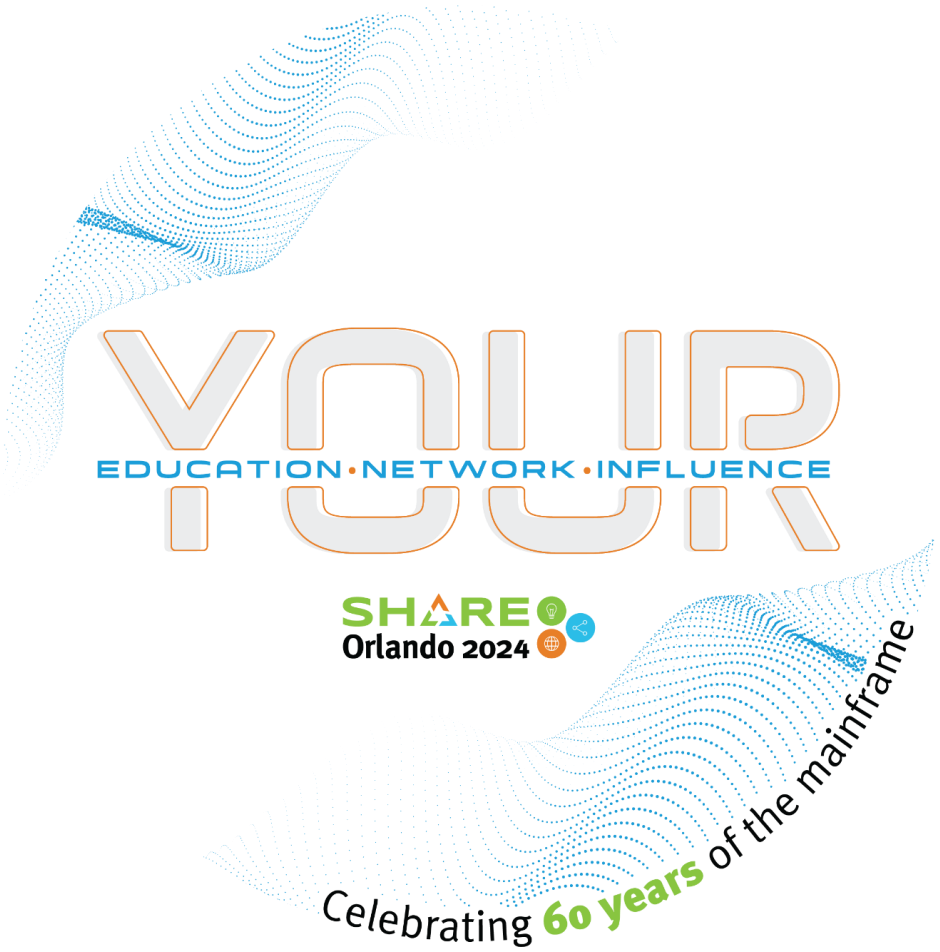

Configuring LPARS to Optimize Performance

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

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 (why you're here!)

PR/SM has been around for nearly 40 years now. PR/SM technology has advanced and evolved over these last 40 years. This also means that the most efficient configuration strategies haven also changed. What may have been a recommended configuration years ago may no longer be the best configuration today.

During this presentation, Scott Chapman will discuss PR/SM LPAR configurations to avoid. Scott will review these configurations and then explain why these configurations are not recommended. You will learn more about PR/SM, LPAR configurations, and processor measurements during this presentation.

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

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
  - September 16-20, 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!

EPS presentations this week

What	Who	When	Where
CPU Critical: A modern revisit of a classic WLM option	Peter Enrico Scott Chapman	Mon 4:00	Salon 12
30th Anniversary of Parallel Sysplex: A Retrospective and Lessons Learned	Peter Enrico	Tue 10:30	Salon 21
z/OS Performance Spotlight: Some Top Things You May Not Know	Peter Enrico Scott Chapman	Tue 1:00	Salon 15
The Highs and Lows: How Does HyperDispatch Really Impact CPU Efficiency?	Scott Chapman	Thu 10:30	Salon 21
Configuring LPARs to Optimize Performance	Scott Chapman	Thu 2:30	Salon 21

Things we'll talk about

- Why do you care?
- Memory
- LPAR Weight
- CPs
- Misc. Parameters
- Summary

Throughout we'll talk about some prior (historical) thinking about the topics as well as give some examples of problems we've seen in customers' configurations.

Why do you care about LPAR Configuration?

- Performance:
  - Improper configuration can impact performance
  - Poor performance can of course impact your business
- Financial
  - Improper configuration can increase your CPU consumption, potentially increasing your software bill
  - This is potentially even more important with TFP vs. RHA
- Understanding
  - Improper configuration can limit the data available for understanding and tuning your system
  - “You can’t manage what you can’t measure”

Memory

Old Thinking -> New Thinking

- Old:
  - Memory is expensive, we need to be careful how we hand it out
  - Hold back any memory that's not actively needed by the LPARs in case we have an emergency need for more memory
  - Be stingy in what we allow the address spaces to use
  - We have to precisely set LFAREA
- New:
  - Only hold back minimal amount of memory to account for growth
  - The LPARs have generous cushions and we monitor for growth to avoid surprises
  - Memory is a lot cheaper than CPU, can we improve performance and potentially reduce CPU consumption by being generous with memory?
  - LFAREA can be set much more generously because of z/OS 2.3 changes
- Still:
  - Don't be reckless: IEFUSI (and/or SMFLIMxx) are still a good idea!

CEC Average LPAR Storage Online

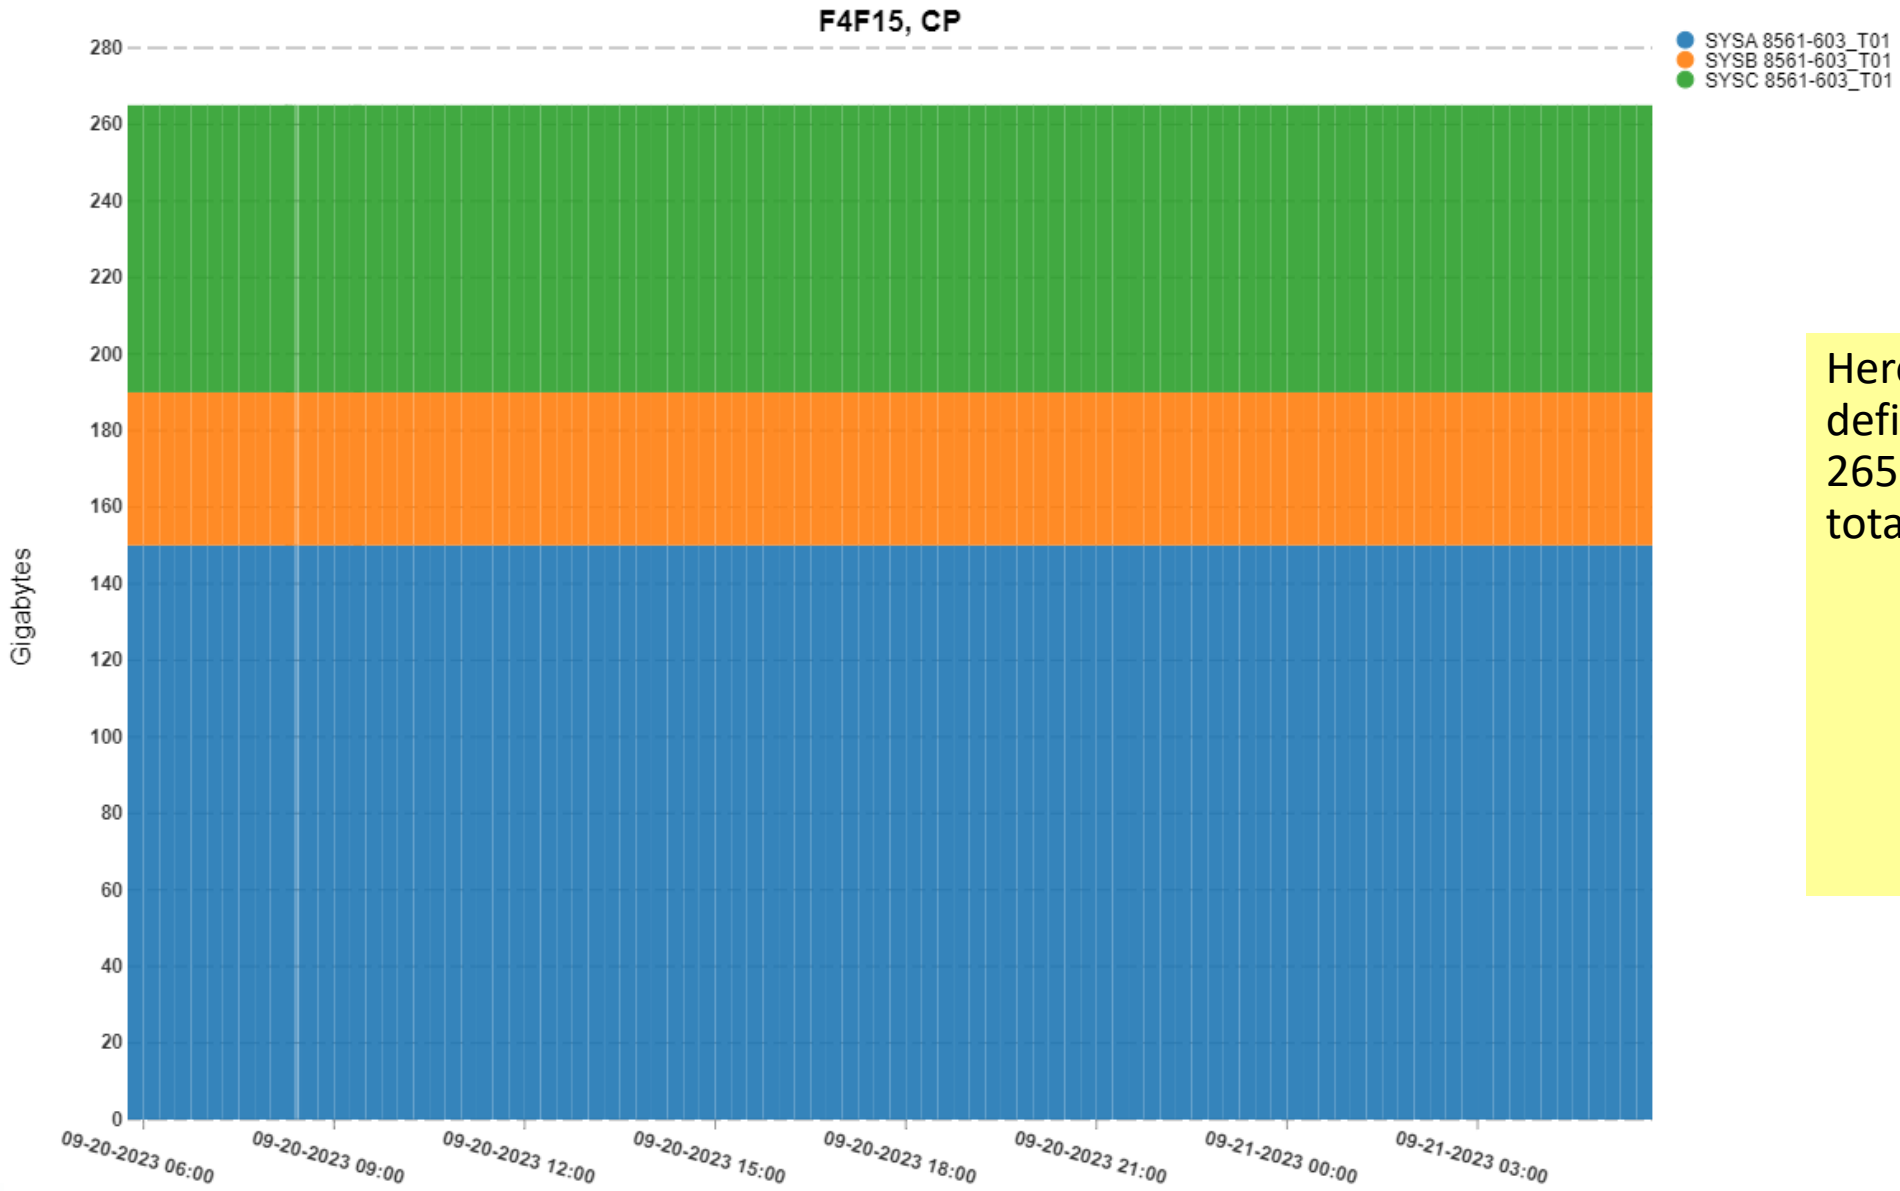

Here the z/OS LPARs are defined with about 265GB of memory in total.

CEC Average LPAR Storage Online

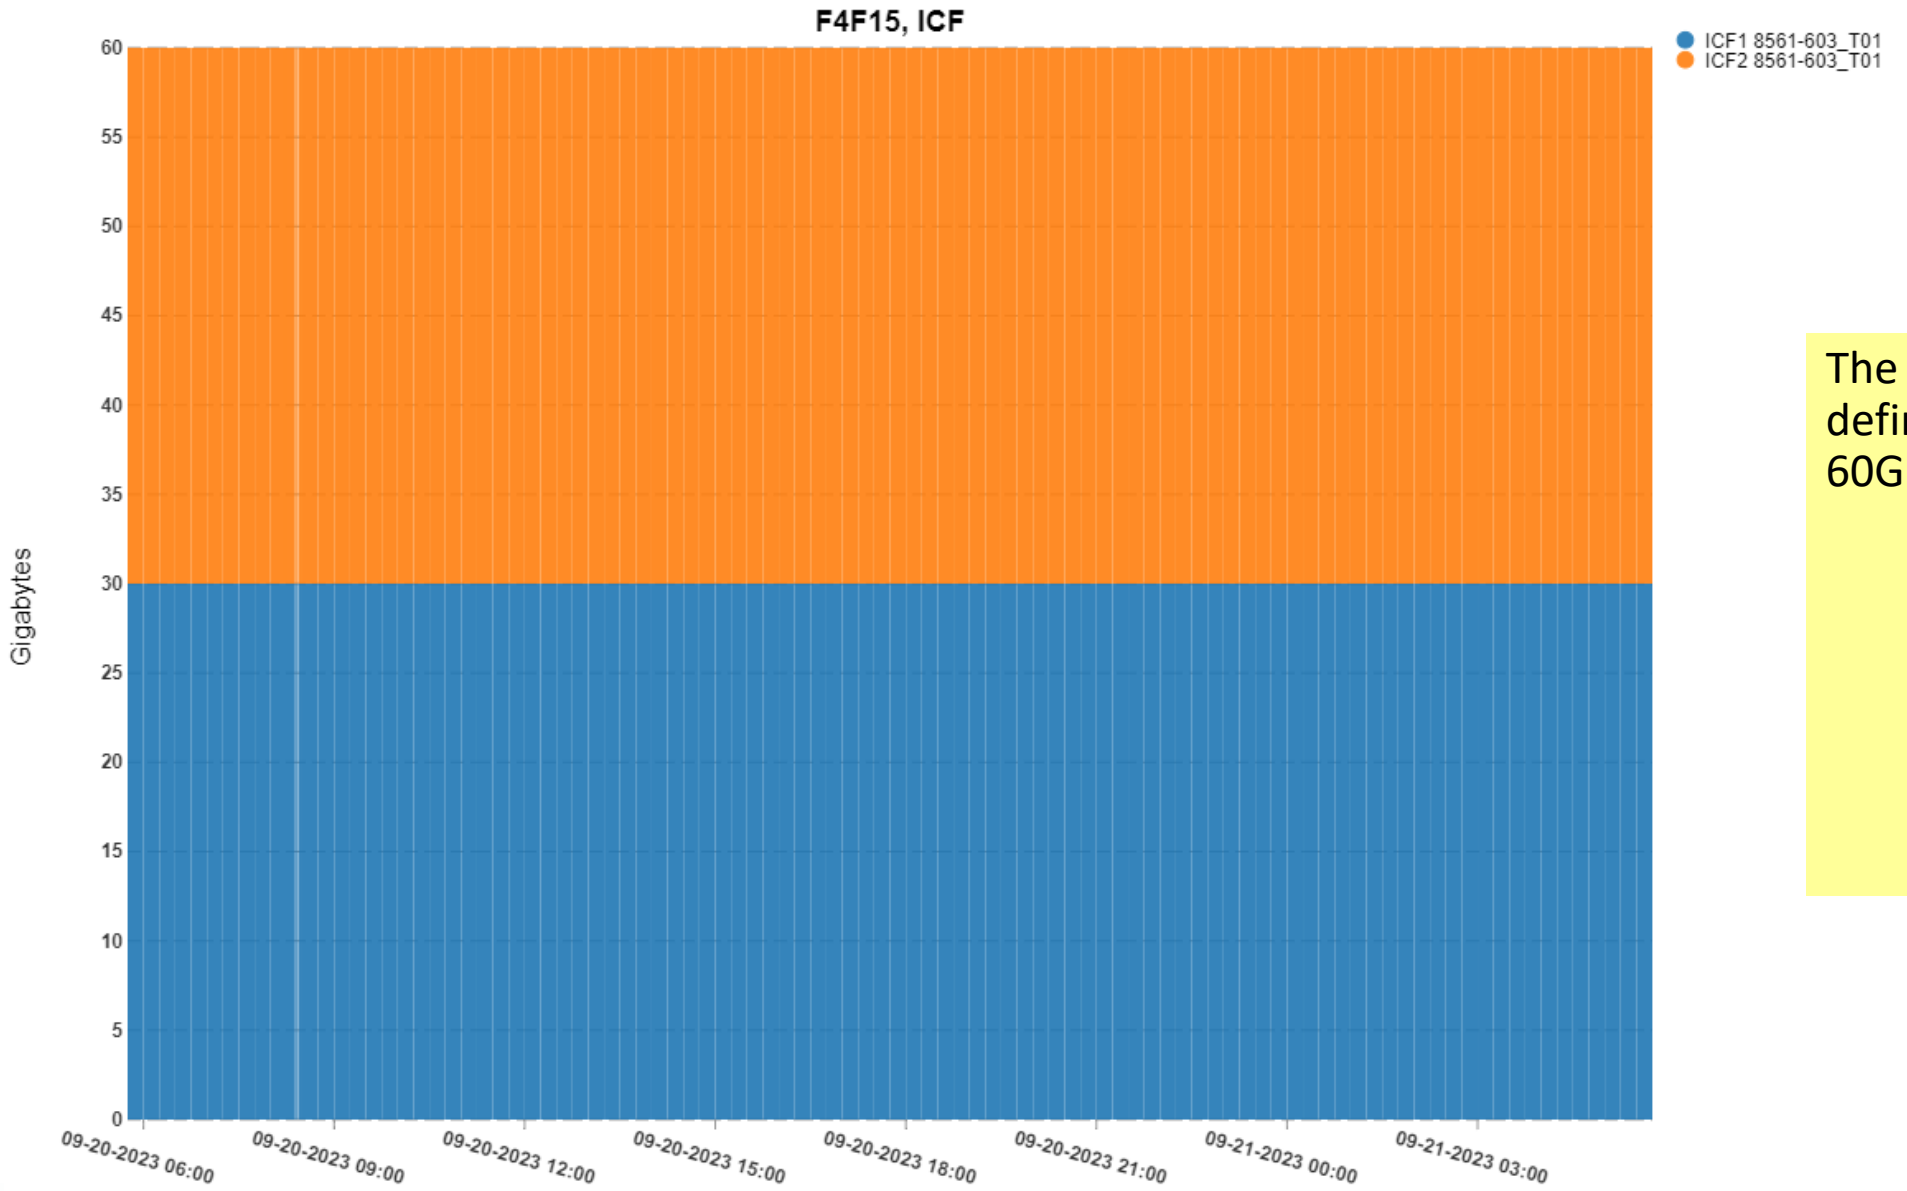

The CF LPARs are defined with another 60GB of memory

Observed Problem

- Some sites have held back way more memory than likely makes sense
  - We can't always (even usually?) tell this because the total installed memory is not in the SMF data, but…
- The minimal orderable memory on the z15 T01, z16 A01 machines is 512GB
  - So when we total up the LPARs and come to 325GB we do wonder if the rest of that memory could be put to effective use!
- If you have a plan for using this (e.g. you're planning on standing up some additional LPARs), then fine. But otherwise: use what you have!
- Scott's ROT re. reserved memory:
  - Old: keep 10% of installed memory in reserve
    - 10% of some machines is … a whole lot!
  - New: keep 10% (or less) of largest LPAR in reserve
  - Subject of course to site-specific situations

z/OS Configuration

- Once the LPARs have been given the memory: allow the address spaces to use it!
- Be generous in your LFA-area size for fixed 1MB pages
  - With z/OS 2.3 there were changes in how z/OS manages the storage areas and the fixed 1MB limit is now treated simply as a limit; i.e. a block of fixed 1MB pages is no longer set aside in the same manner as prior releases
  - 2GB pages more rarely used and still should be more closely managed (they do get a set-aside block of memory that can't be decomposed to smaller pages)
- See also:
  - https://www.pivotor.com/library/content/Chapman_MemoryMgtEvolutionWebinar.pdf

LPAR Weight

CPU Weight / Fair Share

- LPARs all have an assigned weight that effectively defines how much of the machine's capacity the LPAR should be allowed to use if the machine was 100% busy and all LPARs had demand
  - I sometimes call it the "fair share"
  - Separate weights by processor types (GP, zIIP, ICF, IFL)
- LPAR's "fair share" = LPAR's weight / sum(all *activated* LPARs weights)
  - E.G. an LPAR with a weight of 220, all LPARs total weight = 1000, means the LPAR's fair share is 22% of the machine's capacity
- Most customers configure the LPARs to allow them to "borrow" weight from other LPARs that aren't using it
  - I.E. if there's not demand from the other LPARs the LPAR can use more than it's "fair share"

“Difficult” math is not a problem

- It’s easy on us humans if the weights add up to 1000 because it’s then easy to infer the weights as a percentage
  - But PR/SM doesn’t care: it’s going to do the math regardless of the total
- I’ve seen some sites make their weights represent “MSUs” or “MIPS” and then ensure that all the weights add up to the total capacity of the machine
  - Again, PR/SM doesn’t care
- One thing that does trip people up sometimes:
  - Shutting down z/OS without deactivating the LPAR means that LPAR’s weight is still active and contributing to the total
  - Conversely, deactivating an LPAR will remove it from the total
  - So we see sometimes active totals that are something like 1075 or 935 because somebody left an LPAR activated or forgot to change weights when deactivating one

Old Thinking -> New Thinking

- Old:
  - Weights were almost always static (setup once, then rarely adjusted)
  - Changing weights required making manual changes on the HMC
  - Broadly defined, set rather coarsely
- New:
  - Weights probably should be changed dynamically to support changing workload
  - Automation (including REXX scripts) can change weights
  - Fine weight adjustments may have a measurable impact
- Scott's opinion: most customers probably should be dynamically changing their weights, but very few do

Observed Problems

- We often see LPAR with sub-optimal weights
  - LPARs should have weight to match their capacity requirements
    - While most sites allow LPARs to borrow unused capacity from other LPARs, this is not optimal
  - Some LPARs may benefit from small adjustments to get an additional high-pool CP
  - Sometimes this is because LPARs have been added or removed
    - Sometimes DR or test LPARs are shutdown but not deactivated meaning their weight is still influencing the active LPAR's share
- LPARs consuming more than their weights are at risk of being stolen from
  - This could have a significant negative impact on the work
- Weight concerns also relate to processor efficiency under HiperDispatch

CEC Assigned CP LPAR Weights

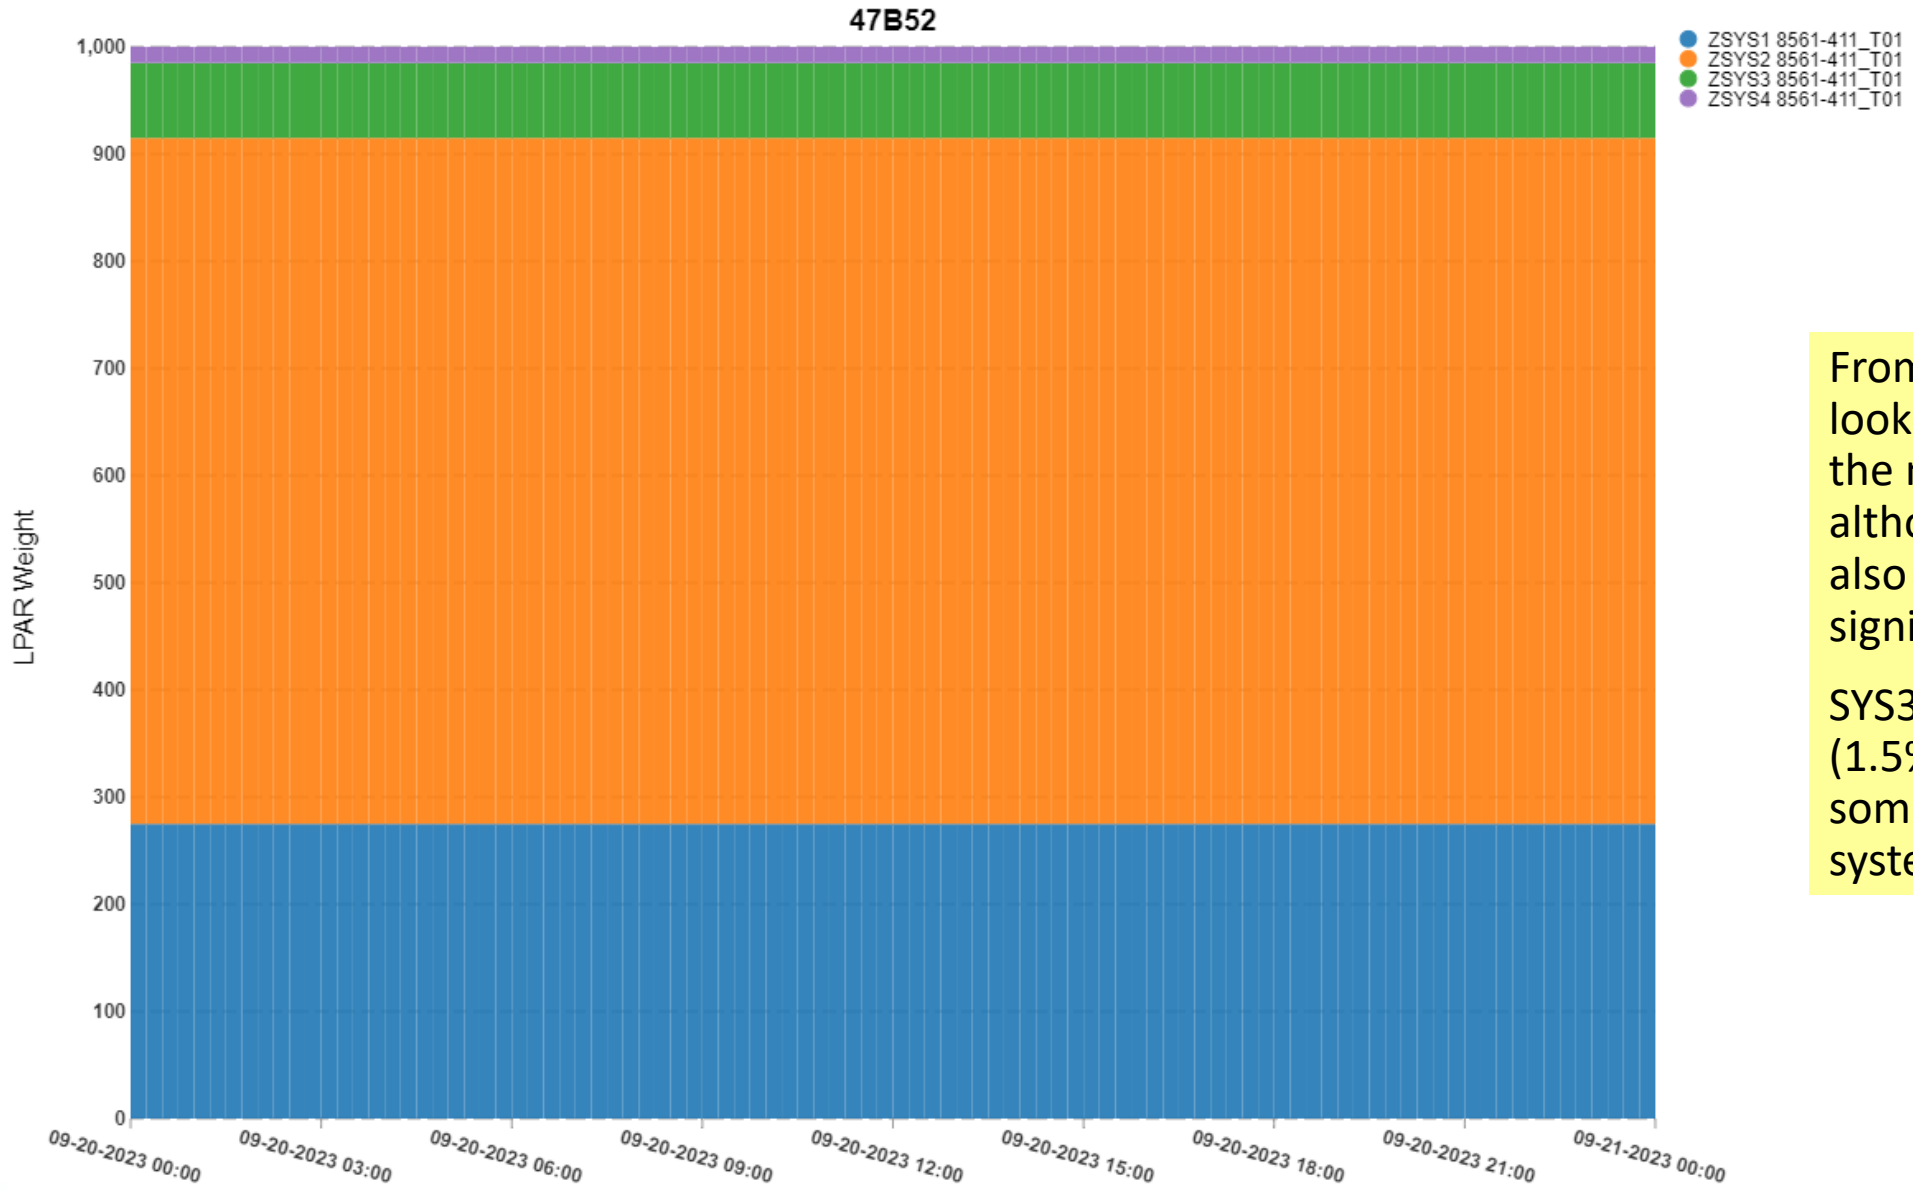

From the LPAR weights it looks like SYS2 (64%) has the most work to do, although SYS1 (27.5%) also seems quite significant.

SYS3 (7%) and SYS4 (1.5%) are probably some sort of dev/test systems.

CEC Percent CP Weight Used

This shows how much capacity each LPAR is consuming relative to its weight.

Note SYS1 rarely even consumes 50% of its weight, while SYS2 in the early evening needs much more than its weight.

SYS3 regularly well above its weight.

Risk mitigation

- What happens if SYS1 has an increase in demand in the early evening?
  - SYS2 (and SYS3) are going to lose access to that weight they're borrowing from SYS1
  - This will add CPU delay to SYS2 and SYS3, impacting work running there
- It may be better to give some of SYS1's weight to SYS2 and (maybe) SYS3
  - Unless SYS1 is in fact really important, and we want it to guarantee access to that capacity
- In many cases it may make sense to move weight between LPARs at different times of the day or different times of week
  - This can be done with BCPii via some REXX (or C or Assembler) code
  - See also: https://github.com/IBM/zOS-BCPii
  - In theory, just-in-time weight movement might be an interesting ML opportunity
    - Actually: probably don't need ML at all

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 essentially dedicated to the LPAR (100% share)
  - Medium – Share between 0% and 100%
  - Low – Unneeded to satisfy LPAR’s weight
- 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

Highs and Lows

- In the prior example, SYS2 or SYS3 might be able to get an additional high pool processor based on getting additional weight from SYS1
- But the bigger problem may be that the capacity used above the LPARs' weights are run on low pool processors
- Generally speaking:
  - High pool processors may be more efficient
  - Low pool processors will likely be somewhat less efficient
- "Efficient" = less CPU consumption required to do a given amount of work
  - Efficient = better performance
  - Efficient = less capacity consumption
  - Less capacity consumption may mean lower software costs
- See my session from this morning for more details!

CPs

Logical vs. Physical CPs

- LPARs use “logical” CPs which are allocated some fraction of a physical CP
  - Highs: 100%
  - Lows: 0%
  - Medium: >0%, <100%
  - Especially for medium and lows, there will be times when the logical CP is not dispatched by PR/SM to a physical CP
- LPAR cannot have more logicals defined than physicals enabled
  - But can have reserve CPs that can be brought online during CoD/CBU events
- If all the logical CPs are busy, it doesn’t matter how much additional physical capacity is available, the LPAR won’t be able to get to it
  - This can be used as an intentional limiting factor (although there may be better ways)

Old Thinking -> New Thinking

- Old:
  - Carefully assign just enough logical CPs to the LPAR
  - Get the fastest CPs you can!
- New:
  - Having 2-3 extra (VL) CPs is not a penalty and allows the LPAR to consume more capacity if it needs to (VL efficiency issues not withstanding)
  - More/slower CPs (sub-cap engines) is often better than fewer/faster CPs
- Scott's opinion: more/slower is almost always better than fewer/faster
  - More/faster would be even better of course
  - Just because you can run on a 1-way or 2-way machine doesn't mean you should

Observed Problem

- Sometimes we find LPARs that have too many or too few logical CPs
  - "CPs" here can mean either GCPs or zIIPs
- Too many implies too many low-pool processors
  - IBM recommendation: have no more than 2 low-pool processors
  - Scott's recommendation: it depends…
    - Unused low pool processors don't hurt
    - Used low pool processors imply a need for weight adjustments
    - Occasional brief use of low pool processors is probably fine
- Too few implies possible limitations imposed on the work
  - Not enough CPs to dispatch on
  - Also: avoid single-engine z/OS LPARs
- Because too few is more impactful than too many, better to over define if possible
  - *This was not true prior to HiperDispatch*

For large systems: limit CPs/zIIPs to less than the number in a drawer

HiperDispatch - Parked / Unparked CPs

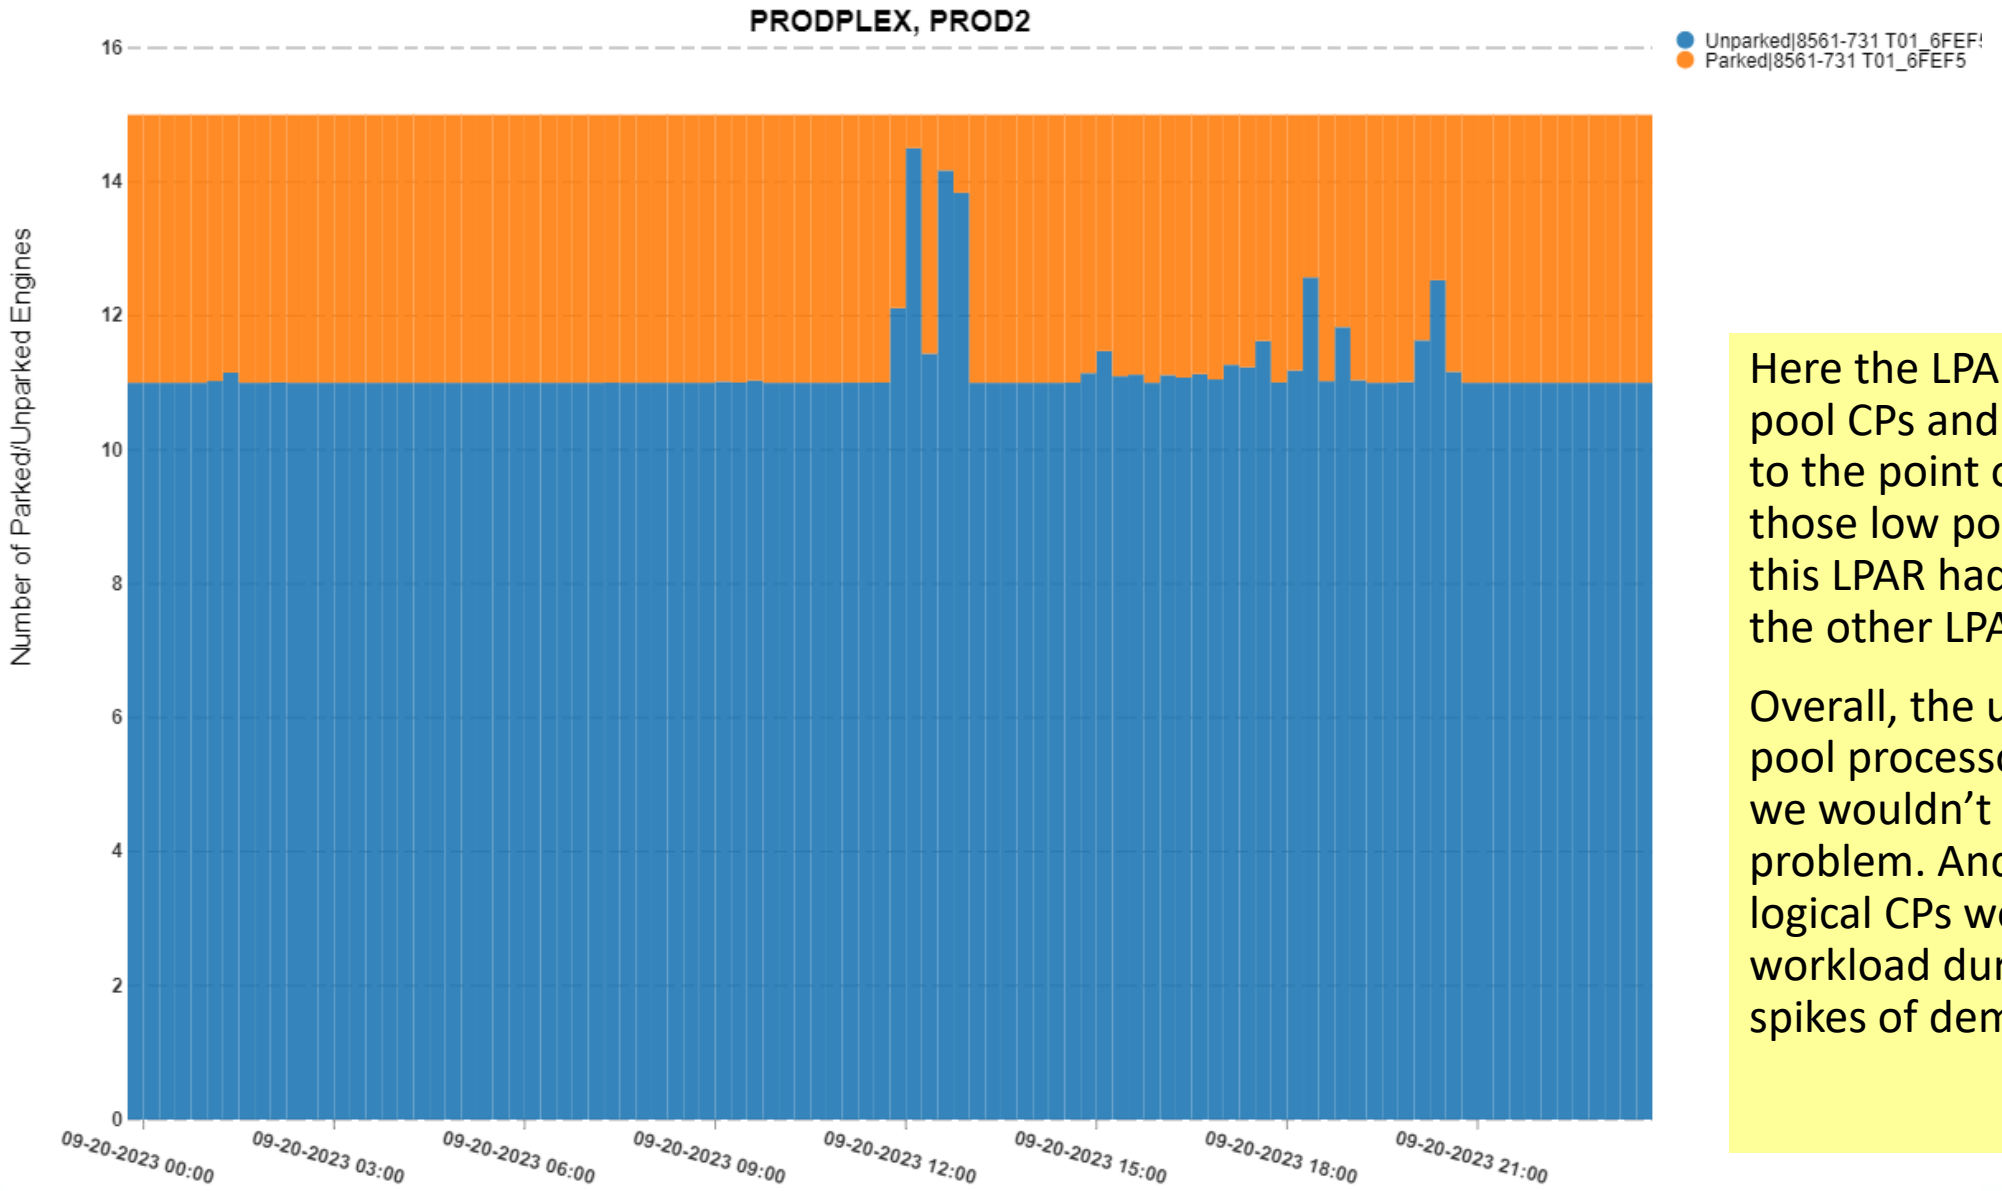

Here the LPAR has 4 low pool CPs and did (rarely) get to the point of using all 4 of those low pool CPs when this LPAR had demand and the other LPARs did not.

Overall, the usage of low pool processors is limited, so we wouldn't consider this a problem. And taking away logical CPs would limit the workload during those spikes of demand.

LPAR, MVS, and Workload CP Busy% with Capture Ratio

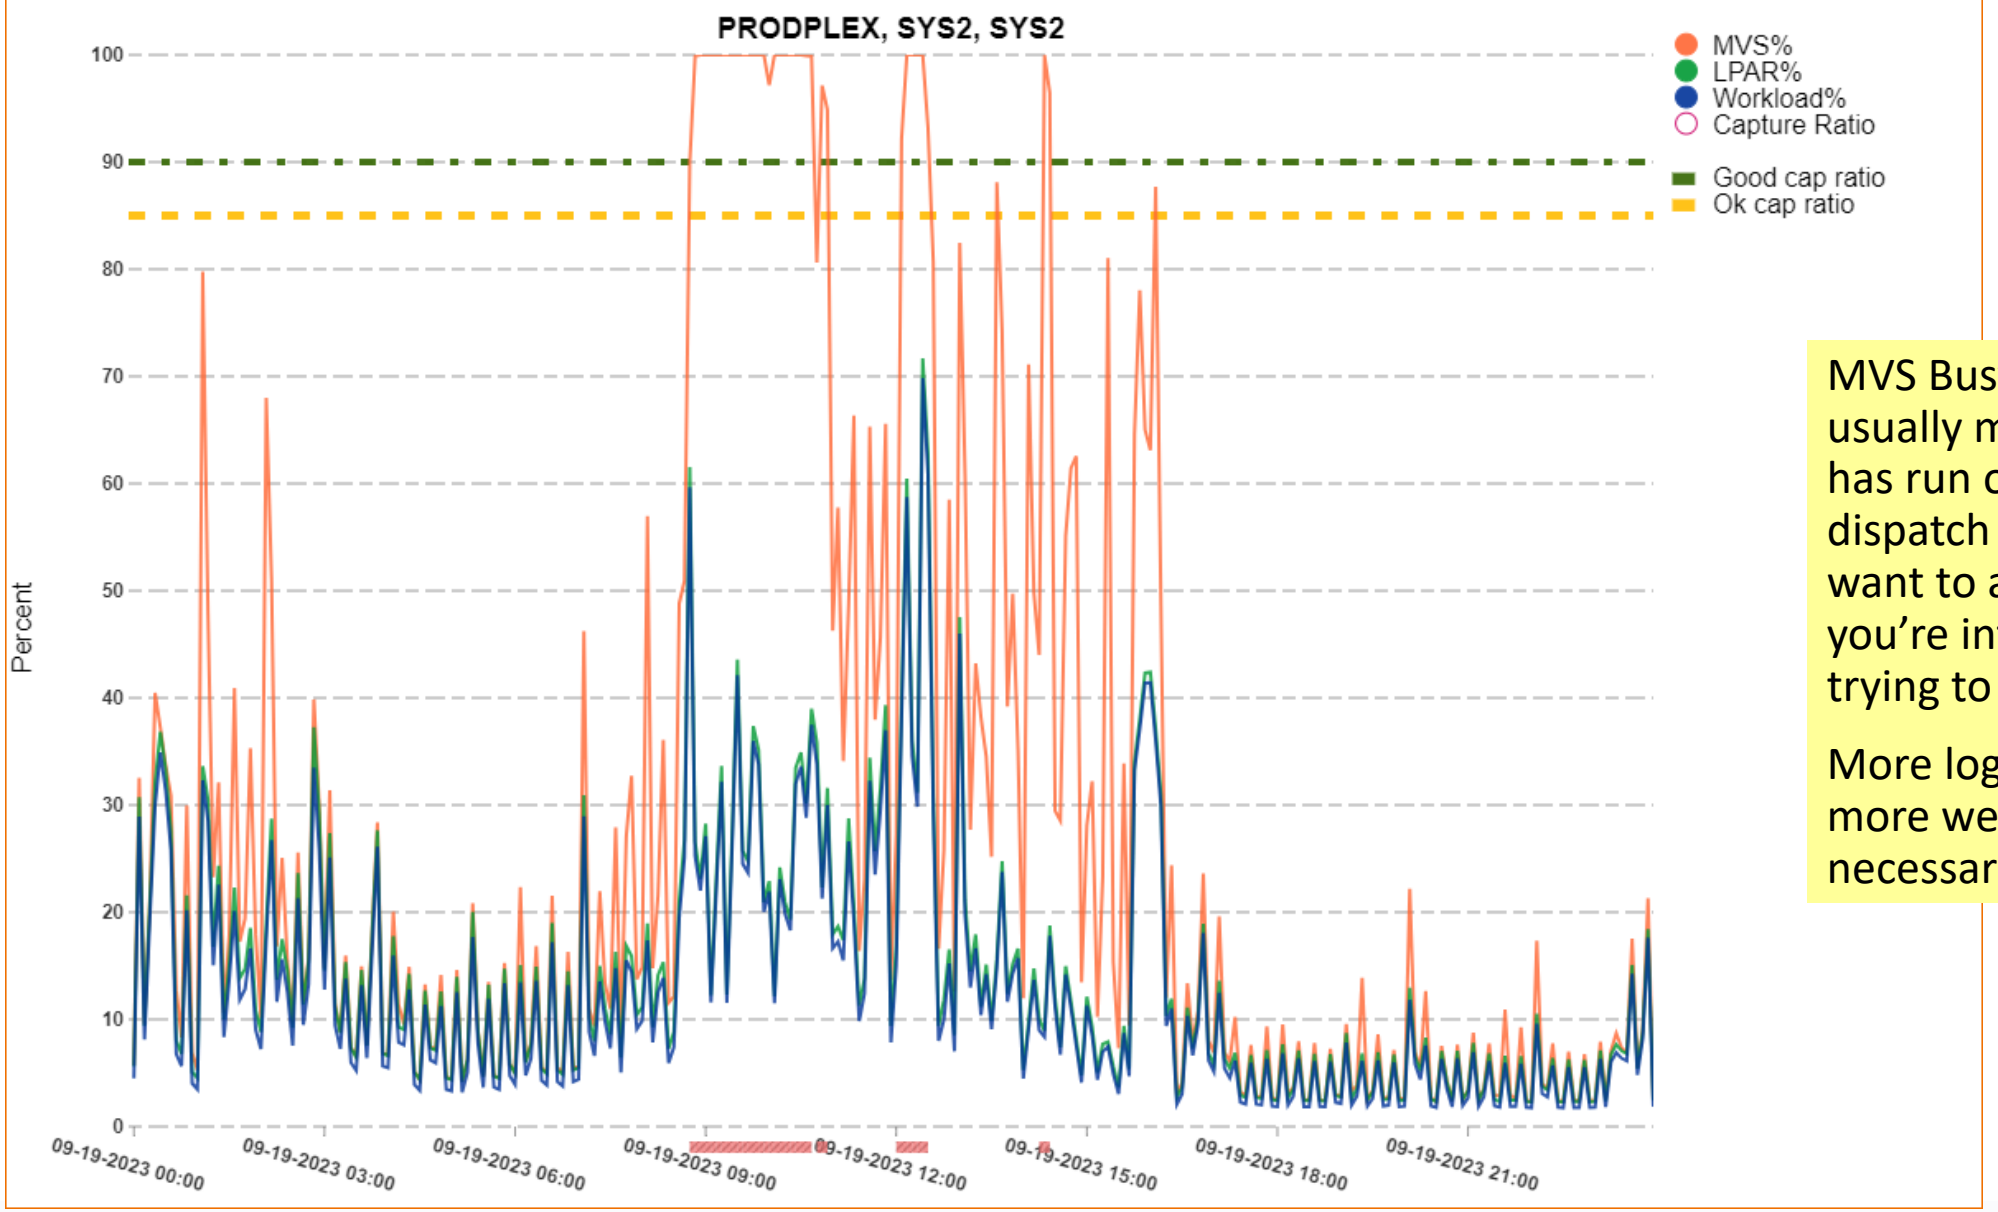

MVS Busy going to 100% usually means the LPAR has run out of logicals to dispatch on. Generally want to avoid this unless you're intentionally trying to limit the LPAR.

More logicals and/or more weight may be necessary.

Observed Problem: Physical CPs

- **You need enough physical CPs to dispatch on**
  - We've seen multiple cases where customers have had issues after going to fewer/faster CPs
- In most cases, more/slower is better than fewer/faster CPs
  - Most systems have multiple LPARs share those physical CPs
  - Most LPARs have many tasks, often trying to run at the same time
- Sometimes do have single-TCB task issues
  - Most often CICS regions that are constrained by QR TCB
  - If you have this situation, you should look at application changes to resolve it
    - Thread-safe
    - Splitting the work across multiple AORs
    - Application tuning

Recent case in point

- Customer upgraded from z14 to z16
  - From: z14 3907-T05, 5 CPs, 450 MSUs
  - To: z16 3932-W03, 3 CPs, 455 MSUs
- Reported “some issues” after upgrade
- Comparison of similar days follows…

WLM RT Goal - Top Ended Transactions

(Period of Study for RT Goal Periods)

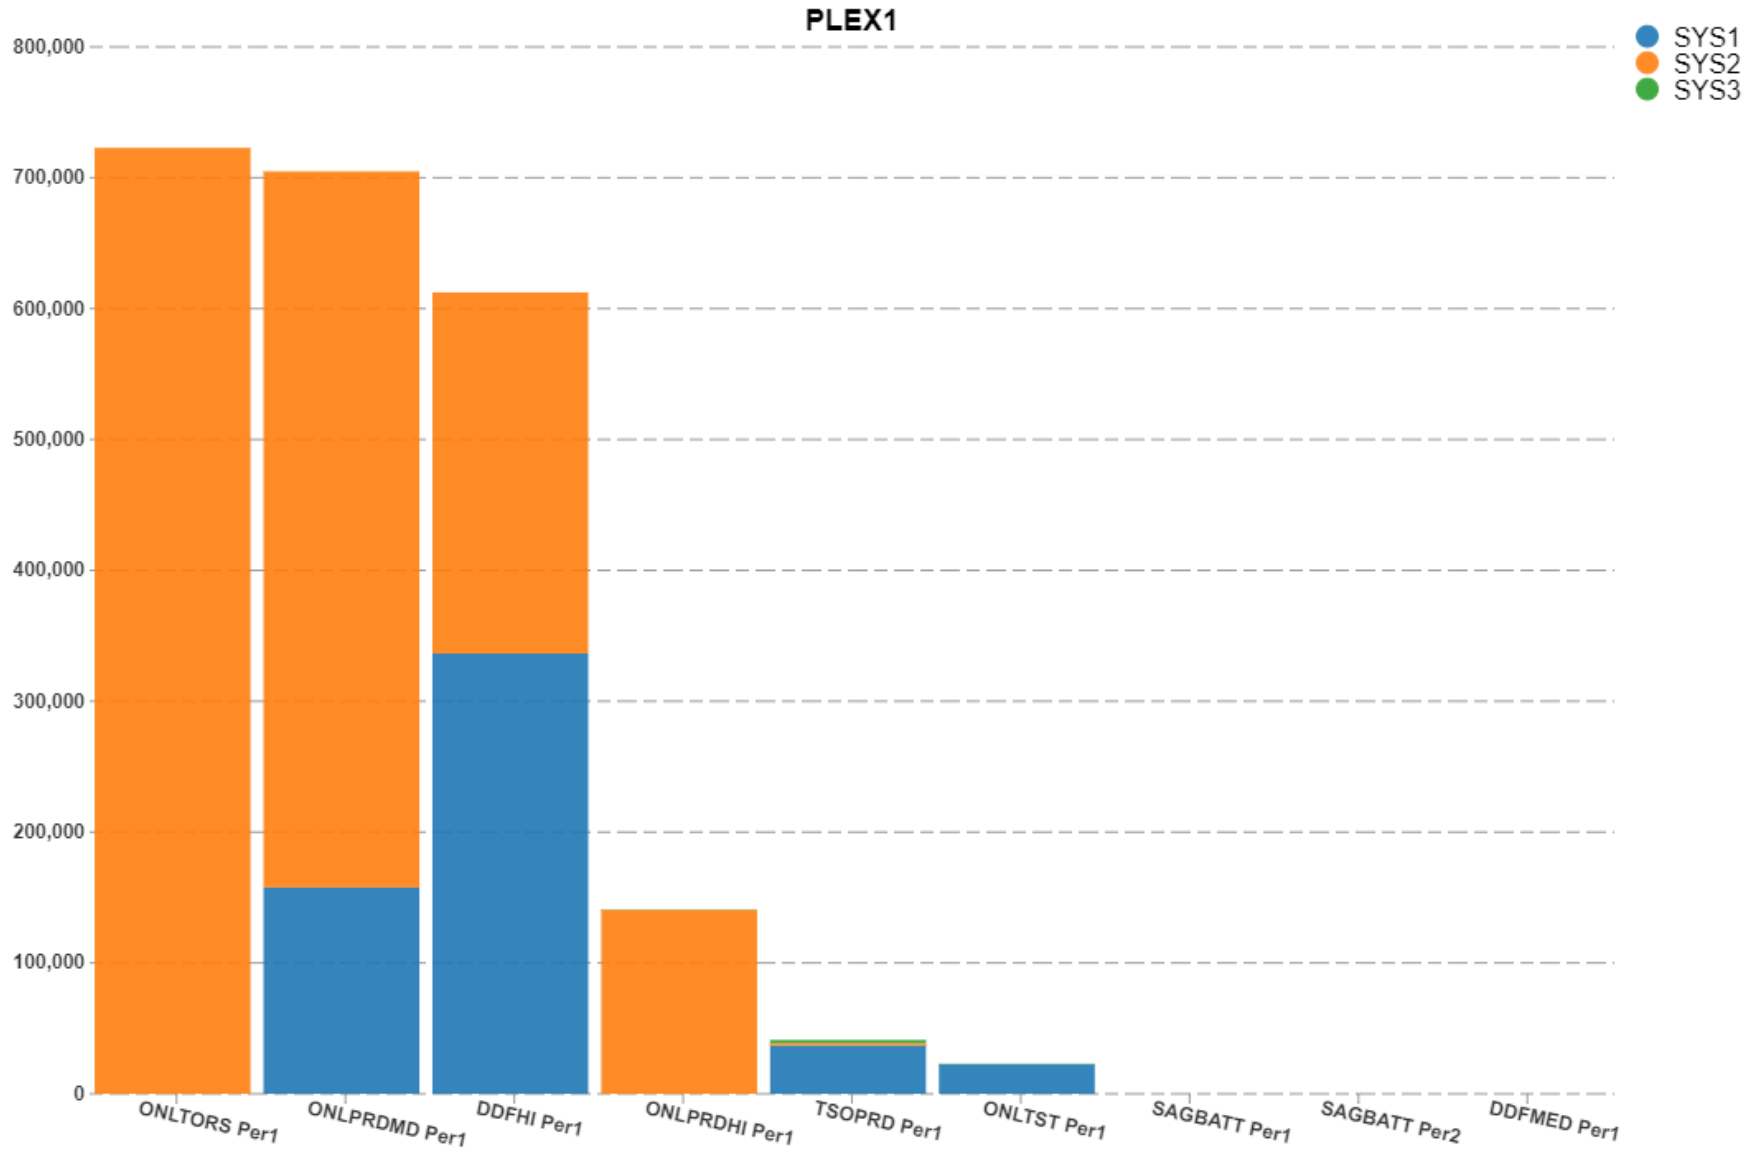

Before

WLM RT Goal - Top Ended Transactions

(Period of Study for RT Goal Periods)

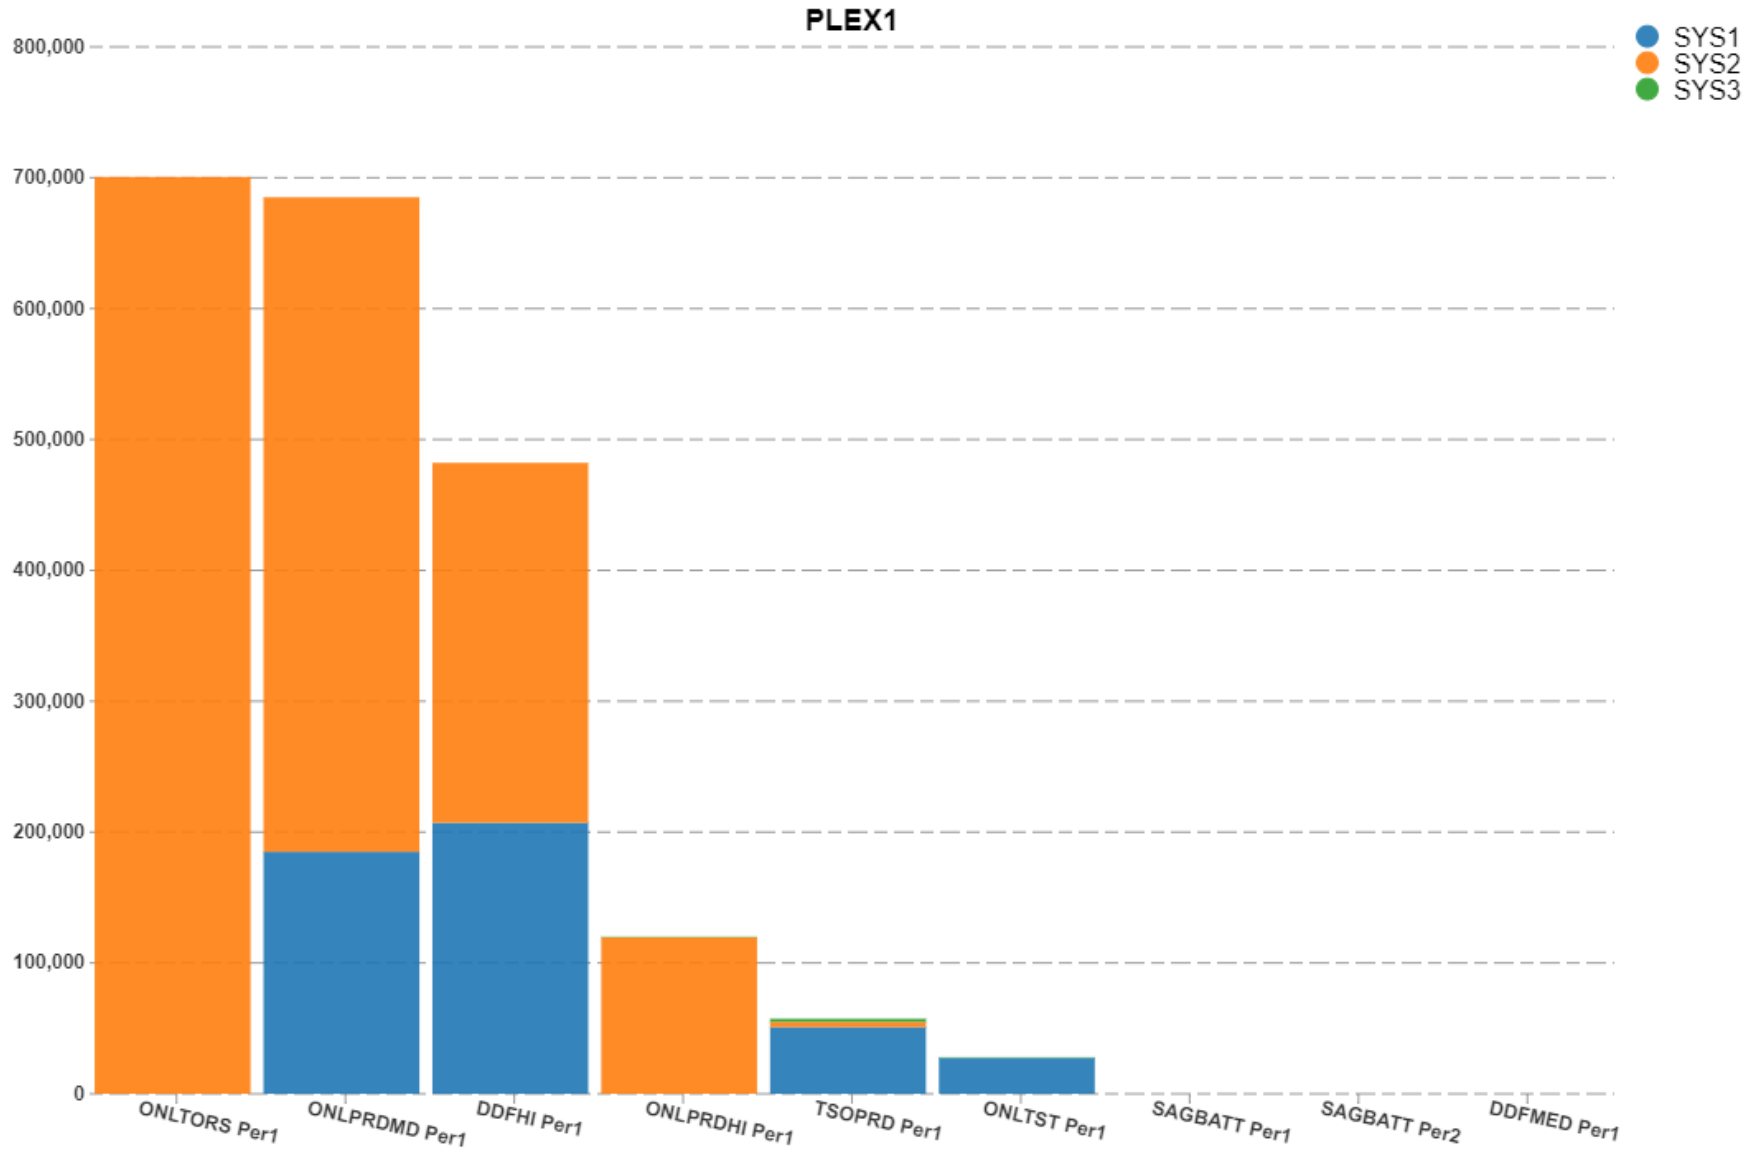

After

Note very similar (but slightly lower) transaction volumes for the day

WLM CPU - CP CPU Delay Samples By Period

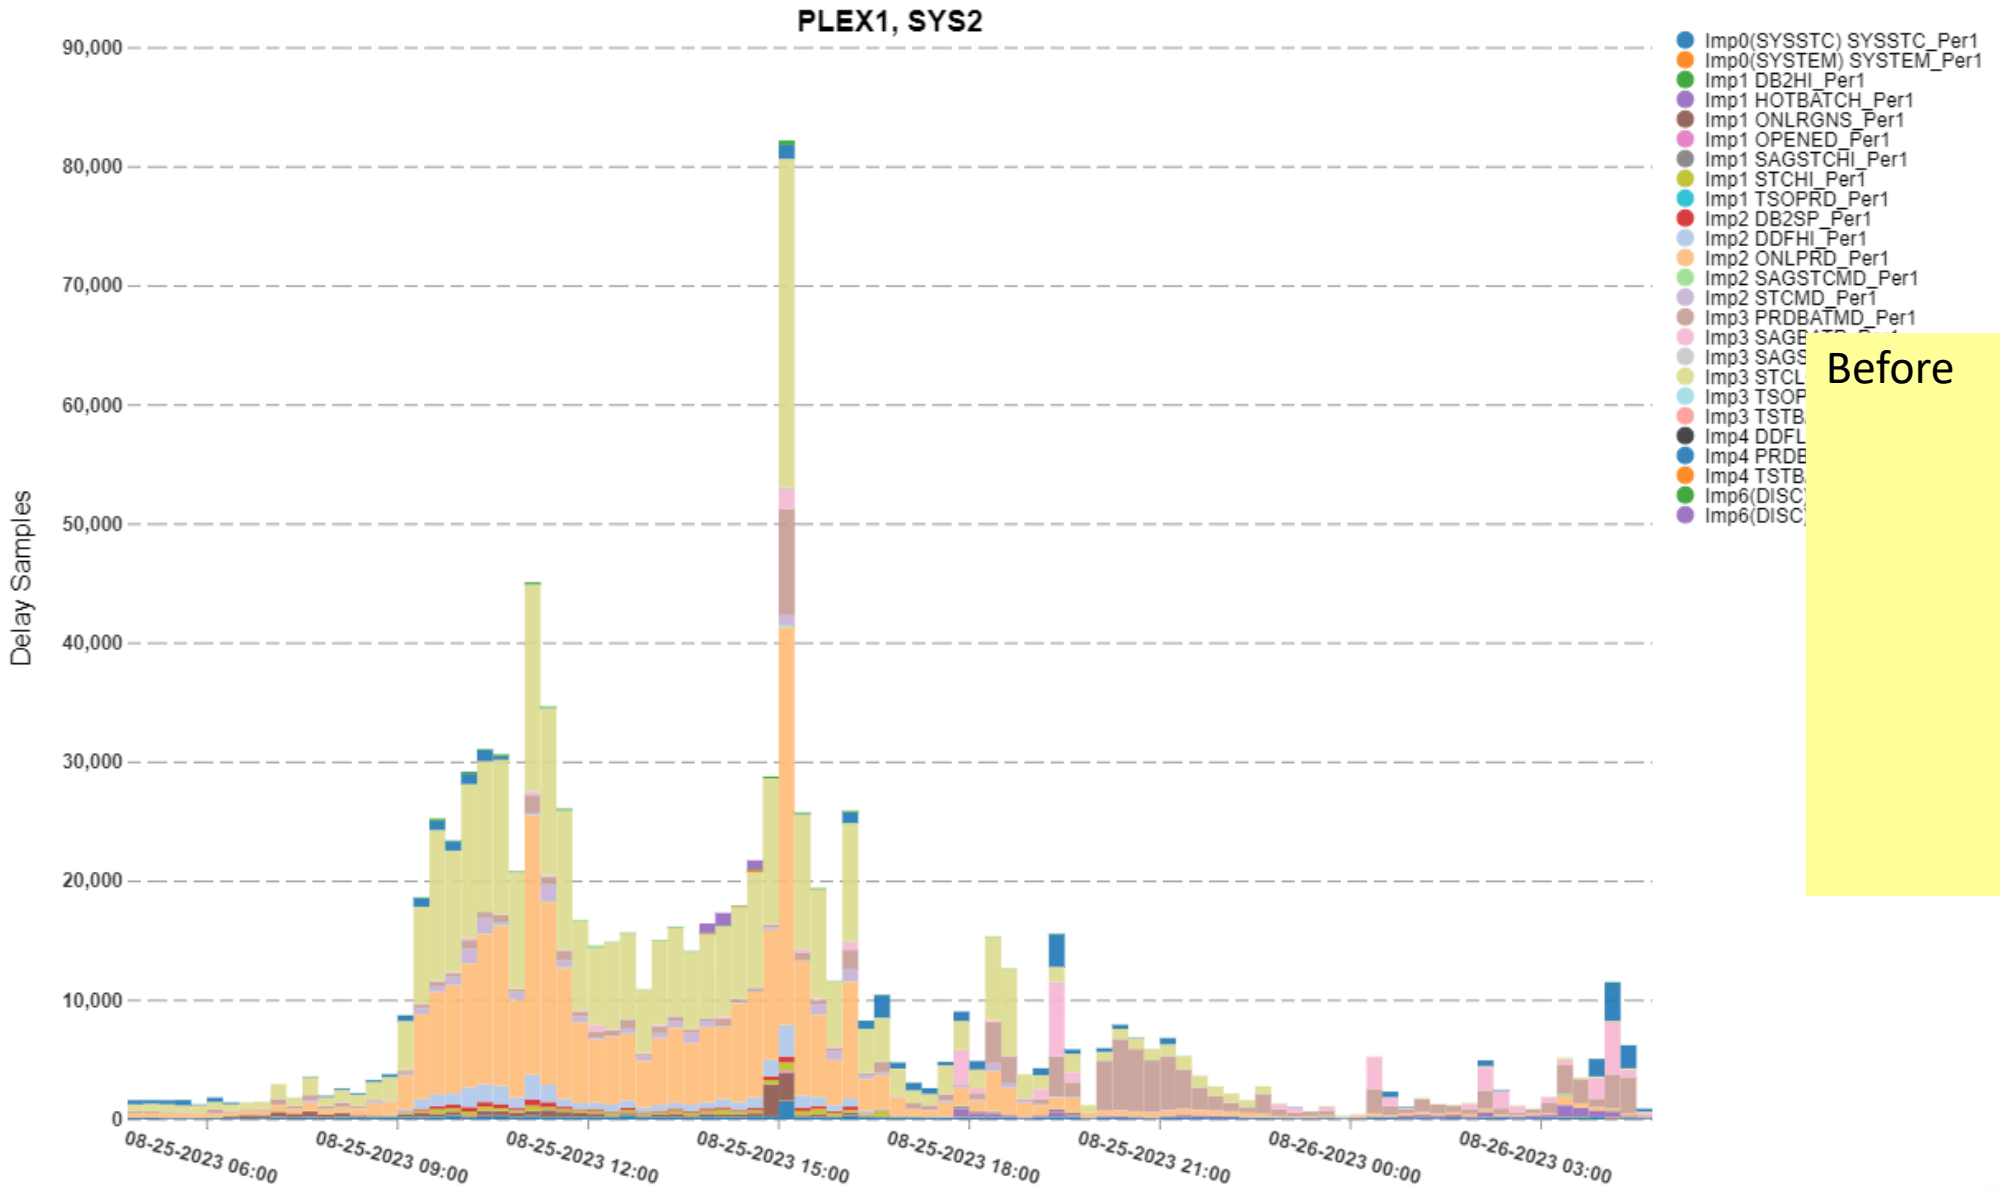

Before

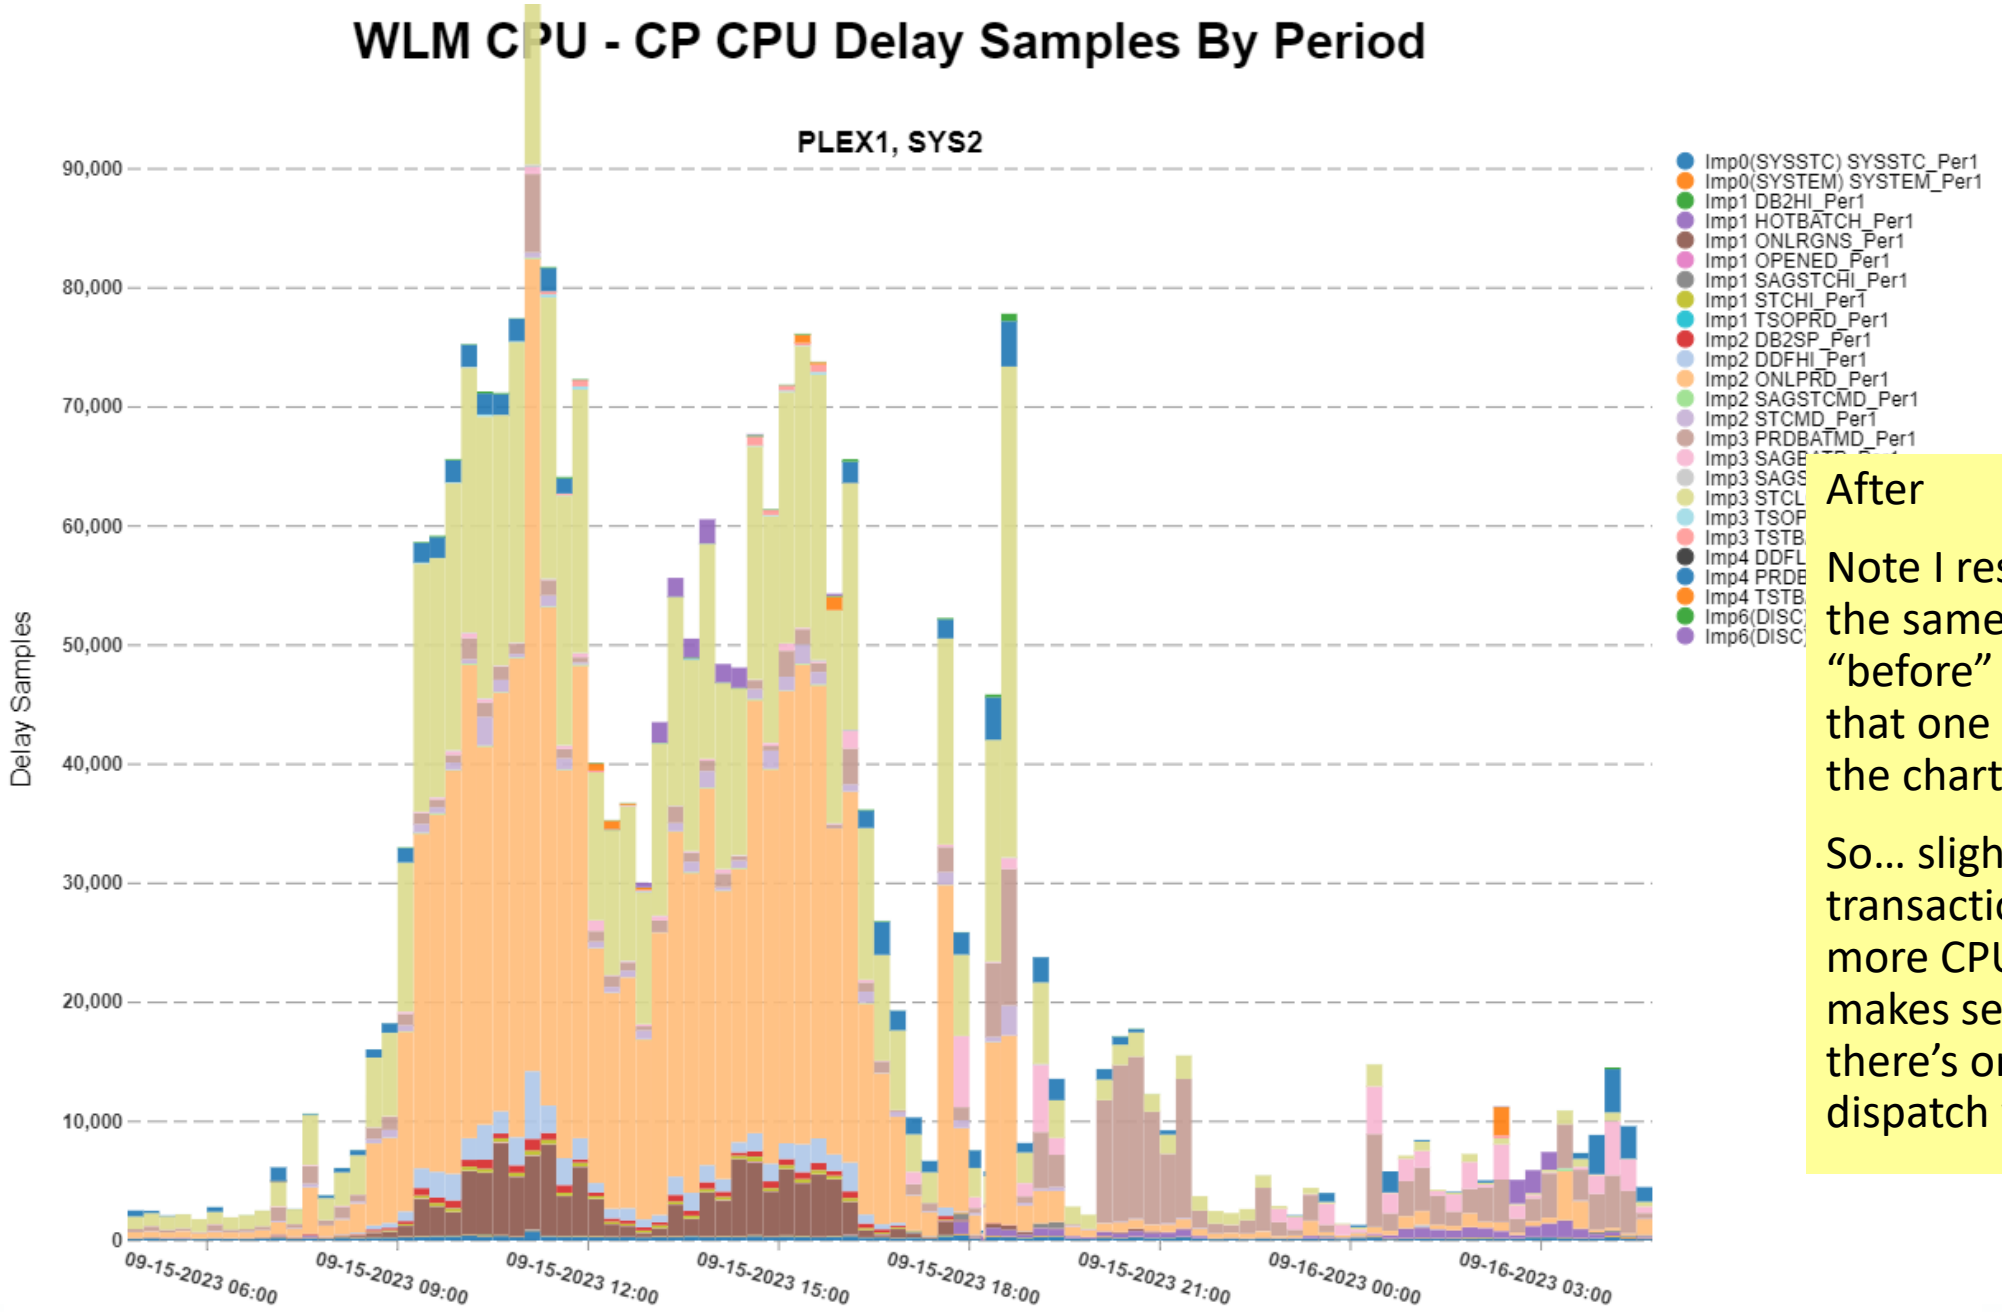

After

Note I rescaled this to the same scale as “before” which is why that one interval is off the chart.

So… slightly lower transaction load, but more CPU delay. This makes sense given that there’s only 3 CPs to dispatch work on.

Long-term (R4HA) Average MSUs for All LPARs

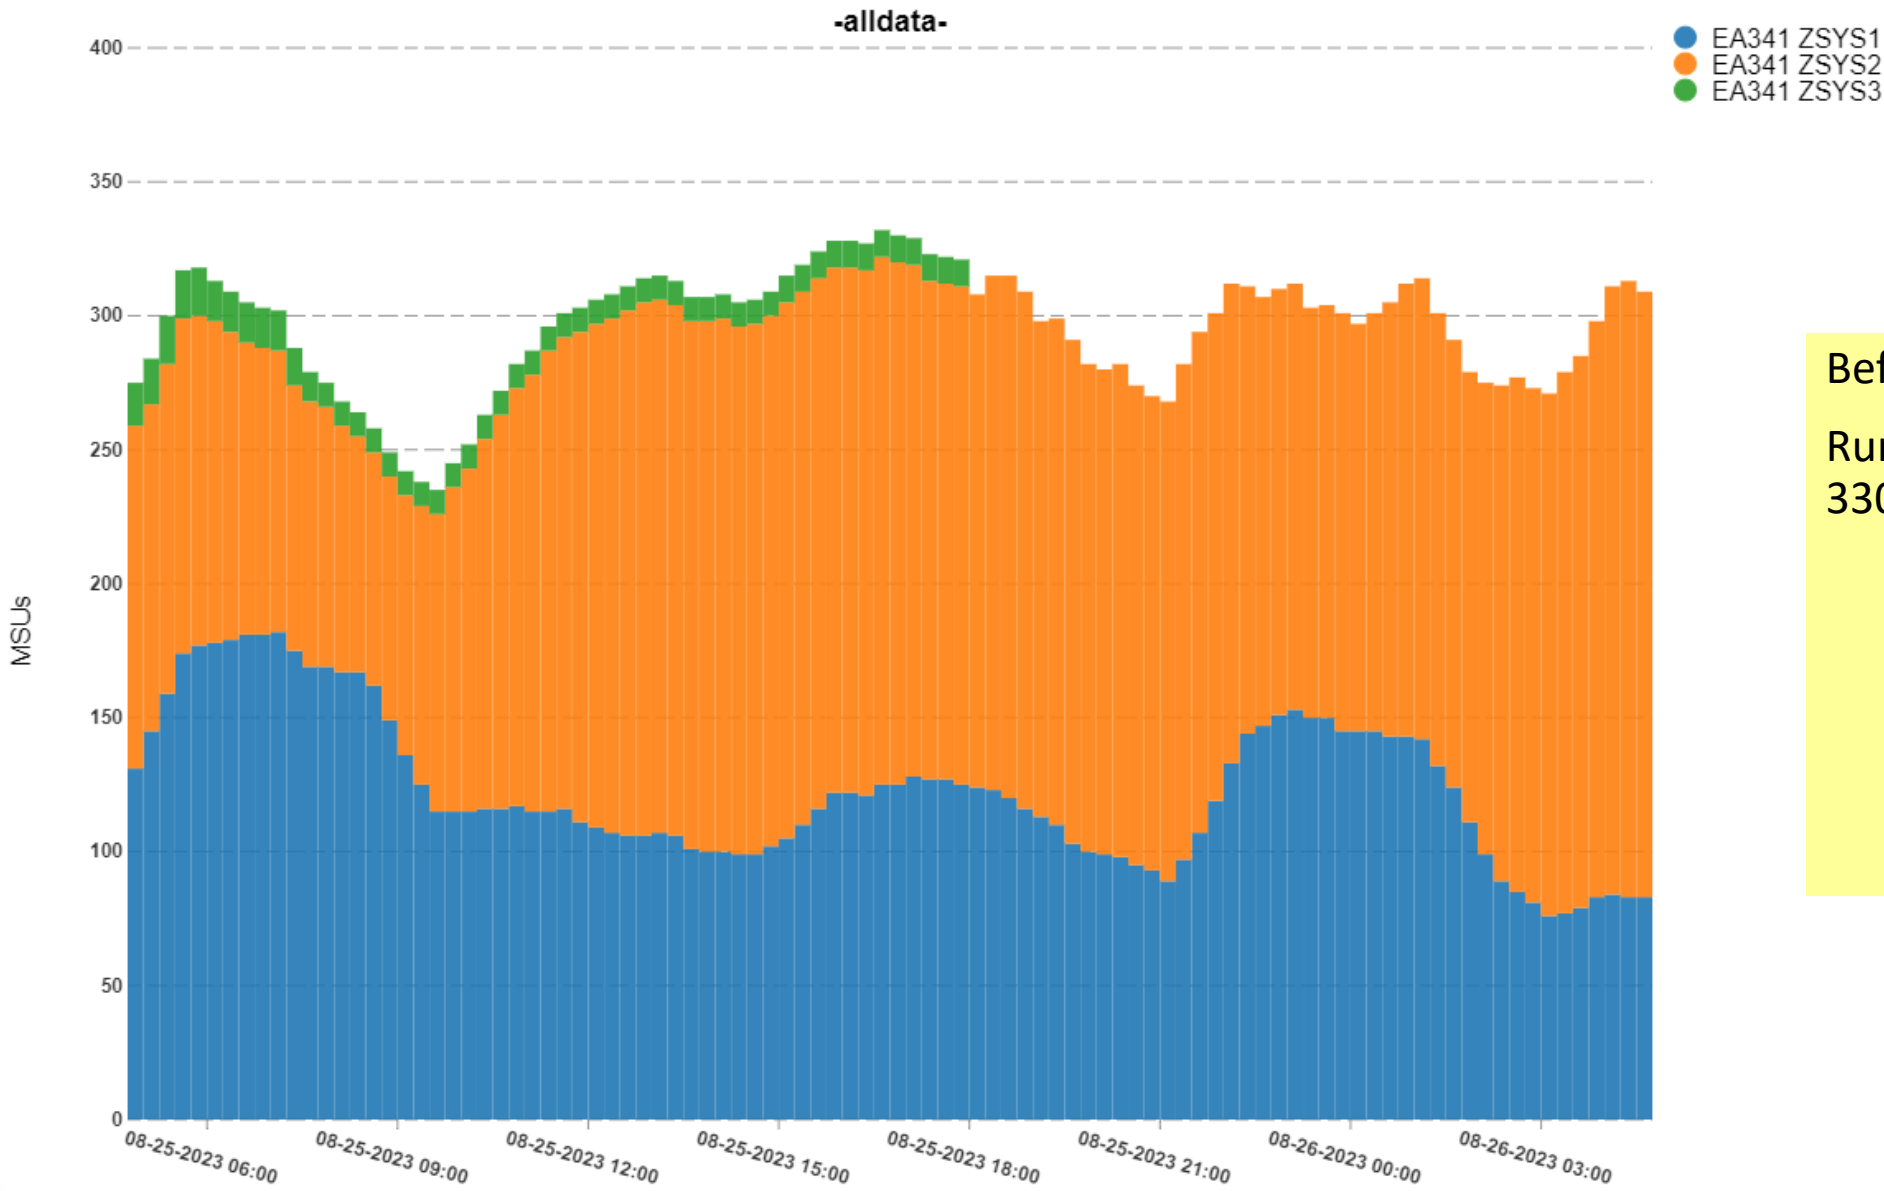

Before

Running below about 330 MSUs

Long-term (R4HA) Average MSUs for All LPARs

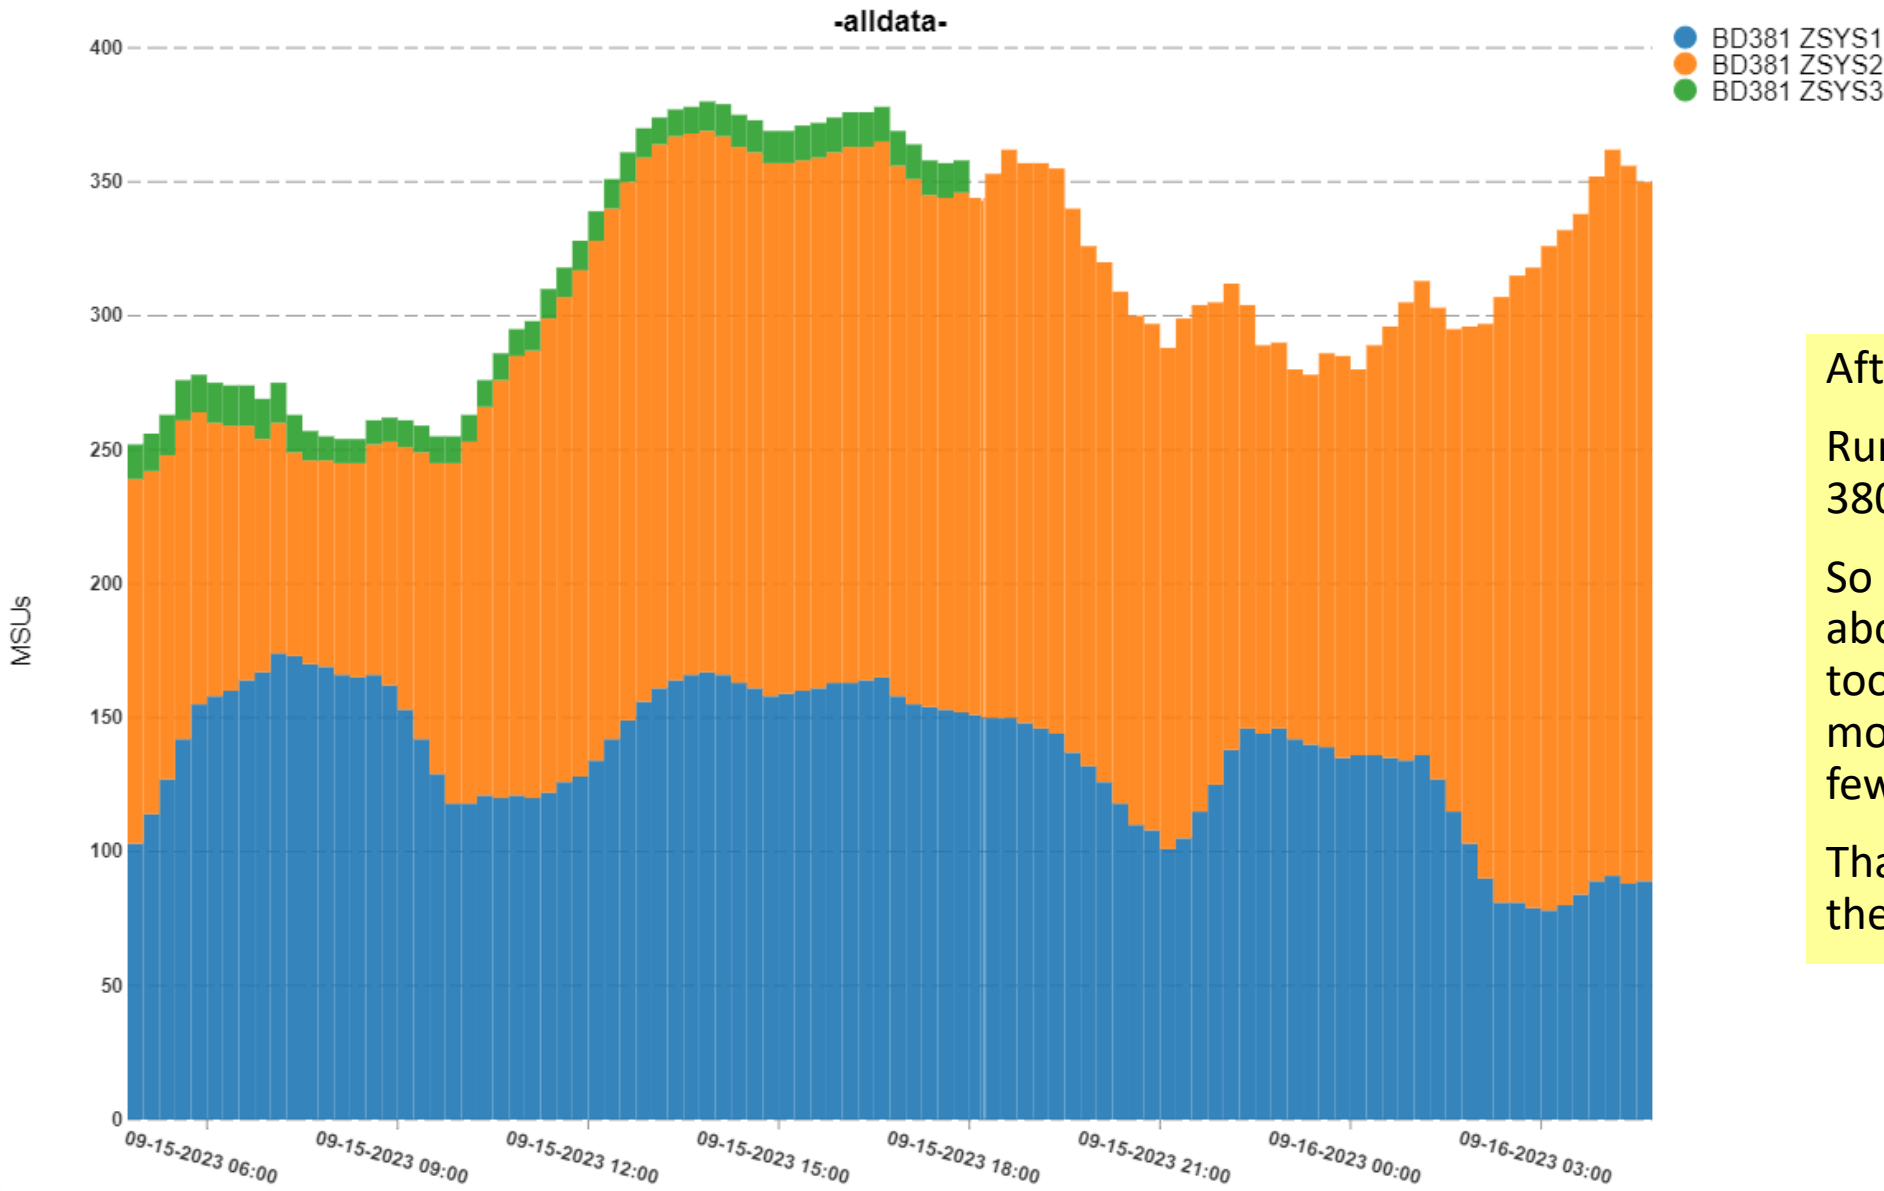

After

Running around 370 – 380 MSUs

So it appears that to do about the same work it took upwards of 15% more MSUs on fewer/faster engines.

That “could” impact their software bill.

Important notes about this example

- I didn't try to examine all the workloads to identify which ones exactly went up vs. down
- The chosen before/after days were a couple of weeks apart, so there might have been application changes
  - But very similar results for comparing days just 1 week apart
- WLM did mostly maintain response times for most important workloads
- **I don't believe this is a z14 to z16 issue**
  - In fact, I would have assumed the z16 could have faired better going to fewer/faster
  - Unfortunately, not all counters were enabled on the z16
- **I do believe this is a more/slower to fewer/faster issue**
  - Largest LPAR went from 2H, 2M, 1L to 1H, 1M, 1L
    - Much more sharing of CPs both between and within LPARs

LPAR Configuration Parameters

Observed Problem

- CPU Measurement Facility Counter sets not enabled
  - Basic
  - Problem state
  - Crypto activity
  - Extended counters
  - Coprocessor group
  - SMT diagnostics

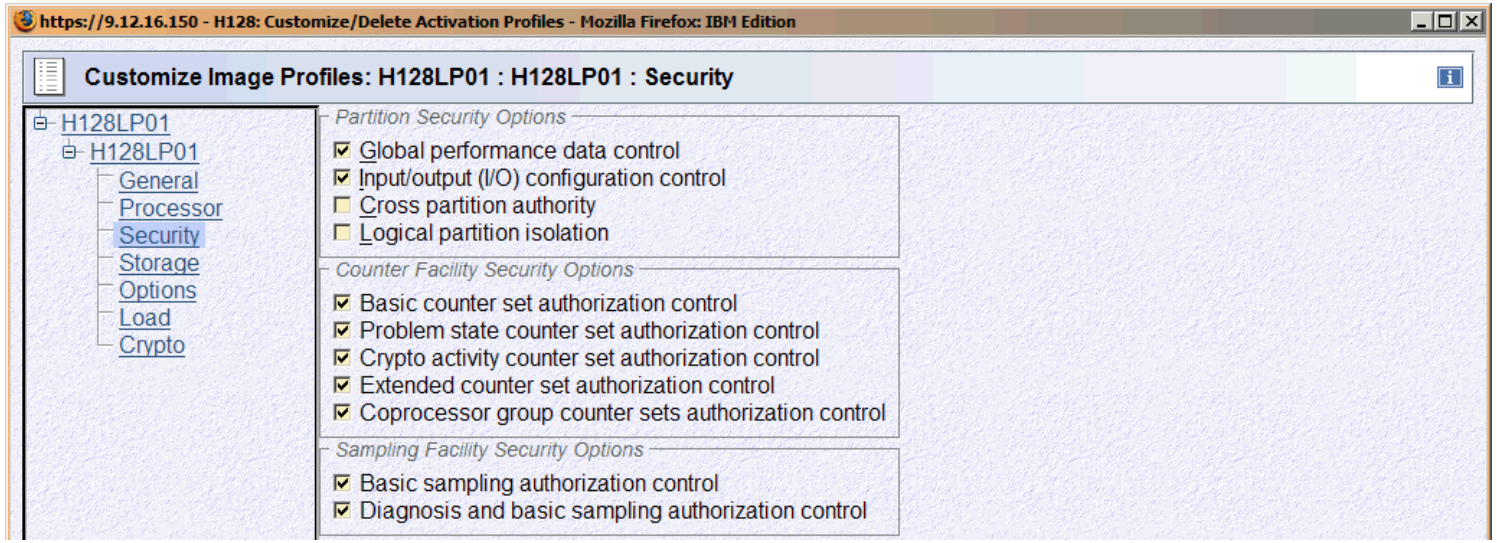

- Basically, enable them all to get all the SMF 113 measurements
  - Can provide detailed CPU hardware metrics that can be useful, especially around upgrades (would have been nice to have that data in the prior example)
  - No measurable impact by having these enabled
- Sampling more rarely used, but no harm in enabling it too

Observed Problem (rare)

- Global Performance Data Control Authority disabled
  - This lets LPARs see all other LPARs' CPU utilization
  - Generally recommended and usually enabled
    - Except sometimes in outsourced environments
- With this disabled an LPAR doesn't know how many other LPARs are on the machine or how much of the capacity that they are using
  - Can impact WLM's ability to manage low-pool processors as expected
  - Can make it difficult to determine the reason for certain performance problems
- Don't make your (or WLM's) job harder!

Summary

- Don’t forget to enable all CPU MF Counter sets
- Don’t starve your LPARs or address spaces for memory
- Don’t use weights that leave important LPARs at risk from being stolen from
- Don’t under-define the number of logical CPs: with HiperDispatch too many logical CPs is much less of a problem than too few
- Don’t assume similar sized machine with fewer/faster engines is going to perform better than one with more/slower engines