3x more CPUs to dispatch on.

Comparing options

- I like to consolidate the options in Excel
- Note some options add less capacity than MSUs
- The z17 505 would be my “least risk” pick, but it does increase the installed MSUs by 13%: if there are full-capacity licenses, that may be a non-starter
- The z17 603 saves barely any MSUs while adding even less effective capacity
- The z17 418 (or even 417) adds more capacity than MSUs, potentially reducing software costs
  - But need to investigate if those slower CPs will really be ok
- The z16-S06 adds no MSUs, but shows a -2% in effective capacity, which is a bit concerning
  - Maybe should wait a few months for the smaller z17 machines to see how they look

		Relative to baseline		
	MSUs	MSUs	Capacity	Speed
z15-506	512	0%	0%	0%
z17-505	578	13%	13%	36%
z17-603	568	11%	10%	120%
z17-418	517	1%	7%	-64%
z16-S06	511	0%	-2%	-2%
z17-417	493	-4%	2%	-64%
z16-T05	493	-4%	-6%	13%

What about changing engine speeds?

- Do you have CPU time limits that need to change?
- If going to more/slower engines:
  - CPU time will increase, but CPU wait will likely decrease
  - Do you have single-TCB tasks that will be limited by the CPU speed?
  - CF Sync requests tie up less capacity
  - Performance may be more consistent
  - MSUs/unit of work may increase
- If going to fewer/faster engines:
  - CPU time will decrease, but CPU wait will likely increase
  - Misbehaving tasks & LPARs can dominate more of the total capacity
  - CF Sync requests tie up more capacity
  - Performance may be more variable
  - MSUs/unit of work may decrease

If you can, more slower is almost always more better!

Highest Task CPU Percent Programs by Hour

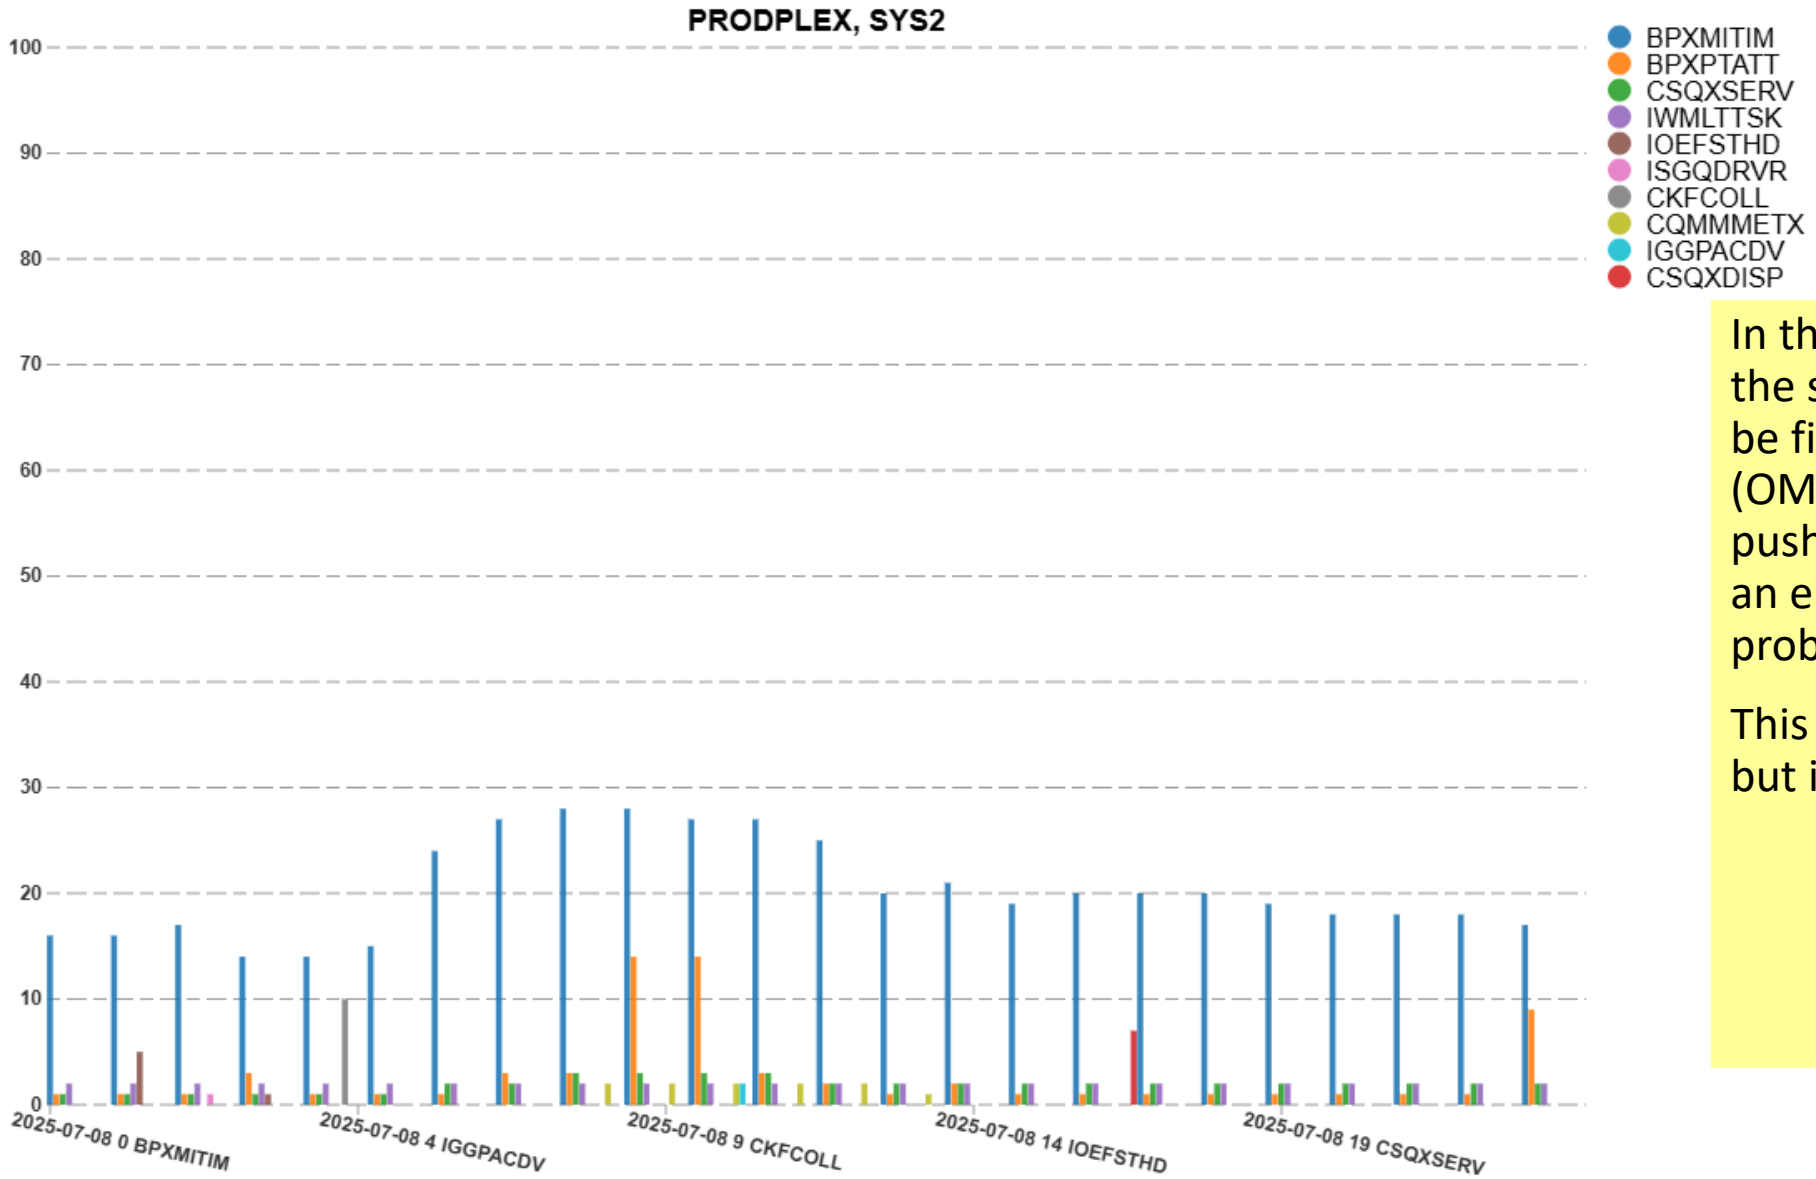

In this case it appears that the slower engines would be fine. That BPXMITIM (OMVS task) would be pushing up around 90% of an engine, but that's probably fine.

This is not always the case, but it's not uncommon.

Confirming your success

Set expectations early

- Why are you upgrading?
  - If it's to solve a problem, measure that "problem" before and after!
- Make sure you have good measurements from your important workloads
  - Understand the normal variability in those measurements
  - You may need multiple days of activity before drawing solid conclusions
- Understand the relative CPU speed change so you can set expectations for new CPU time measurements
  - See next slide
- Understand that some workloads will over- and some under-perform
  - Under-performers may be older programs that may just need a recompile
  - Update your compiler options for the oldest of production and DR machines

CPU Time Conversion

- $cpu_{new} = cpu_{old} / (1 + delta)$

Expected new CPU time

Old CPU time measurement

Single-CP delta from zPCR

- E.G. for a job using 500 CPU seconds on the z15-506
- On z17-505:
  - 500 / (1+.359) = 368 seconds
- On z17-418:
  - 500 / (1-.642) = 1397 seconds
- On z16-T05:
  - 500 / (1+.134) = 441 seconds

LPAR Configuration			Single-CP Capacity (b		
Identity	Hardware	...	GP*	zAAP	zIIP
#1 z15-506 512 MSUs	8561-T01(Max34)/500: GP=6 zIIP=6 ICF=1		1.18	n/s	2.95
#2 z17-505 578 MSUs	9175-ME1(Max43)/500: GP=5 zIIP=6 ICF=1		1.60	n/s	3.69
	Percent Delta from "z15-506 512 MSUs"		+35.9%		+25.0%
#3 z17-603 568 MSUs	9175-ME1(Max43)/600: GP=3 zIIP=6 ICF=1		2.60	n/s	3.80
	Percent Delta from "z15-506 512 MSUs"		+120.4%		+28.5%
#4 z17 418 517 MSUs	9175-ME1(Max43)/400: GP=18 zIIP=6 ICF=1		0.42	n/s	3.30
	Percent Delta from "z15-506 512 MSUs"		-64.2%		+11.7%
#5 z16 S06 511 MSUs	3932-A02(Max16)/S00: GP=6 zIIP=6 ICF=1		1.16	n/s	2.86
	Percent Delta from "z15-506 512 MSUs"		-1.8%		-3.3%
#6 z17 417 493 MSUs	9175-ME1(Max43)/400: GP=17 zIIP=6 ICF=1		0.43	n/s	3.32
	Percent Delta from "z15-506 512 MSUs"		-64.0%		+12.5%
#7 z16 T05 493 MSUs	3932-A02(Max16)/T00: GP=5 zIIP=6 ICF=1		1.34	n/s	2.90
	Percent Delta from "z15-506 512 MSUs"		+13.4%		-1.9%

Content Control

Show Capacity Deltas

Based on "z15-506 512 MSUs"

Incremental

Show capacity as

Full CPC

Single-CP

+/- 5% of these values is spot-on
Remember to consider the normal variability!

What to evaluate when

- Mostly looking for things whose CPU time is significantly higher than expected, but elapsed time can be important too
  - Elapsed more interesting when upgrading I/O subsystems
  - Faster channels / faster network may help some workloads though
- Can probably start to get some indications first business day for high-volume transactional workloads
- For once-a-day batch, you'll probably need several days to account for the normal variability
- Remember: it's normal for some work to over-perform and some to under-perform! The expectation is that you'll come out close to the expectation on average across all workloads.
  - Some work may be more/less impacted by processor architectural changes

Wrap-up

In summary…

- Don’t just accept the first offered solution: understand the pros and cons of different options
- Don’t overlook changing to a “mid-range” machine or to slower engines
  - These options might help save on your software cost
- Build your own zPCR models: it’s free, not difficult, and forces you to think through things
- Once you’ve selected an option, know what the expected change to the CPU time measurements will be
- Evaluate your most important workloads after the migration to make sure they’re within expectations

Your feedback is important!

SHARE
EDUCATE • NETWORK • INFLUENCE

Submit a session evaluation for each session you attend:

www.share.org/evaluation

Session 44850

SHARE Orlando
February 2026
Session TECH_135

Thanks!
Questions?

Email: scott.chapman@epstrategies.com