z/OS Performance Spotlight – Some Top Things You May Not Know

aka Peter and Scott's Tips and Tidbits

Peter Enrico & Scott Chapman

Enterprise Performance Strategies, Inc.

performance.questions@epstrategies.com

z/OS Performance Education, Software, and Managed Service Providers

Creators of Pivotor®

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

- During this session, Peter Enrico and Scott Chapman will discuss a variety of z/OS performance measurement, analysis, and tuning techniques that may not be commonly known or are not often discussed.

- The key objective of this presentation is to provide the attendee with information they can bring back to their shop and conduct some analysis or tuning exercises. A secondary objective of this session is to help the attendee learn more about the z/OS environment, and how things work. This session is sure to be highly educational!

EPS: We do z/OS performance…

- We are z/OS performance!

- Pivotor
  - Performance reporting and analysis of your z/OS measurements
  - Example: SMF, DCOLLECT, other, etc.
  - Not just reporting, but cost-effective analysis-based reporting based on our expertise

- Performance Educational Workshops (while analyzing your own data)
  - Essential z/OS Performance Tuning
  - Parallel Sysplex and z/OS Performance Tuning
  - WLM Performance and Re-evaluating Goals

- Performance War Rooms
  - Concentrated, highly productive group discussions and analysis

- MSU reductions
  - Application and MSU reduction

z/OS Performance workshops available

During these workshops you will be analyzing your own data!

- Essential z/OS Performance Tuning
  - March 20-24, 2023
- Parallel Sysplex and z/OS Performance Tuning
  - May 2-3, 2023
- WLM Performance and Re-evaluating Goals
  - October 2-6, 2023
- Also… please make sure you are signed up for our free monthly z/OS educational webinars! (email contact@epstrategies.com)

EPS presentations this week

What	Who	When	Where
PSP: z/OS Performance Tuning – Some Top Things You May Not Know	Peter Enrico Scott Chapman	Tue 13:15	Strand 12A
z/OS WLM – Revisiting Goals Over Time	Peter Enrico	Tue 16:00	Empire C
Sharing CPUs: How z/OS & PR/SM Manage Logical & Physical Processors	Scott Chapman	Wed 08:00	Empire C
Observability Shootout	Scott & other ISVs	Wed 16:00	Empire C
I/O, I/O It's Home to Memory We (Should) Go	Scott Chapman	Fri 09:15	Strand 12A

Like what you hear today?

- Free z/OS Performance Educational webinars!
  - Have been on hiatus for a couple of months but should be coming back soon
  - Let us know if you want to be on our mailing list for these webinars

- If you want a free cursory 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

Like what you see?

- The z/OS Performance Graphs you see here come from Pivotor™

- If you just a free cursory 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!

Pivotor – Intelligent Reporting

- Pivotor is our data reporting tool & service designed specifically for z/OS performance reporting
  - Designed and used by z/OS performance experts
  - Processes data from SMF, DCOLLECT, and customer sources
  - Contains hundreds of z/OS performance reports “out of the box”
  - Designed to be easy to use and manage
  - Reports are organized into logical and searchable report sets
  - Features include intelligent exceptions, drill down, search, canned analysis, and so much more
  - Built in expanded helps to help foster education

Comprehensive Report Sets for Immediate Performance Analysis

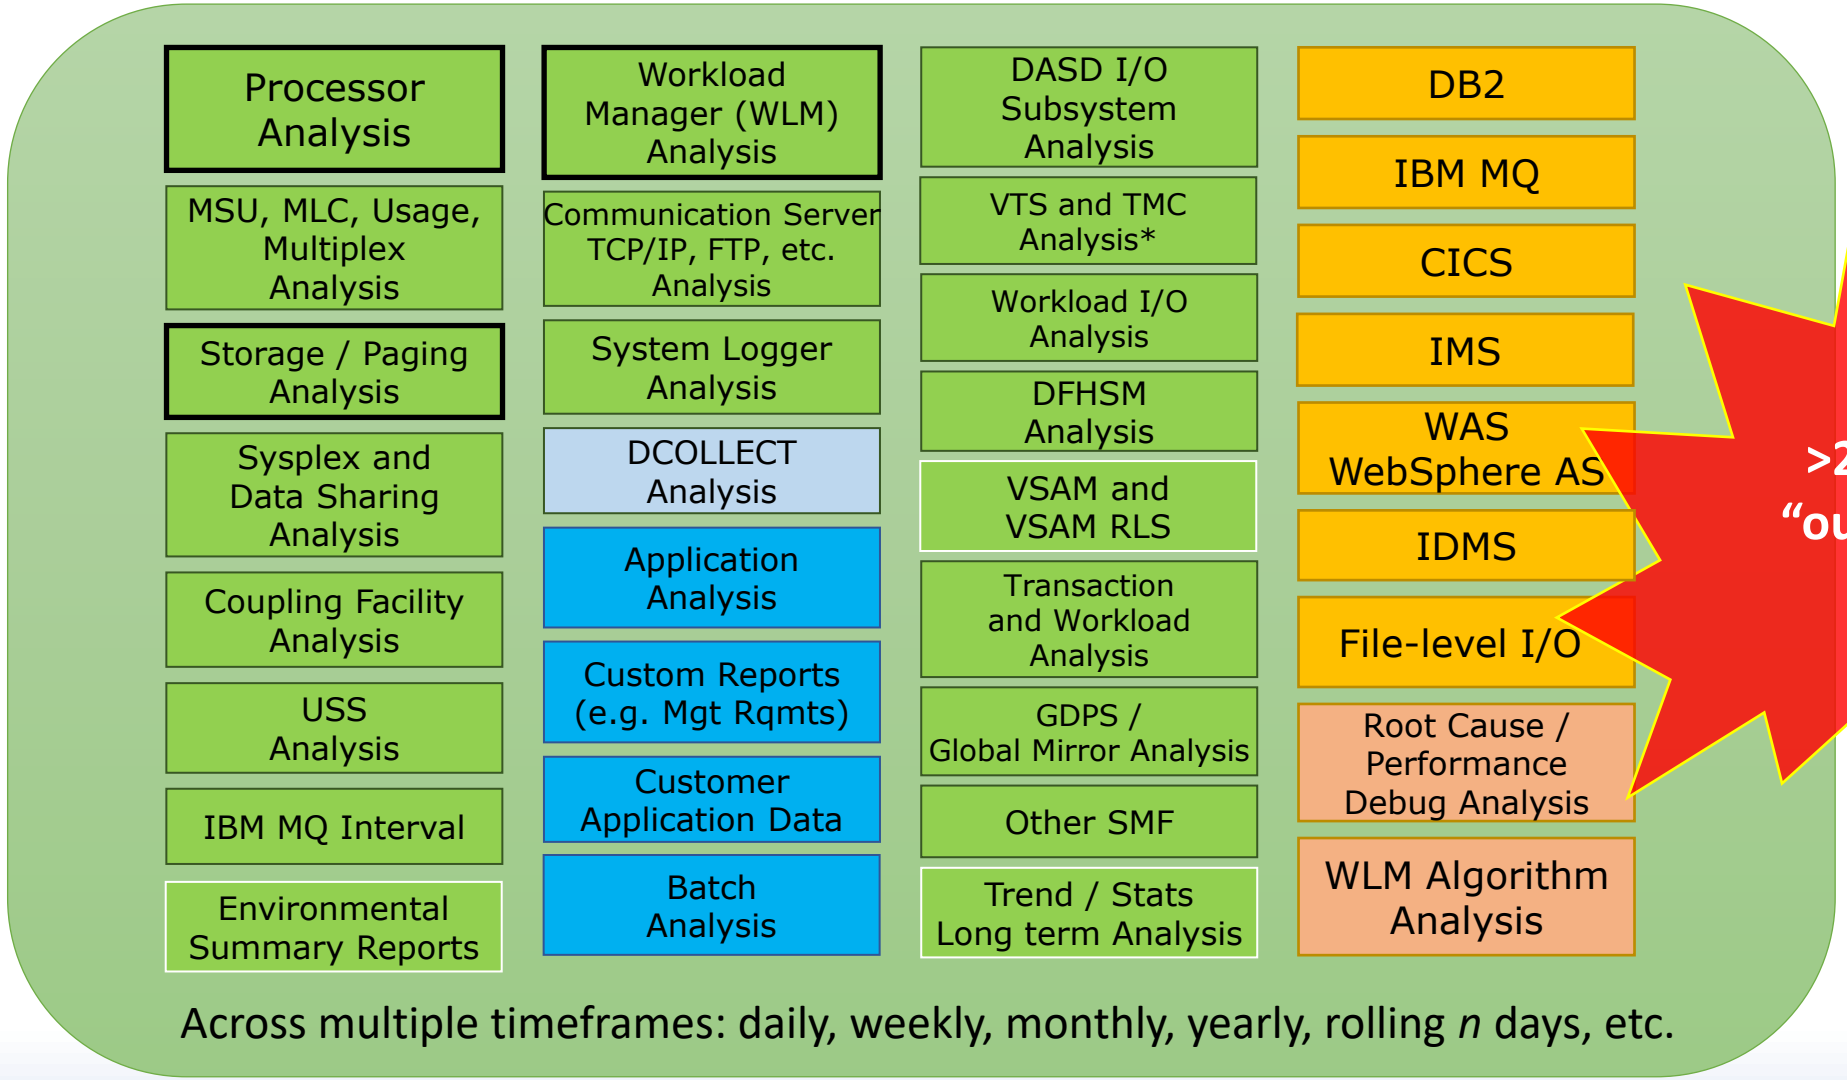

Pivotor Software as a Solution (SaaS)

- Pivotor is offered as both a SaaS or local install
- When SaaS:

SaaS Includes:

- Formal yearly cursory review / discussion
- Ability to ask us performance questions, or for us to look at a particular problem or concern. (support@epstrategies.com)
- We can occasionally look in on your data and performance
- We can participate in performance debug with IBM, or other vendors

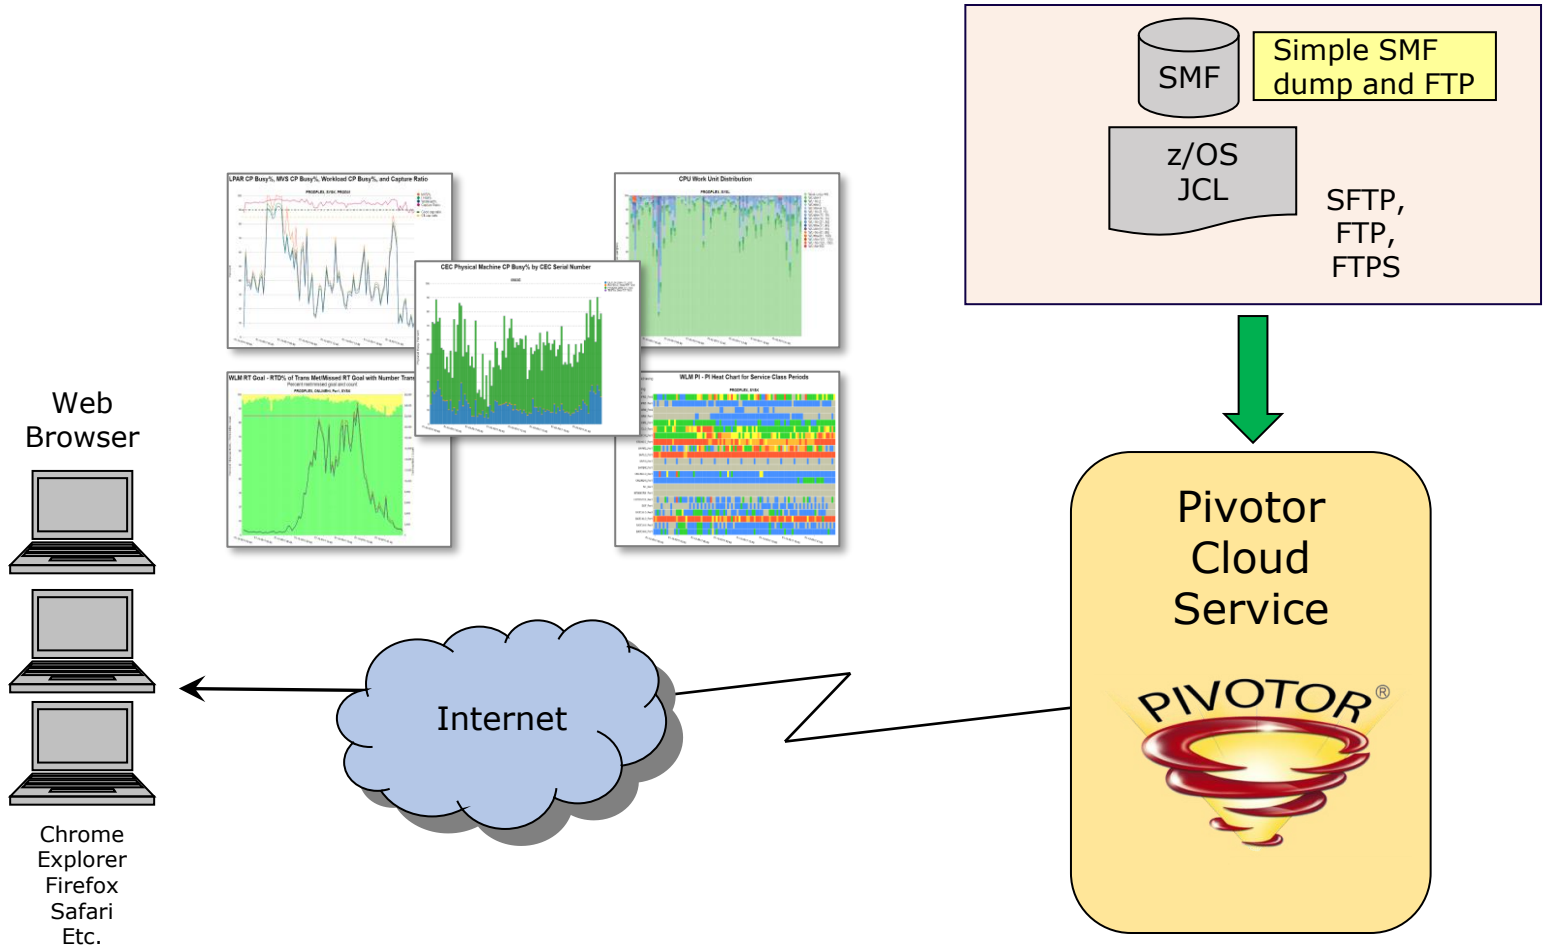

EPS

- Pivotor pricing is clear and affordable

More Free Things!

- On our web site click on Tools & Resources to access:
  - WLM to HTML Tool
    - Get your WLM policy in a useful and usable HTML format
  - Our Presentations
    - Lots of great content from the past few years (now even easier/faster to access!)

https://www.epstrategies.com
https://www.pivotor.com/
(Same site behind both URLs)

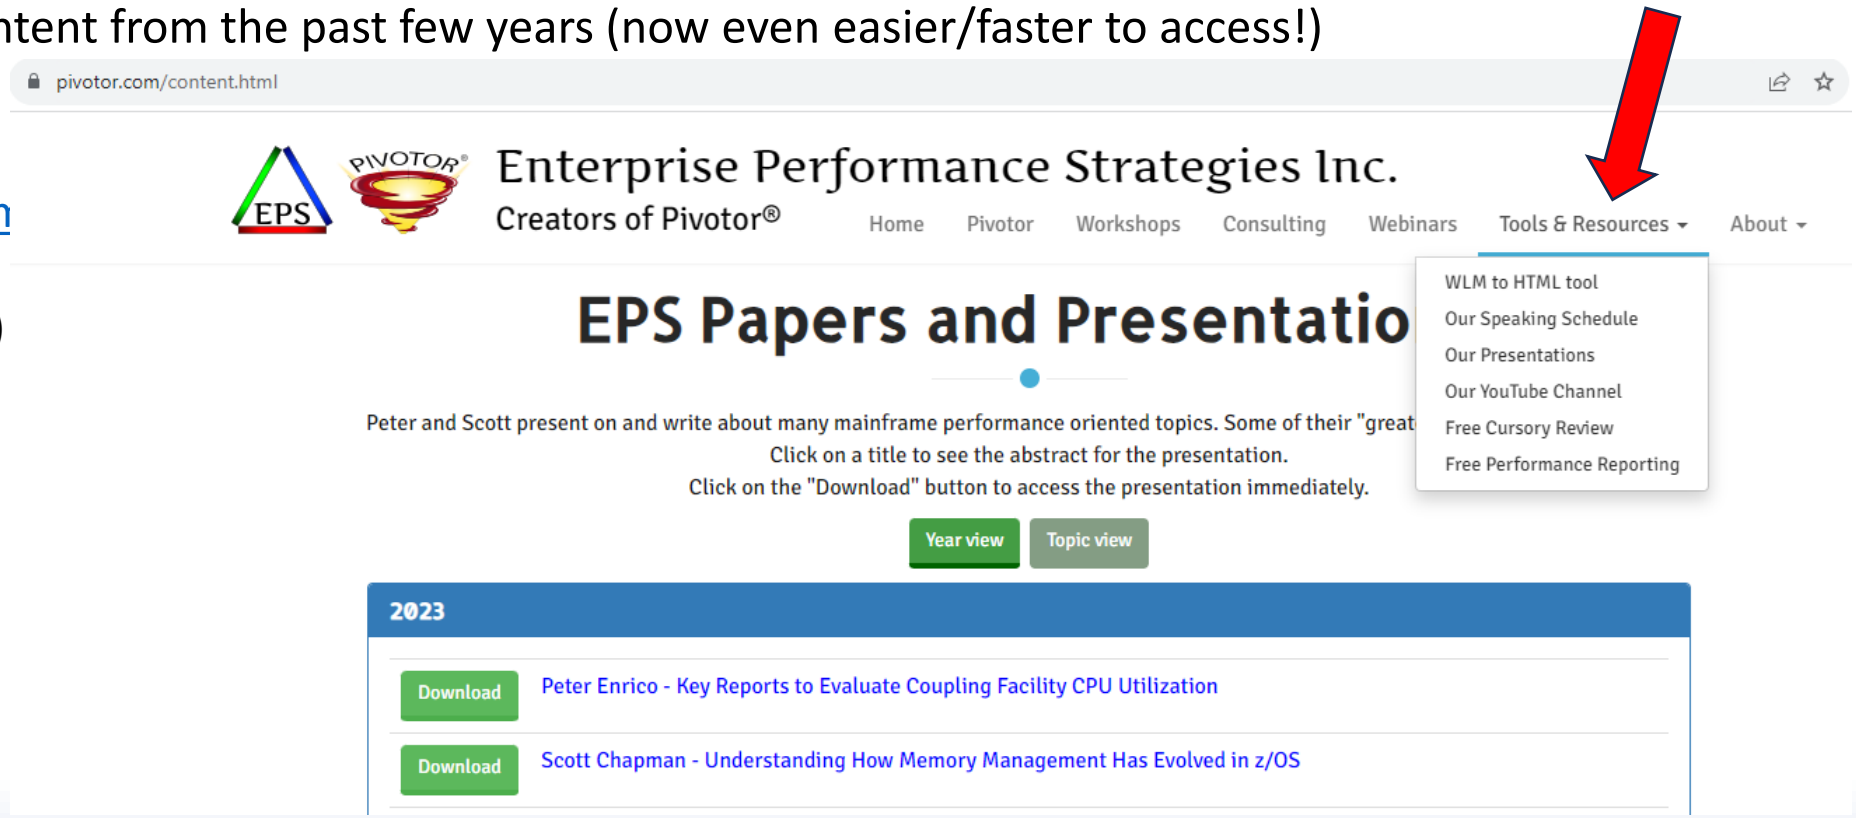

Agenda

- Emerging Areas of Interest
  - z16 Processor Cache
  - CPENABLE and z/OS 3.1
  - Implicit CPU Protection in z/OS 3.1
  - Large memory should mean less I/O?
  - How will AI change what we do?
- Short Reminders of Ongoing Opportunities
  - Re-evaluating goals (see Peter's presentation)
  - SuperPAV
  - Logger?
  - XCF transport class simplification
  - Record the 98s and 99s
  - SMT
  - I/O Priority Management
  - MSO/IOC

Emerging Areas of Interest

New things coming and things we're actively keeping an eye on

z16 Processor Cache Performance

z16 Virtual Caches (slide source: IBM)

- What's different from z15
  - There is no L3 physical cache present on the cores
    - There is a new L1 Shadow Cache that will help manage syncing lines with L2
  - There is no SC chip or physical L4 Cache
    - All CPs L2 are interconnected via buses
- How Virtual Caches work
  - L2 Caches of unused cores or underutilized cores will be converted to be used as virtual caches
    - If the core becomes actives the cache will be returned
  - Virtual cache on the same CP will be seen as additional virtual L3 cache to the core
  - Virtual Cache on a different CP on the same drawer will be seen as L4 Cache

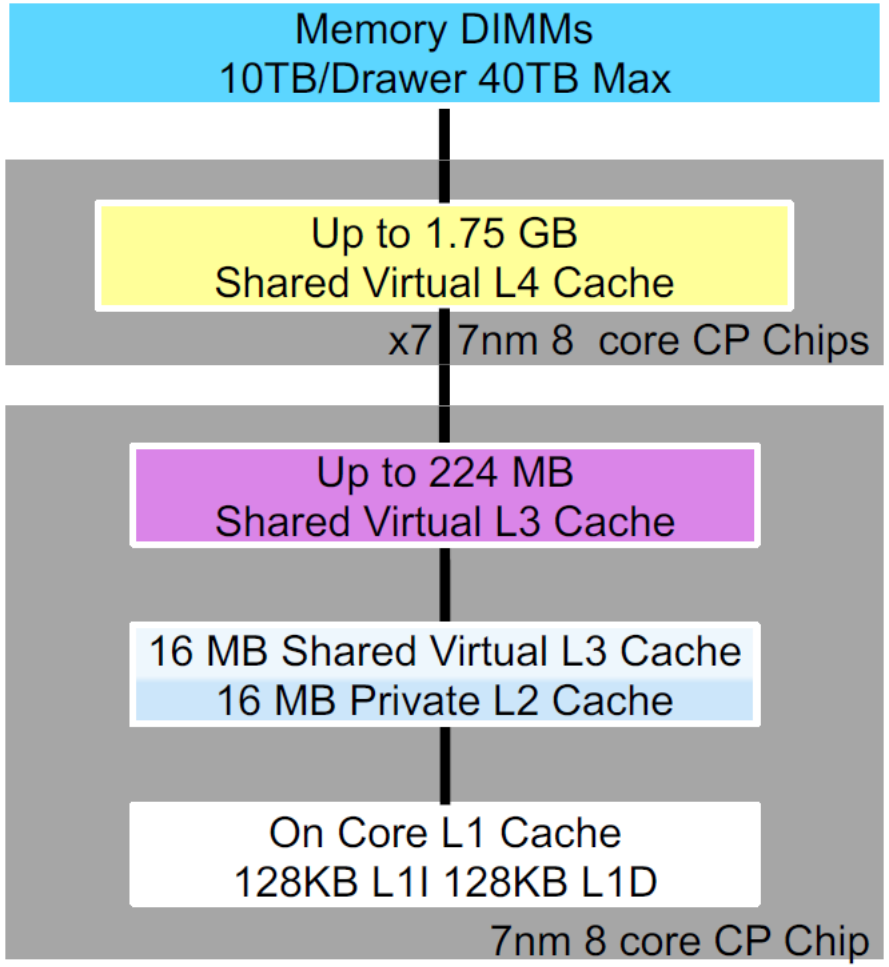

z16 Virtual Cache Provisioning

- One chip example (just to make the point)

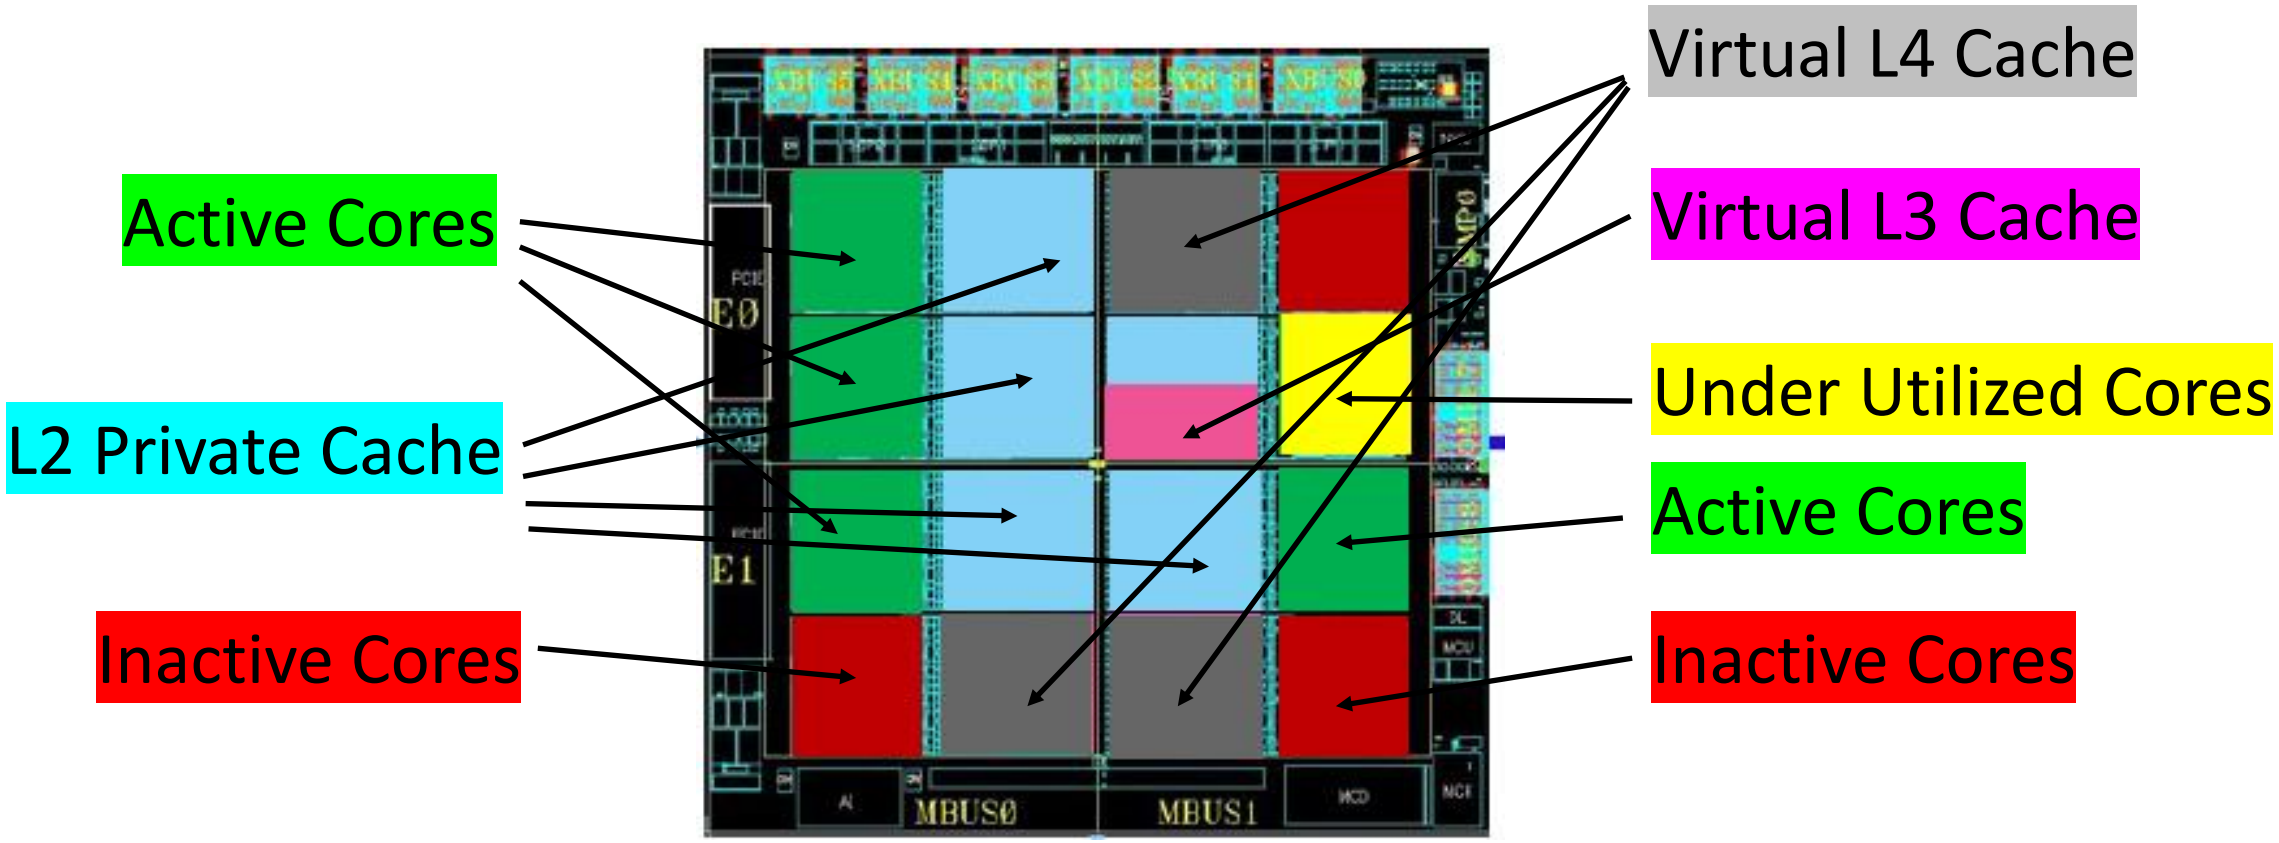

Cache Demo

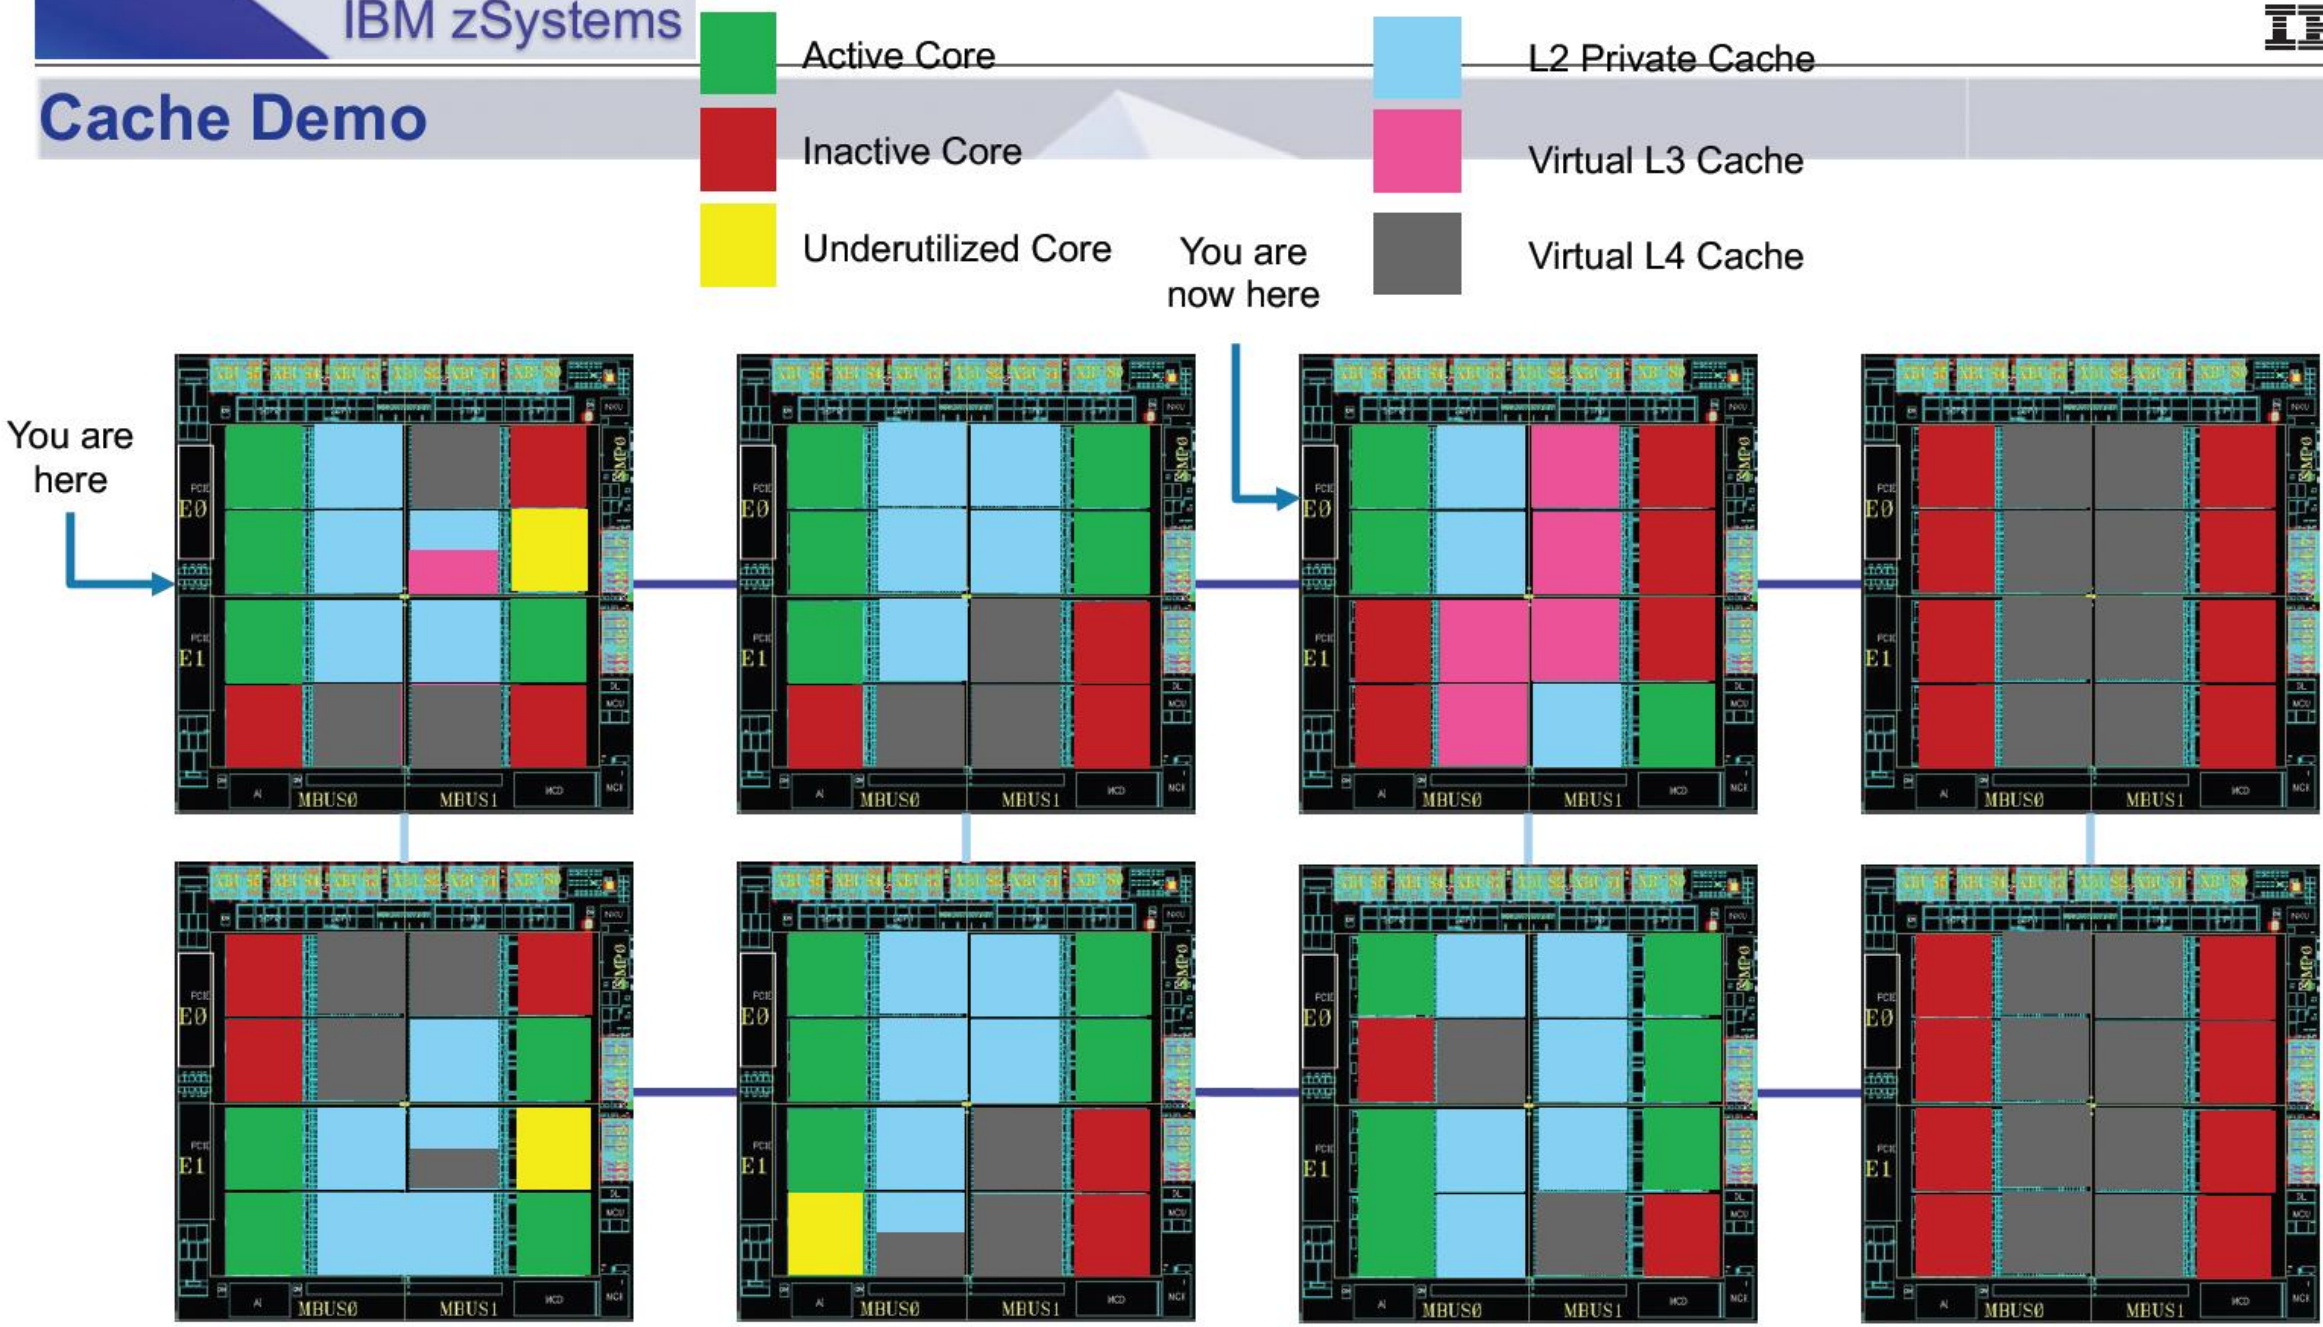

Case Study CEC LSPRs: z14 vs z16

- z14 (3906-609 M02)
- z16 (3931-606 A01)

	Processor	#CP	PCI**	MSU***	Low*	Average*	High*
z14	**3906-609**	9	8142	997	15.99	14.55	12.79
	Processor	**#CP**	**PCI****	**MSU*****	**Low***	**Average***	**High***
z16	**3931-606**	6	8006	980	14.92	14.3	13.01

z14 vs z16 SYS2 config

- z14 (3906-609 M02)
  - 9 CPs, 4 zIIPs
  - SYS2: 4 CPs, 4 zIIPs

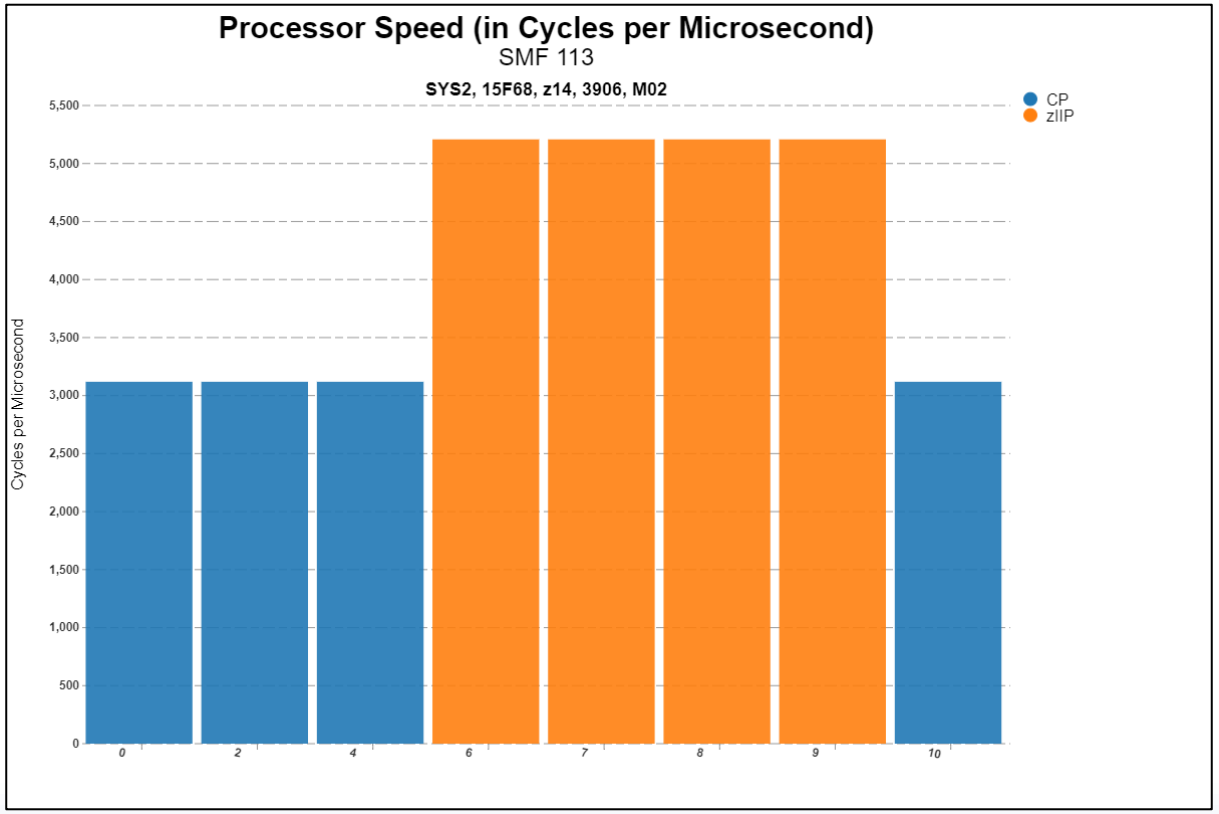

- z16 (3931-606 A01)
  - 6CPs, 4 zIIPs
  - SYS2: 4 CPs, 4 zIIPs

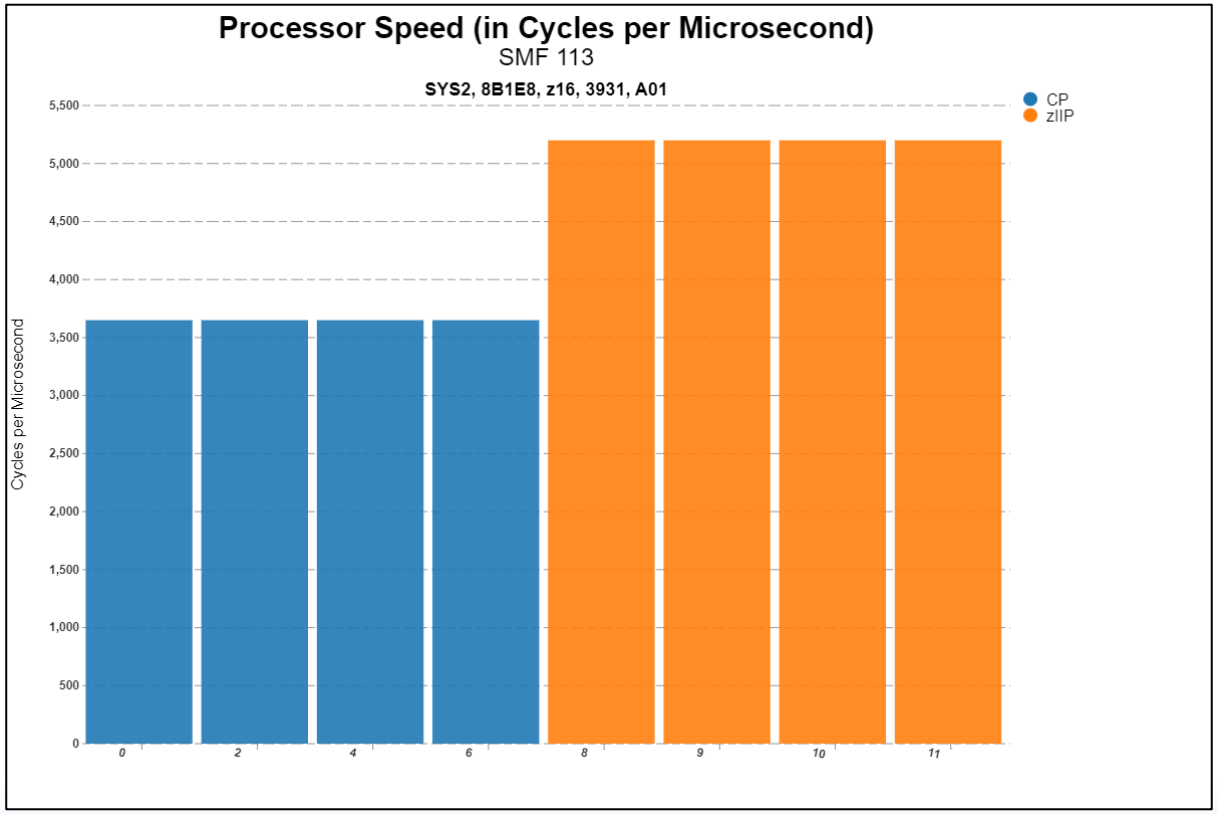

z14 vs z16 – Cache Sourcing

Notice the improved sourcing from L2 since L2 caches are much larger

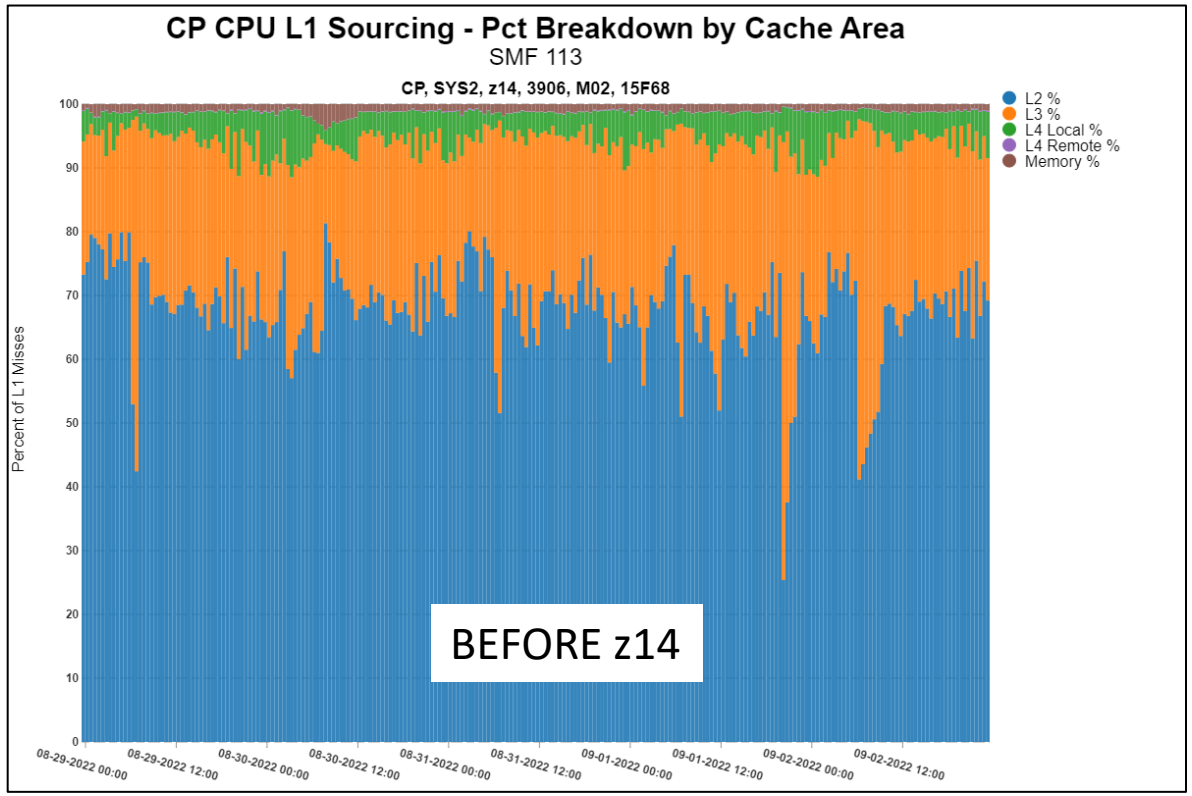

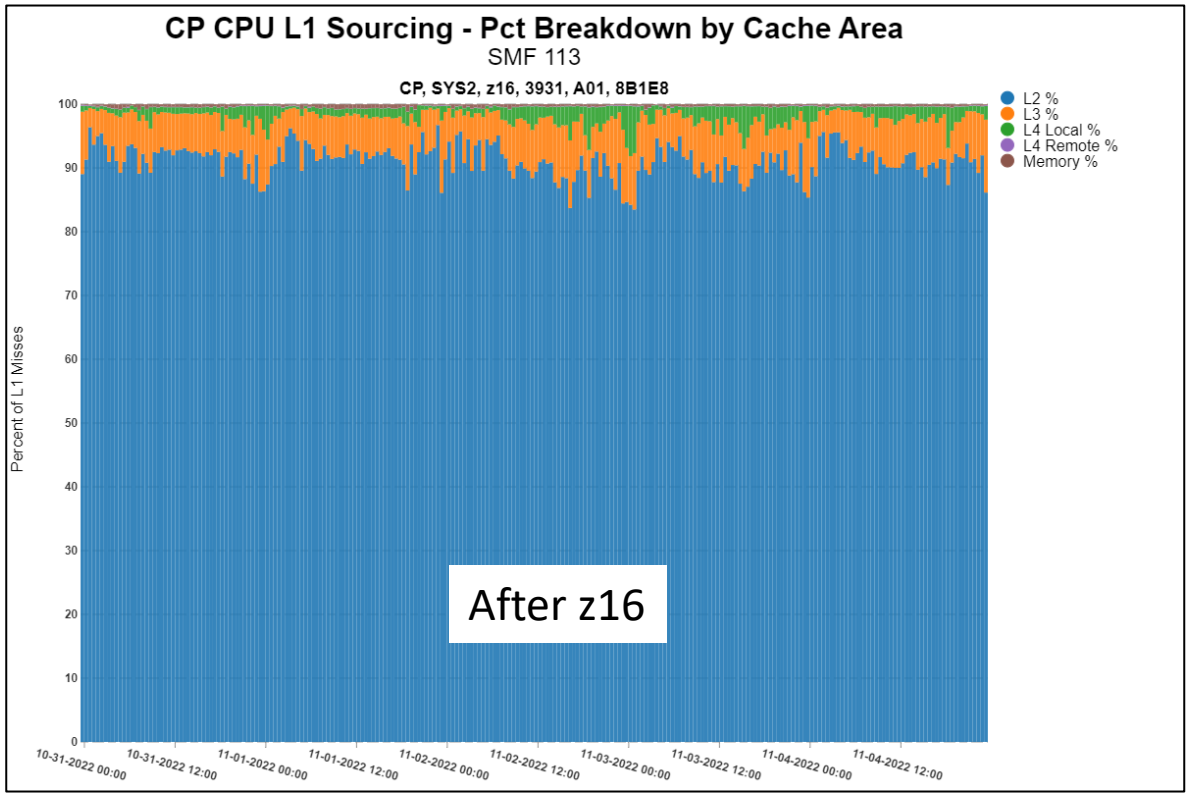

RNI – Breakdown by Cache

Notice the improved Relative Nest Intensity. Reminder, RNI is not a performance metric to be tuned, but rather a 'signature' of a customer's workloads relative to the LSPRs and machine capacity delivered.

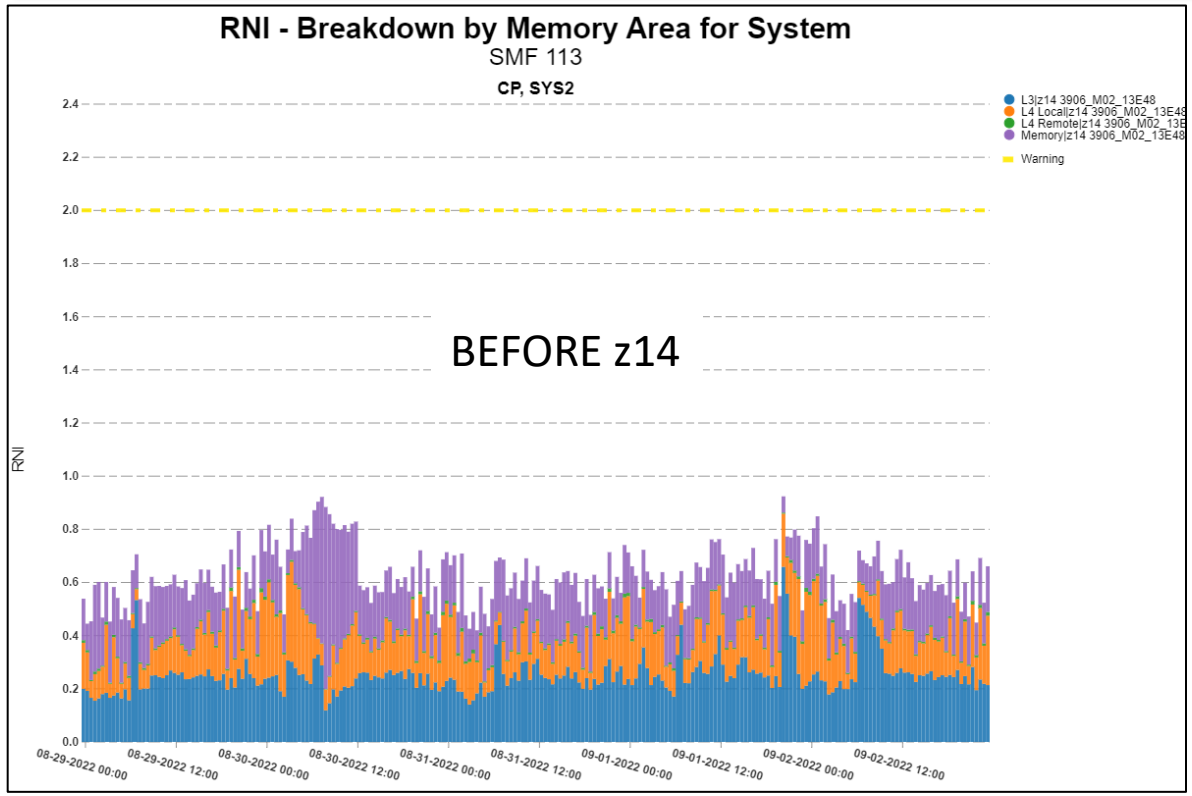

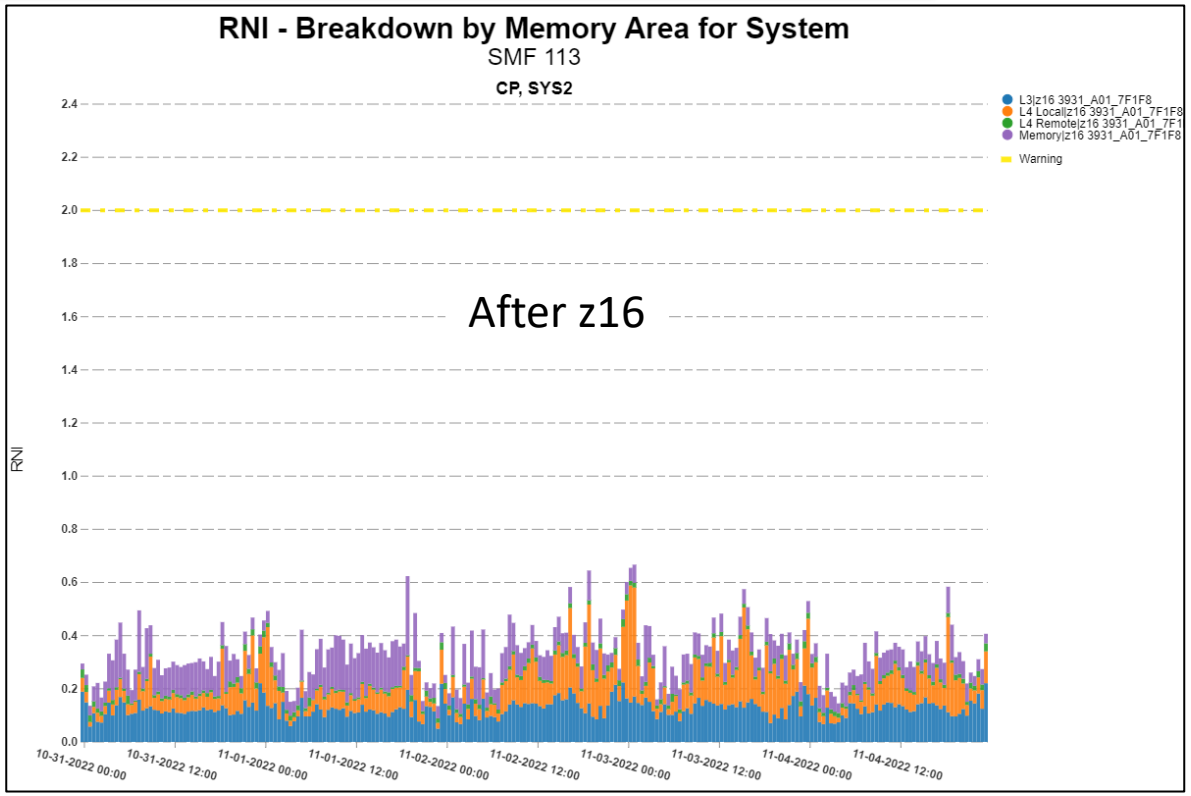

CPENABLE in z/OS 3.1

CPENABLE

- CPENABLE in IEAOPTxx sets the low and high threshold for disabling / enabling processors for handling I/O interrupts
- z13 and below recommendation is (10,30)
- On z/14 and above the recommendation is (5,15)
  - Prior to z/14 all no-work wait CPs were enabled for interrupts
  - z/14+ rely solely on WLM/SRM to set the number of CPs enabled for interrupts
- The goal of this change was to better ensure 2 CPs are enabled for handling I/O interrupts
  - Single CP enabled for I/O interrupts puts LPAR at greater risk of delaying I/O
  - Sometimes with quite problematic results – having 2 is partly risk mitigation
- We've sometimes recommended even more aggressive settings (e.g. 3,10)

I/O Interrupt Analysis

(CPENABLE=(x,y) recommended settings)

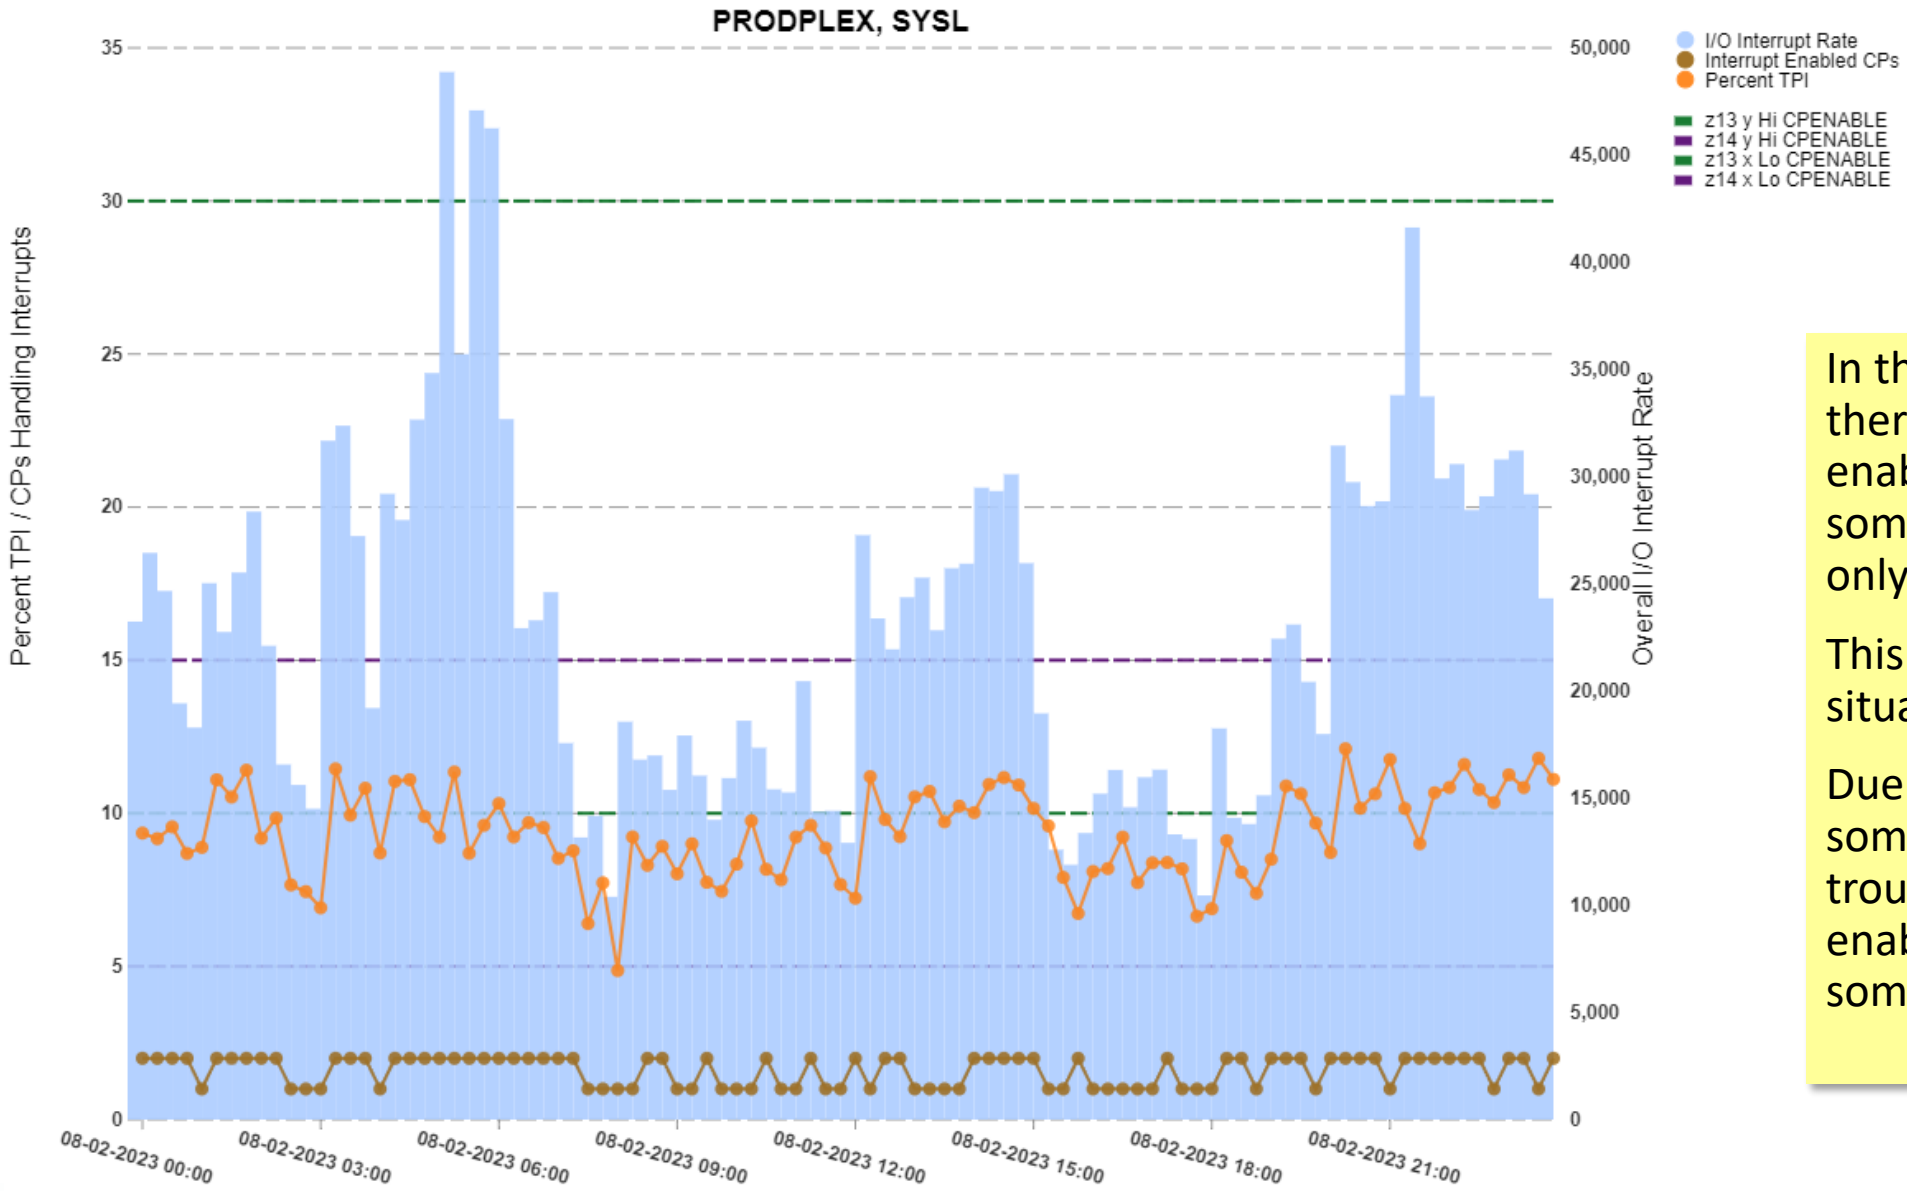

In this case, sometimes there were 2 CPs enabled for interrupts, sometimes there was only a single CP.

This is a fairly common situation.

Due to arrival patterns, some systems have trouble getting a second enabled even with something like (3,10).

CPENABLE in z/OS 3.1

- IBM SoD for 3.1 is that minimum CPs enabled will raise from 1 to 2
- New CPENABLE option of SYSTEM will take IBM's recommendation for the generation of hardware the system is running on
- Evaluation of enabled CPs will change from 20 seconds to 2 seconds

- We think this is a great change!
  - Will be able to specify CPENABLE=SYSTEM and probably not worry about it
  - A lot of I/O can happen in 20 seconds so changing to every 2 seconds (same as HiperDispatch cycle) makes sense
    - Extra path length seems like it would be pretty minimal

Implicit CPU Protection in z/OS 3.1

CPU Critical aka Long-term CPU Protection

- Long-time option in your WLM service definition
- Enabled by setting YES for CPU Critical on a Service Class
  - Must be a single-period SC and cannot be discretionary
- Ensures that the CPU Critical SC always has a dispatching priority that's greater than the DP of lower importance service class periods
- Note some small amount of lower-importance work may still get higher DP:
  - Due to promotion for locks, resource contention, etc.
  - Small consumers
- General recommendation has been to avoid this option
  - Allows WLM to make better decisions about balancing overall work throughput to best meet the goals of all work

> **Important:** The use of these options limits WLM's ability to manage the system. This may affect system performance and/or reduce the system's overall throughput.

IBM SoD for z/OS 3.1

- New option for “Implicit” Long-Term CPU Protection
  - In other words, CPU Critical without having to specify it on every SC definition
- Default is said to be “On” for importance 1 service classes
  - Optional but “Off” for importance 2 service classes
- We think “On” for importance 1 workloads is a bad default
  - Could significantly change the dispatching priority of work in the system
  - Goes against historical practices of not changing defaults that change behavior
- DP/Importance inversions are common
  - I.E. Lower Importance work running with a DP above higher importance work
  - Not all such inversions are problematic
  - Not all importance 1 work really should be importance 1

We did an analysis of 116 systems

- Covered a variety of sizes from large to small, “IPO” to “Prod”, a couple of dozen customers
- Evaluated a day’s worth of 99.6 data from each system (over 17M records)
- Came up with 2 new metrics to help understand the risk/benefit:
  - For SCPs that would be bumped down:
    Inversion Risk Ratio – relative amount of CPU that would move above the SCP
  - For SCPs that would move up in priority:
    Protection Benefit ratio – relative amount of CPU that would move below this SCP
  - Higher numbers means more potential risk/benefit
    - Can be very high if there’s a relatively large difference in the consumption of the workloads

Study findings

- 78% of systems had at least one interval with an inversion
- 39% of systems had inversions in at least 25% of their intervals
- 82% of “Inverted” Importance 1 workloads were meeting their goal

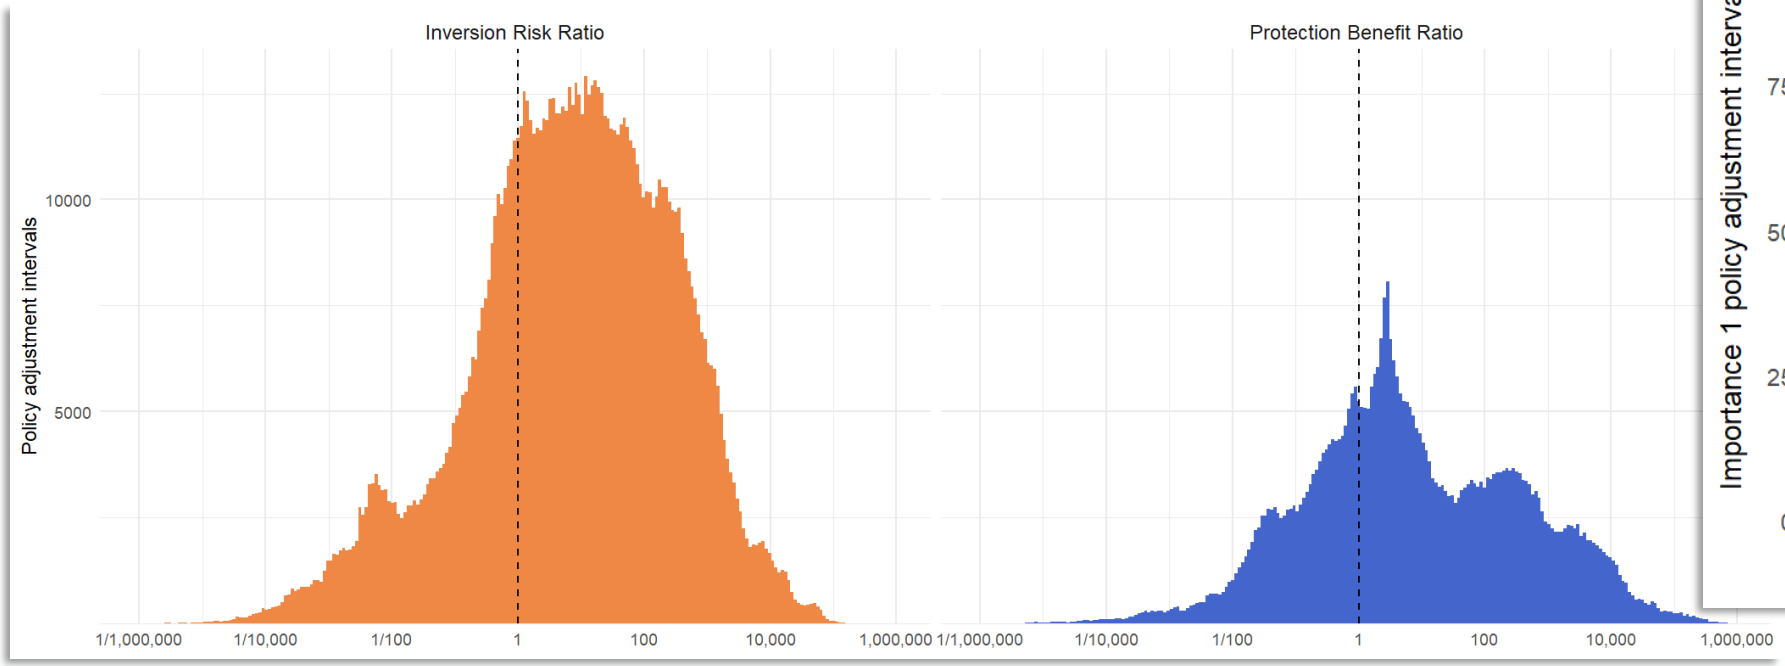

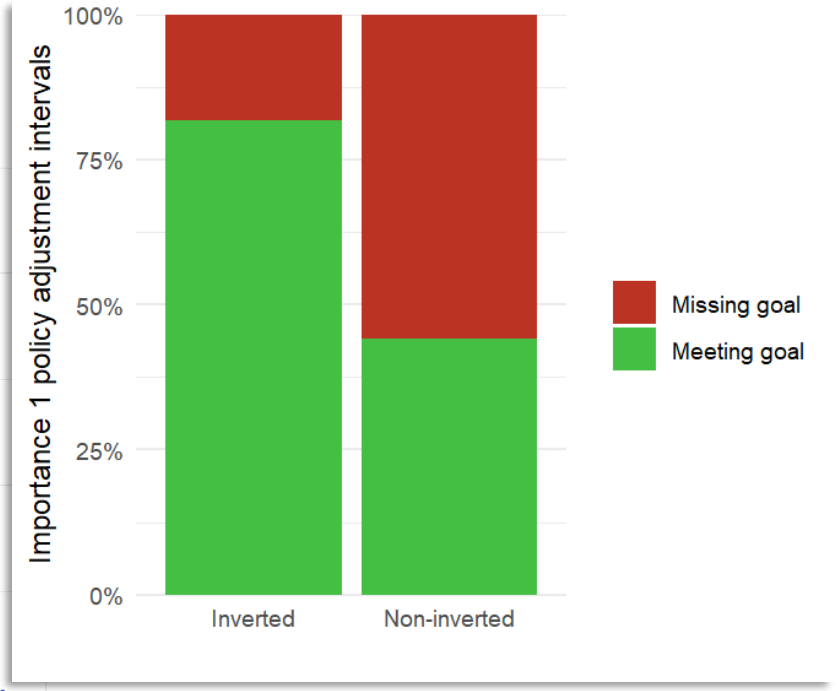

CP CPU Protection Inversion Risk Ratio

Posed by Importance 1 Workloads Running Below This

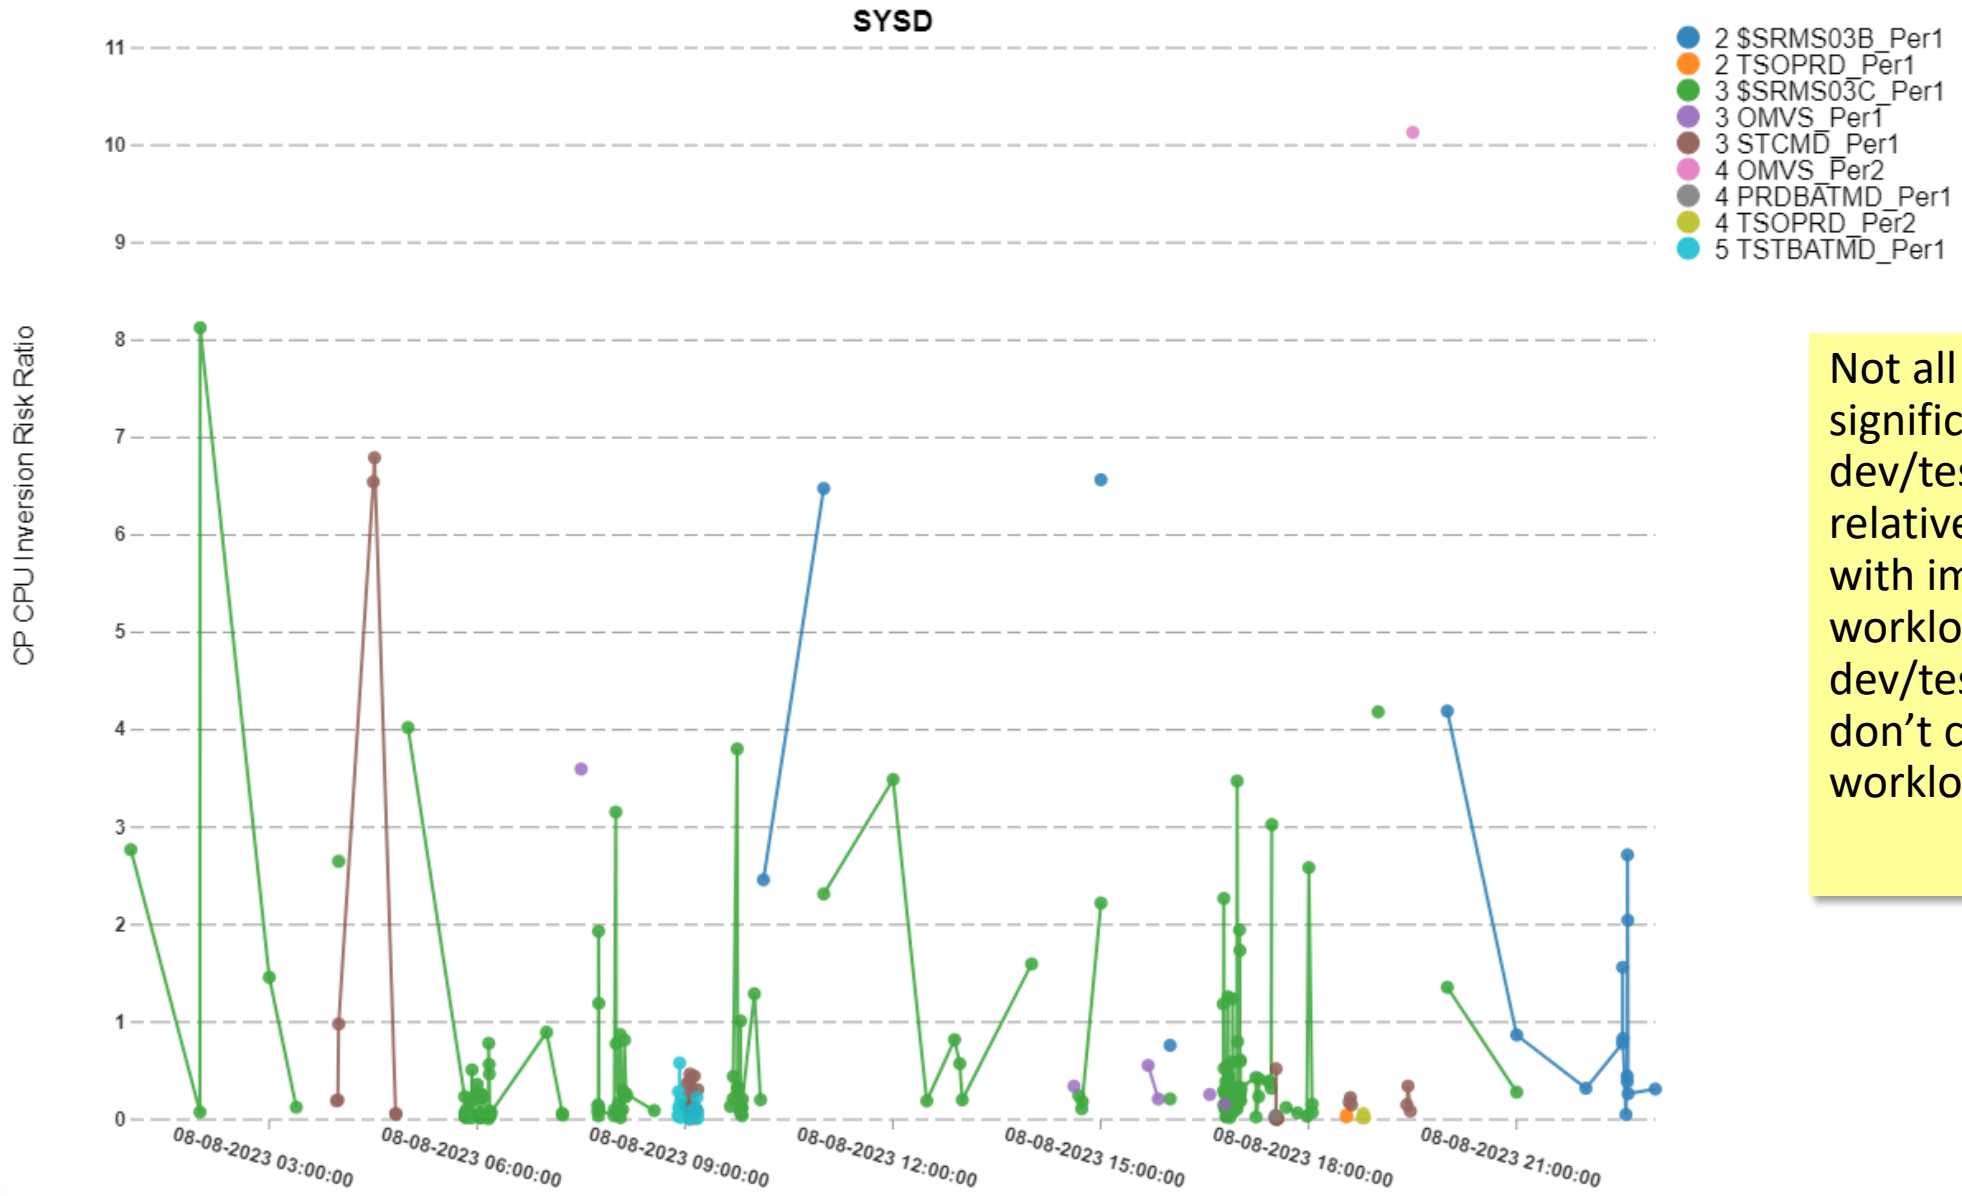

Not all systems will have significant risk. This dev/test system has relative few inversion with importance 1 workloads. And it's dev/test: maybe you don't care if those workloads suffer more.

CP CPU Protection Inversion Risk Ratio

Posed by Importance 1 Workloads Running Below This

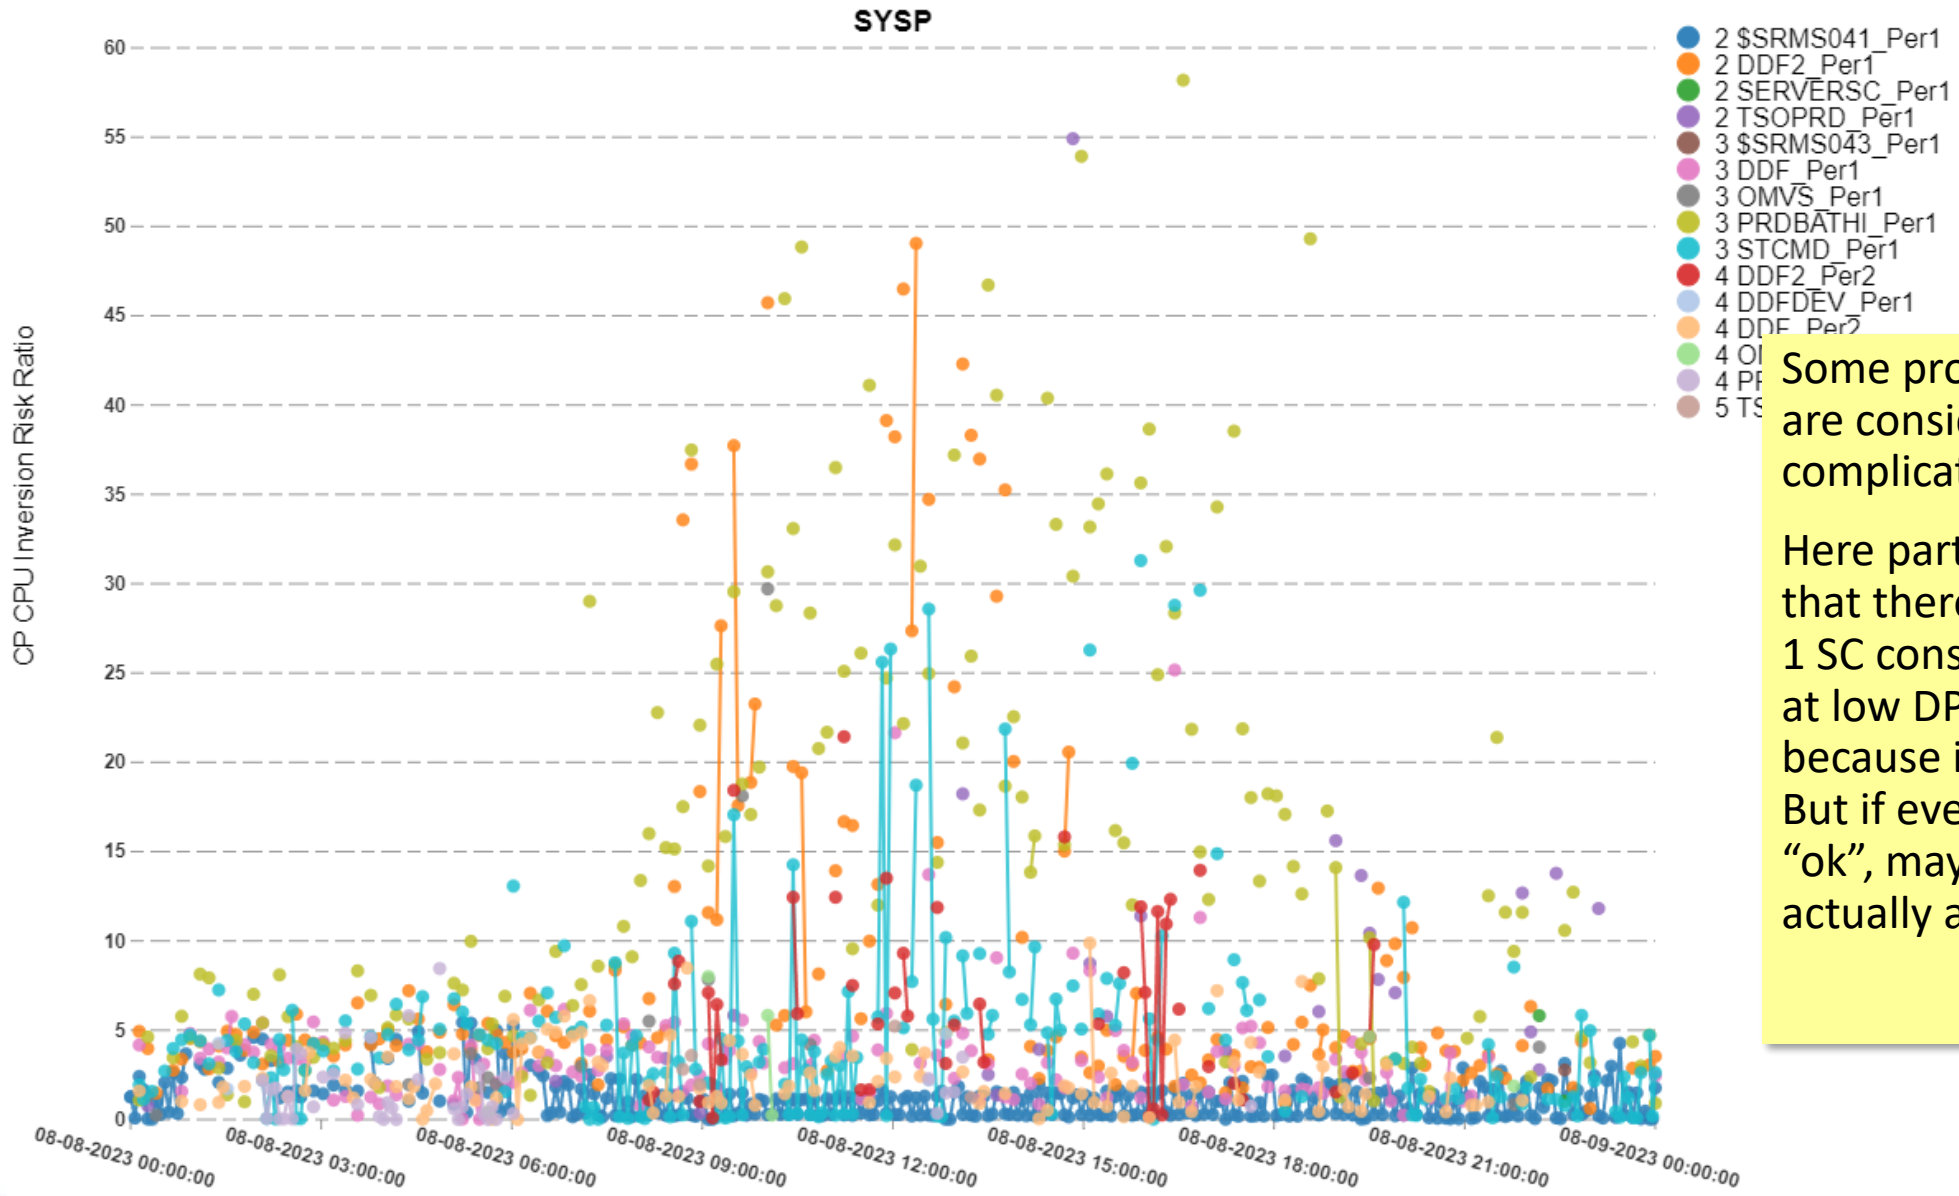

Some production systems are considerably more complicated!

Here part of the issue is that there's an importance 1 SC consistently running at low DPs. That might be because it has a poor goal. But if everything is running "ok", maybe it's not actually a bad goal.

Our thoughts (at this time)

- We don’t see the need for this change
  - A significant part of the premise of WLM was that it would manage dispatching priorities and could intelligently move them in possibly counter-intuitive ways to better balance throughput for diverse workloads
  - If you want, you can make all importance 1 work CPU Critical today
- We’d recommend turning this off for z/OS 3.1 and wish that was the default
- If you want to go to z/OS 3.1 with it on, we might suggest
  1. Evaluate which workloads are at risk
  2. Before 3.1, incrementally add CPU Critical to importance 1 workloads
     - If something goes wrong you can back out your change and z/OS 3.1 doesn’t get the blame
- We do sometimes recommend CPU Critical, but it’s an exception, not the rule
- Again: emerging area of study, we might refine our recommendations over time

More thoughts soon…

- We intend to do a deeper presentation on this topic for a webinar in the coming months
- Peter will be presenting on it at GSEUK

- Special thanks to Ethan Chapman for his statistical and R expertise!

Large memory should mean less I/O

In short: see Scott's presentation on Friday!

Memory and I/O

- We see systems with lots of memory free and yet they're doing significant amounts of I/O
- We've been saying for a long while things like "make your BPs bigger"
- But lately we've been trying to look deeper to point out opportunities
  - How much data is really on those busy volumes?
  - Which specific datasets are getting lots of read I/O

LVs with Highest I/O Rates

(Averaged Over Period of Study)

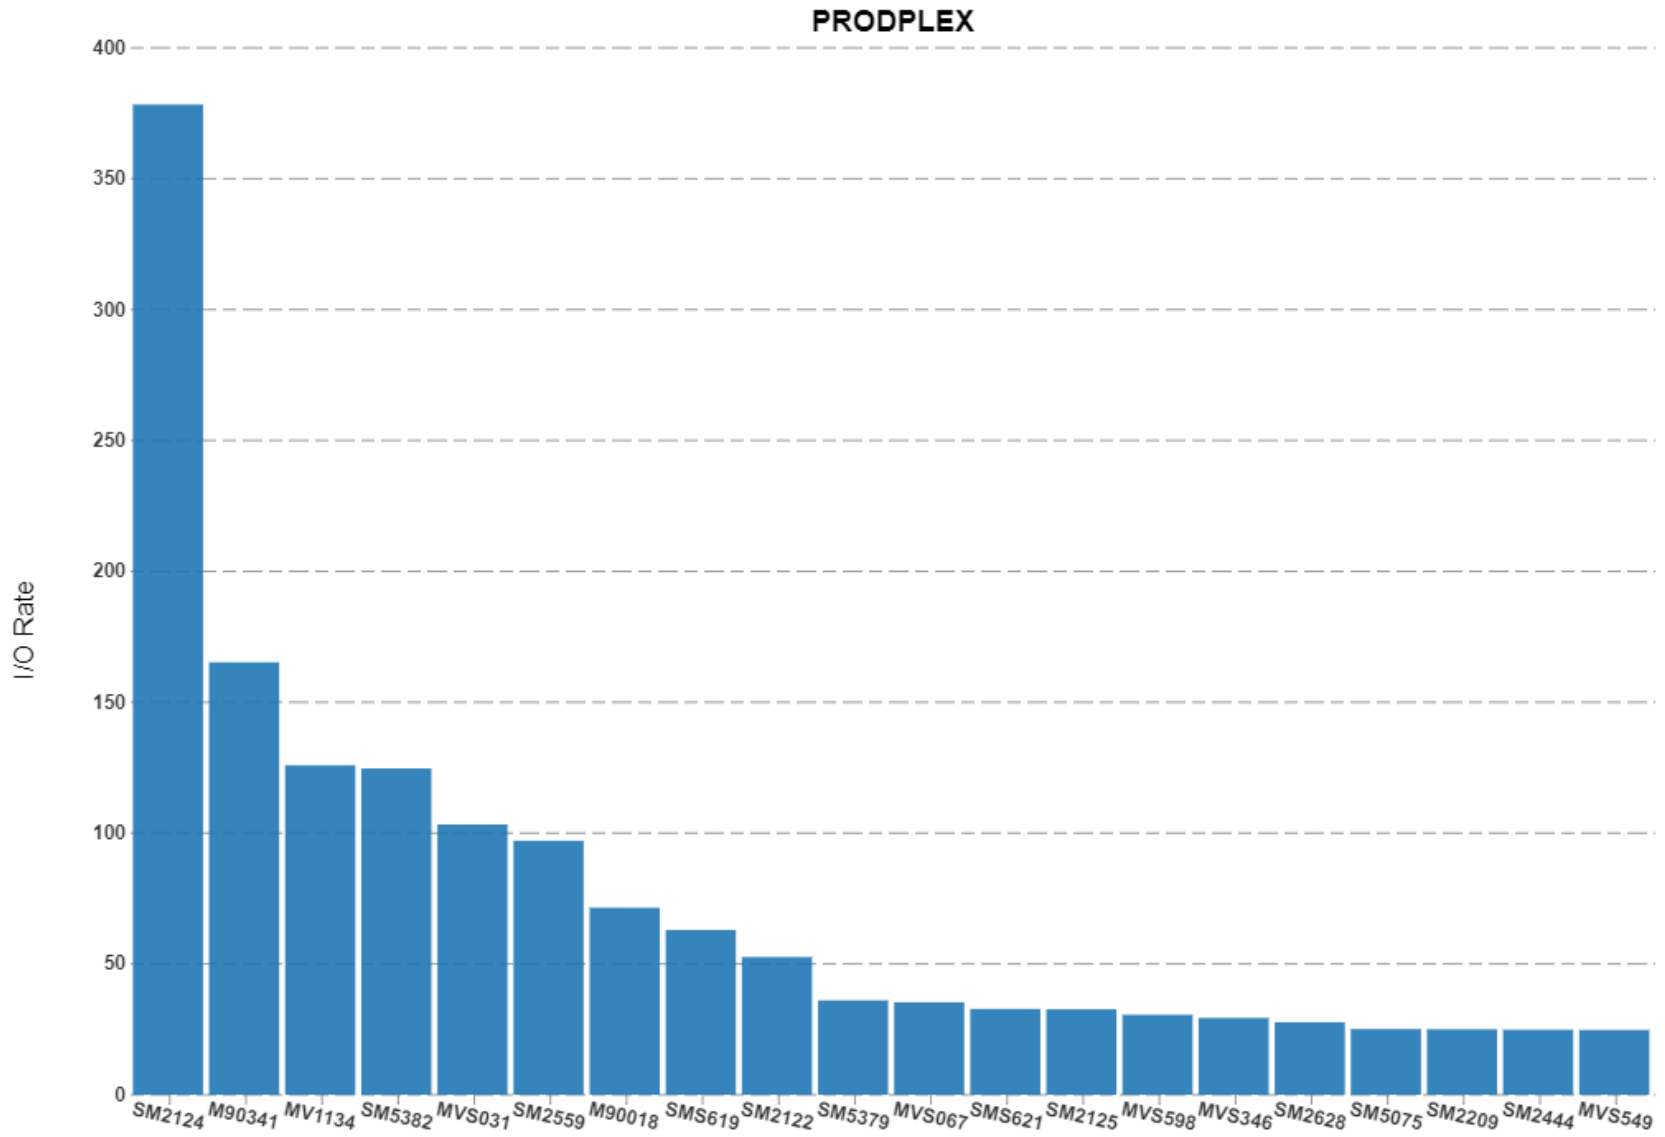

This Pivotor report shows the top volumes by I/O rate over the day.

375 IOPS doesn't sound too interesting but note that is an average I/O rate over 24 hours.

Logical DASD Volume Explorer

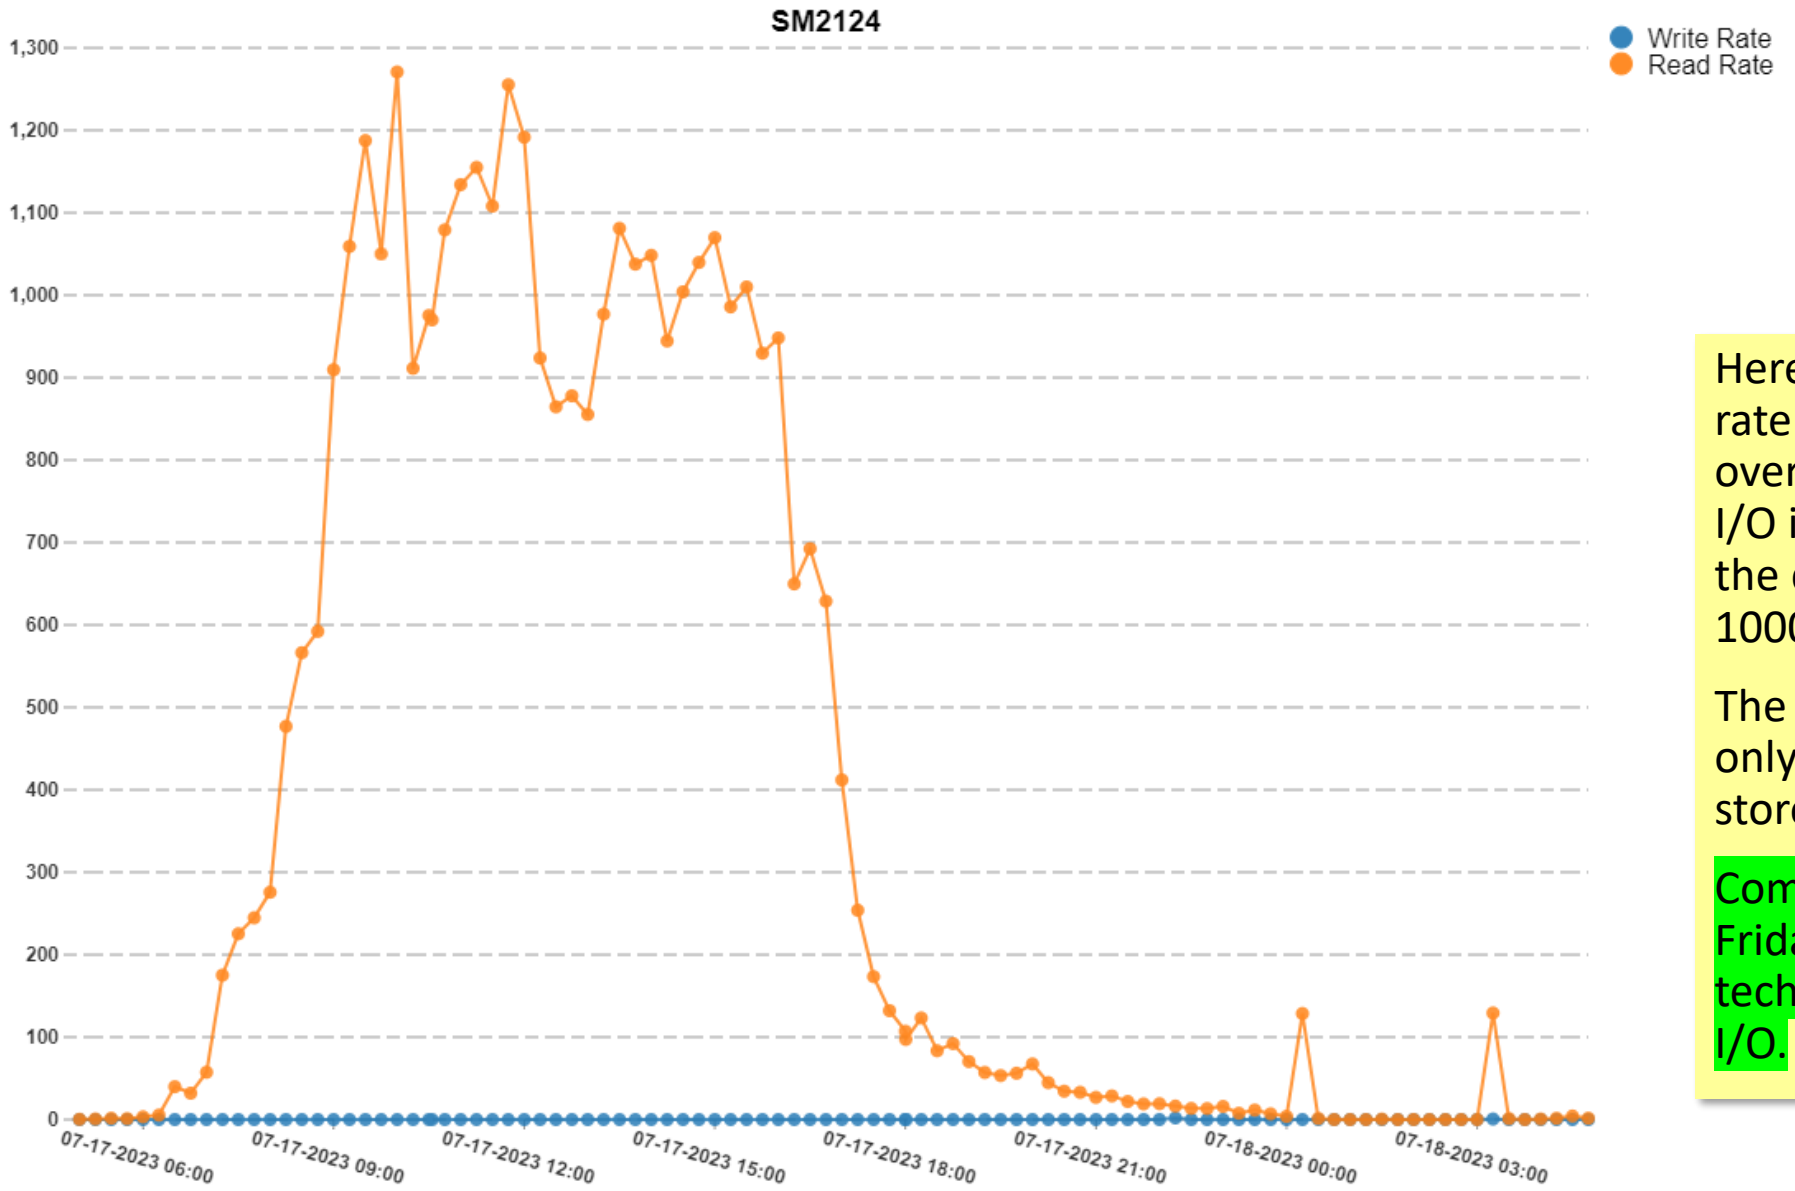

Here's the read and write rate for a particular volume over time. Virtually all the I/O is read I/O, and during the day it is doing over 1000 IOPS.

The kicker: this volume only has 1.5 GB of data stored on it!

Come to Scott's talk on Friday for specific techniques for avoiding I/O.

How will AI change what we do?

humanscript

> A truly natural scripting language

humanscript is an inferpreter. A script interpreter that infers the meaning behind commands written in natural language using large language models. Human writeable commands are translated into code that is then executed on the fly. There is no predefined syntax, humanscripts just say what they want to happen, and when you execute them, it happens.

The humanscript inferpreter supports a wide range of LLM backends. It can be used with cloud hosted LLMs like OpenAI's GPT-3.5 and GPT-4 or locally running open source LLMs like Llama 2.

Example

This is a humanscript called `tidy-screenshots`. It takes an unorganised directory of screenshots and organises them into directories based on the month the screenshot was taken.

```
#!/usr/bin/env humanscript

loop over all files (ignoring directories) in $HOME/Screenshots

move each file into a subdirectory in the format year-month

while the task is running show an ascii loading spinner

show how many files where moved

show the size of each subdirectory
```

It can be executed like any other script.

```
$ ./tidy-screenshots
Moved 593 files.
364K    2023-08
```

Note that I am definitely not suggesting that you actually do this!

github.com/lukechilds/humanscript

```
PROCEDURE DIVISION.
          DISPLAY "CALC Begins." UPON CONSOLE.
          MOVE 1 TO BUFFER-PTR.
          MOVE SPACES TO INPUT-1.
          PERFORM ACCEPT-INPUT UNTIL INPUT-1 EQUAL TO "END".
          DISPLAY "CALC Ends." UPON CONSOLE.
          GOBACK.
      *
      * Accept input data from buffer
      *
       ACCEPT-INPUT.
          MOVE BUFFER-ARRAY (BUFFER-PTR) TO INPUT-1.
          ADD 1 BUFFER-PTR GIVING BUFFER-PTR.
          EVALUATE FUNCTION UPPER-CASE(INPUT-1)     CALC1
            WHEN "END"
              MOVE "END" TO INPUT-1
            WHEN "LOAN"
              PERFORM CALCULATE-LOAN
            WHEN "PVALUE"
              PERFORM CALCULATE-VALUE
            WHEN OTHER
              DISPLAY "Invalid input: " INPUT-1
          END-EVALUATE.
```

This looks pretty human-readable too!

Maybe not everything needs AI. Maybe some advances just need a smart woman in IT?

See also:
https://en.wikipedia.org/wiki/Grace_Hopper

Scott's AI Thoughts

- There's going to be a lot of interesting applications for AI over the next several years
  - Most of which have nothing to do with managing z/OS performance
  - For some forms & uses of AI there's a lot of questions and uncertainty in the realms of ethics, legal liabilities, and potential regulation
- z/OS performance analysts are not going to be put out of a job tomorrow
  - There's a lot of exterior factors that come into play in managing a system that is not captured in the performance data about the system
    - Not all dispatching priority inversions are bad, not all "bad" goals are wrong
    - Sometimes we intentionally restrict performance for various reasons
  - Anybody(thing) evaluating your system should be asking "what" and "why" and explaining "what" and "why" as well!
- On the other hand, there are some machine learning techniques that likely will be valuable tools
  - Some (many) of which really are just refinements to what has been done for years

Batch Management

- "AI-powered Workload Manager (WLM), designed to intelligently predict upcoming batch workload and react accordingly to optimize system resources in a proactive way. This AI capability represents the first use case that leverages the AI Framework for IBM z/OS." (IBM announcement)
- But … predicting upcoming batch workloads and proactively managing initiators has been a thing in the past without AI
  - E.G. ThruputManager Automation Edition
  - And z/OS Performance Analysts have been doing this with Actual Intelligence
- Nonetheless, this is an interesting area to explore and could be useful
  - Given how reluctant people were to move to WLM-managed inits… it will be interesting to see the uptake on AI-managed initiators!

Outlier / Anomaly Detection

- How do you know if something is different?
  - Of course not all anomalies are problems!
- There's been a lot of statistical techniques used for anomaly detection
  - E.G. MASF by Buzen and Shum in 1995
- Some of the new machine learning techniques do look like they might address some of the stats techniques' shortcomings
  - But none of these are perfect
- We're actively working on a ML-based outlier detection system
  - There's lots of non-trivial issues, including just how to best present the outliers!

Ongoing Opportunities

Things we're still talking about with people

Re-evaluating goals

- We continue to come across sites where the WLM policy has been static for many years
- You need to be re-evaluating your policy periodically
- Come to Peter's session this afternoon!

SuperPAV

- IOSQ time is rarely a significant component of I/O response time, but we still sometimes see some
- SuperPAV generally eliminates the little remaining IOSQ time
  - SuperPAV enables sharing of PAVs between LCUs, effectively allowing access to more PAVs for each volume
- If your DASD is less than even 5 years old, it almost certainly supports SuperPAV
  - Check with your DASD vendor and enable in IECIOS: HYPERPAV=XPAV

Record the 98s and 99s

- They provide insights into performance at a sub-minute level
  - 10 second WLM Policy Adjustment interval
  - 2 second HiperDispatch interval
  - 5-60 second High Frequency Throughput Statistics
- Yes, you're not going to look at them every day, but they can be quite useful for problem determination: especially for transient problems!

SMF 98/99 records to Include

- SMF 98 High-frequency Throughput Statistics (HFTS)
  - IBM recommendation is to record on 5 second interval
    - Can use 5, 10, 15, 20, 30 or 60 seconds
    - 5 second interval is about 400MB-500MB/system/day

```
In SMFPRMxx:

HFTSINTVL(15)
```

- SMF 99 SRM/WLM details
  - Our minimum recommended subtypes: 6, 10, 11, 12, 14
    - These will be around 50-150MB/system/day
  - Subtype 1, 2, and 3 can be quite useful, but can be more voluminous
    - These can be 1-1.5GB/system/day
  - Pivotor customers: send them if you're collecting them!
  - Subtype 13 is somewhat voluminous but is undocumented "IBM use only"
    - 150-200MB/system/day
    - We recommend you turn off subtype 13s until/unless IBM asks for them

None of these records represent data you will look at every day, but it's nice to have them available when you need them!

Classic CEC Utilization Transient Performance Problem

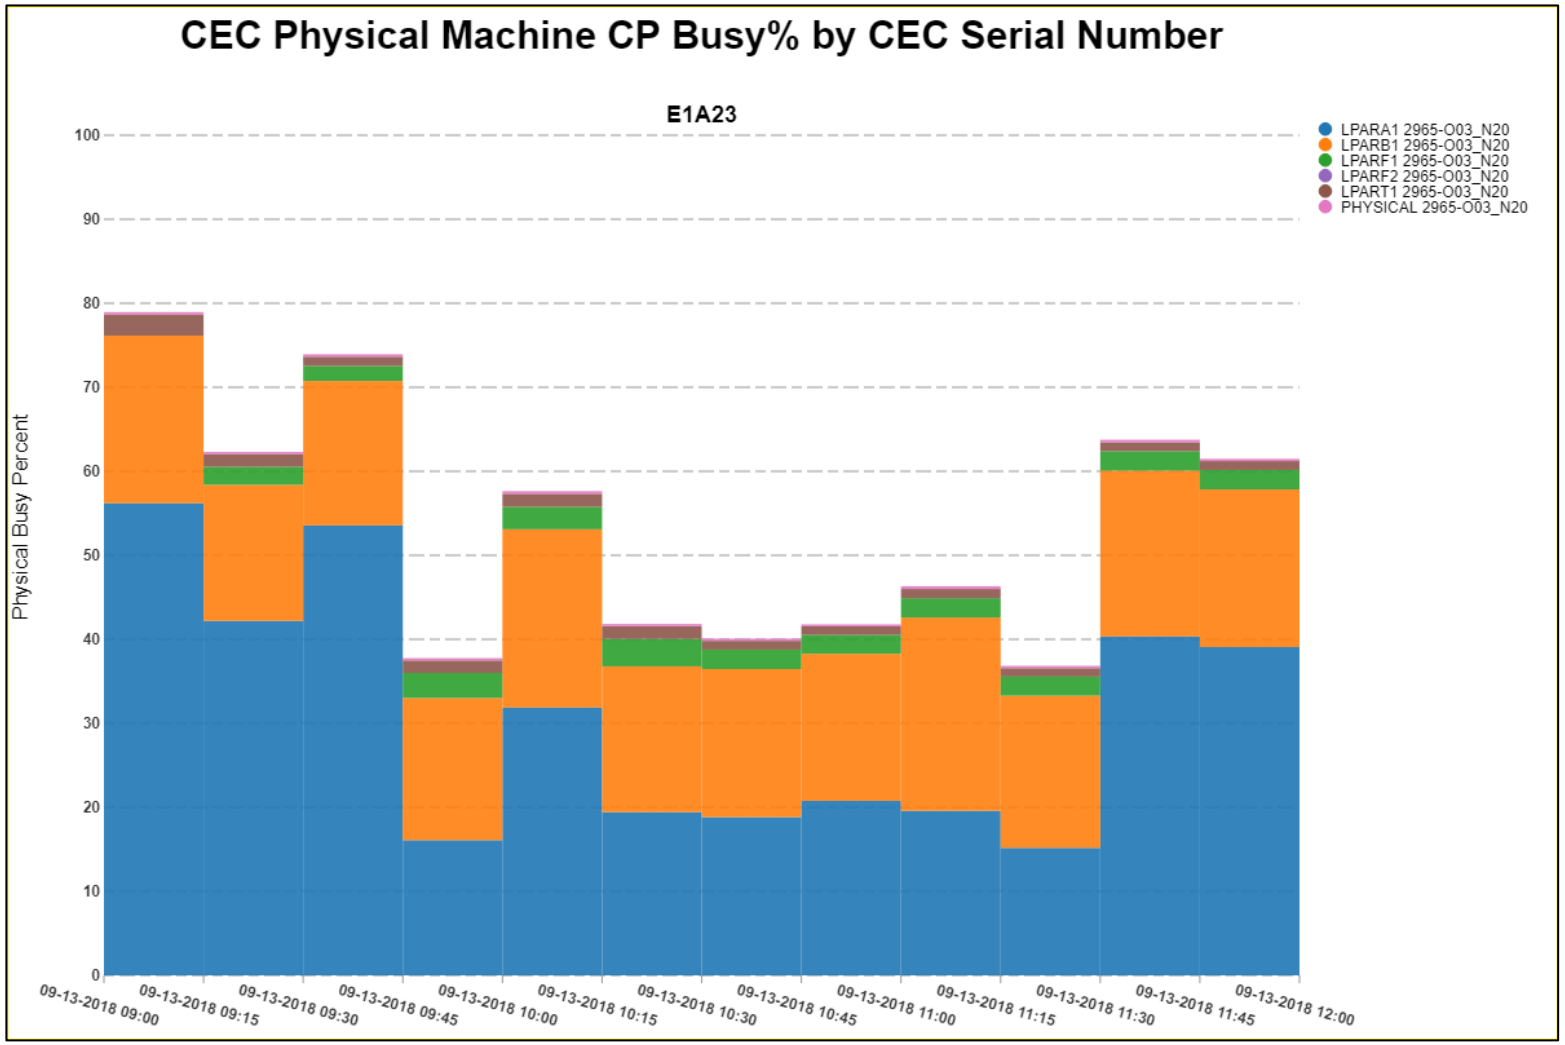

Problem Statement:

System Seemed to Freeze / Stall / things too a long time, but we have lots of available capacity

This is just a standard view of CEC Utilization, here we've narrowed in to just 3 hours in the morning, where it doesn't appear there's really any capacity concerns.

This chart is generated from data that comes from the SMF 70 records. In this example, the measurement intervals are 15 minutes.

Classic CEC Utilization Transient Performance Problem

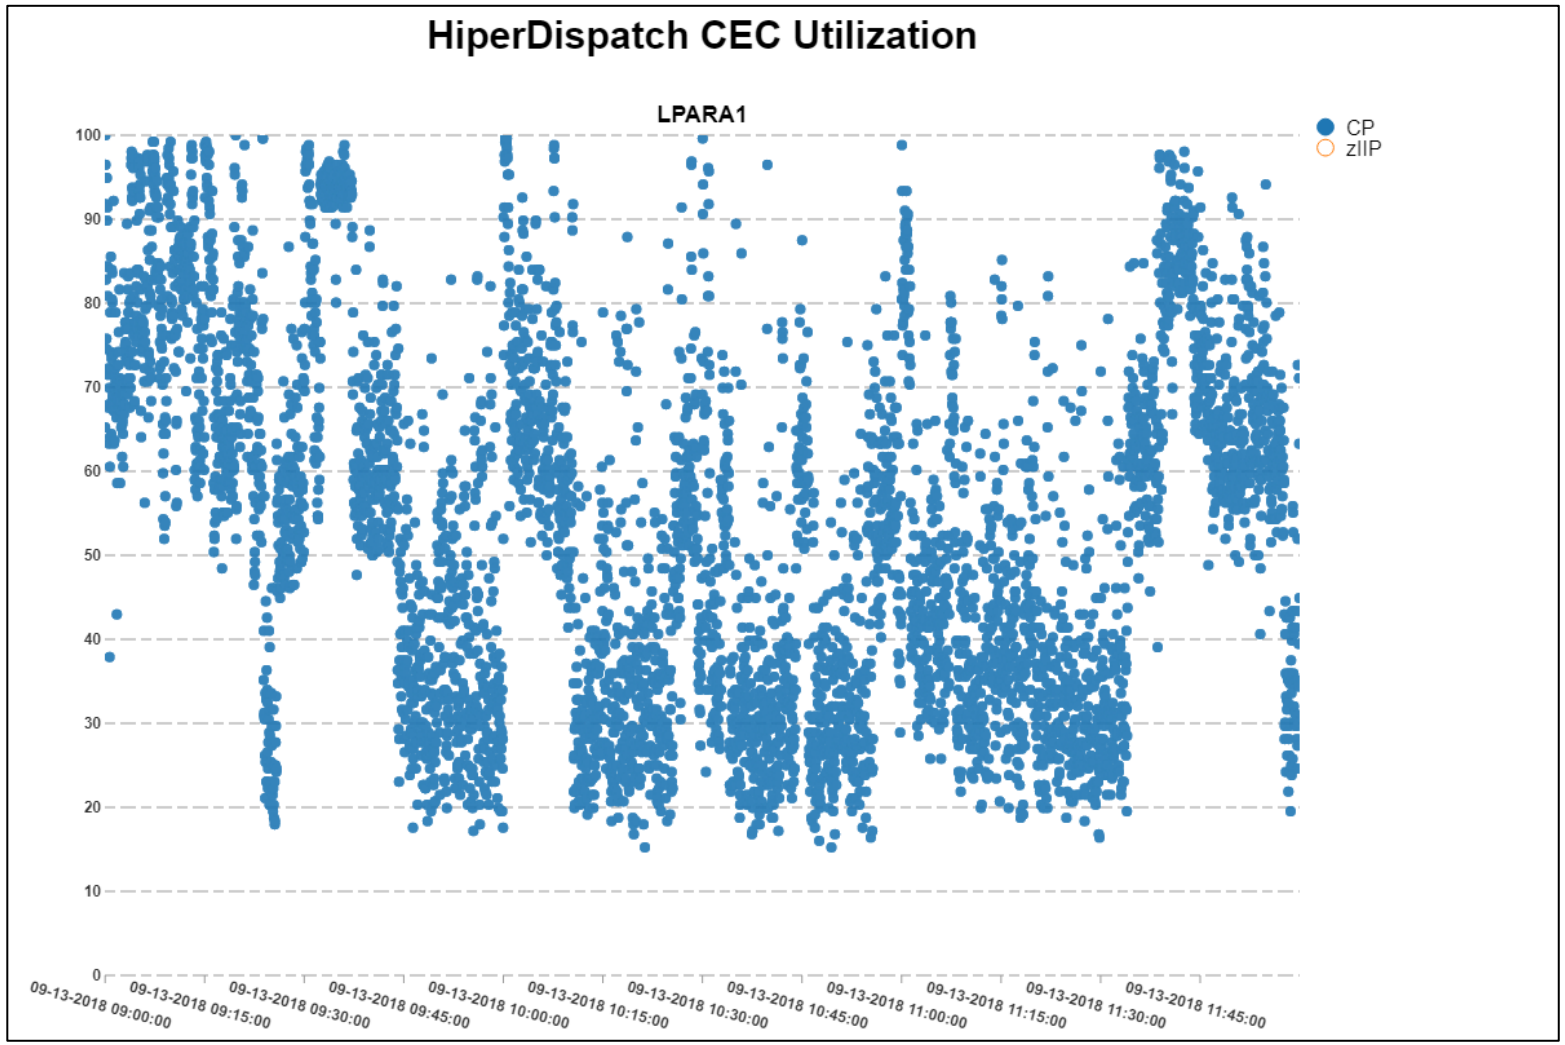

High Frequency CEC Utilization:

This also is a CEC utilization chart for the same 3 hours as the previous chart.

This data comes from the from the SMF 99.12 HyperDispatch records.

The CEC utilization is at 2-second measurement interval.

Note that this tells a different story than the 15-minute RMF intervals.

SMT

- We sometimes see customer with SMT enabled “just because”
  - That’s probably “ok” but it’s probably also unnecessary
- In some cases unnecessary use of SMT might be sub-optimal
  - Remember z/OS densely packs the cores so even if you have a relatively high number of unused zIIP cores, with SMT enabled the work will be assigned to an in-use core first
- Our general recommendation is to only enable SMT when you have a defined need
  - Leave SMT in your bag of tricks to pull out to buy some additional headroom when the need develops
  - SMT also makes detailed capacity planning for zIIPs effectively impossible.
- See also Scott’s SMT presentation on our website https://www.pivotor.com/content.html

SMT Enablement Flowchart

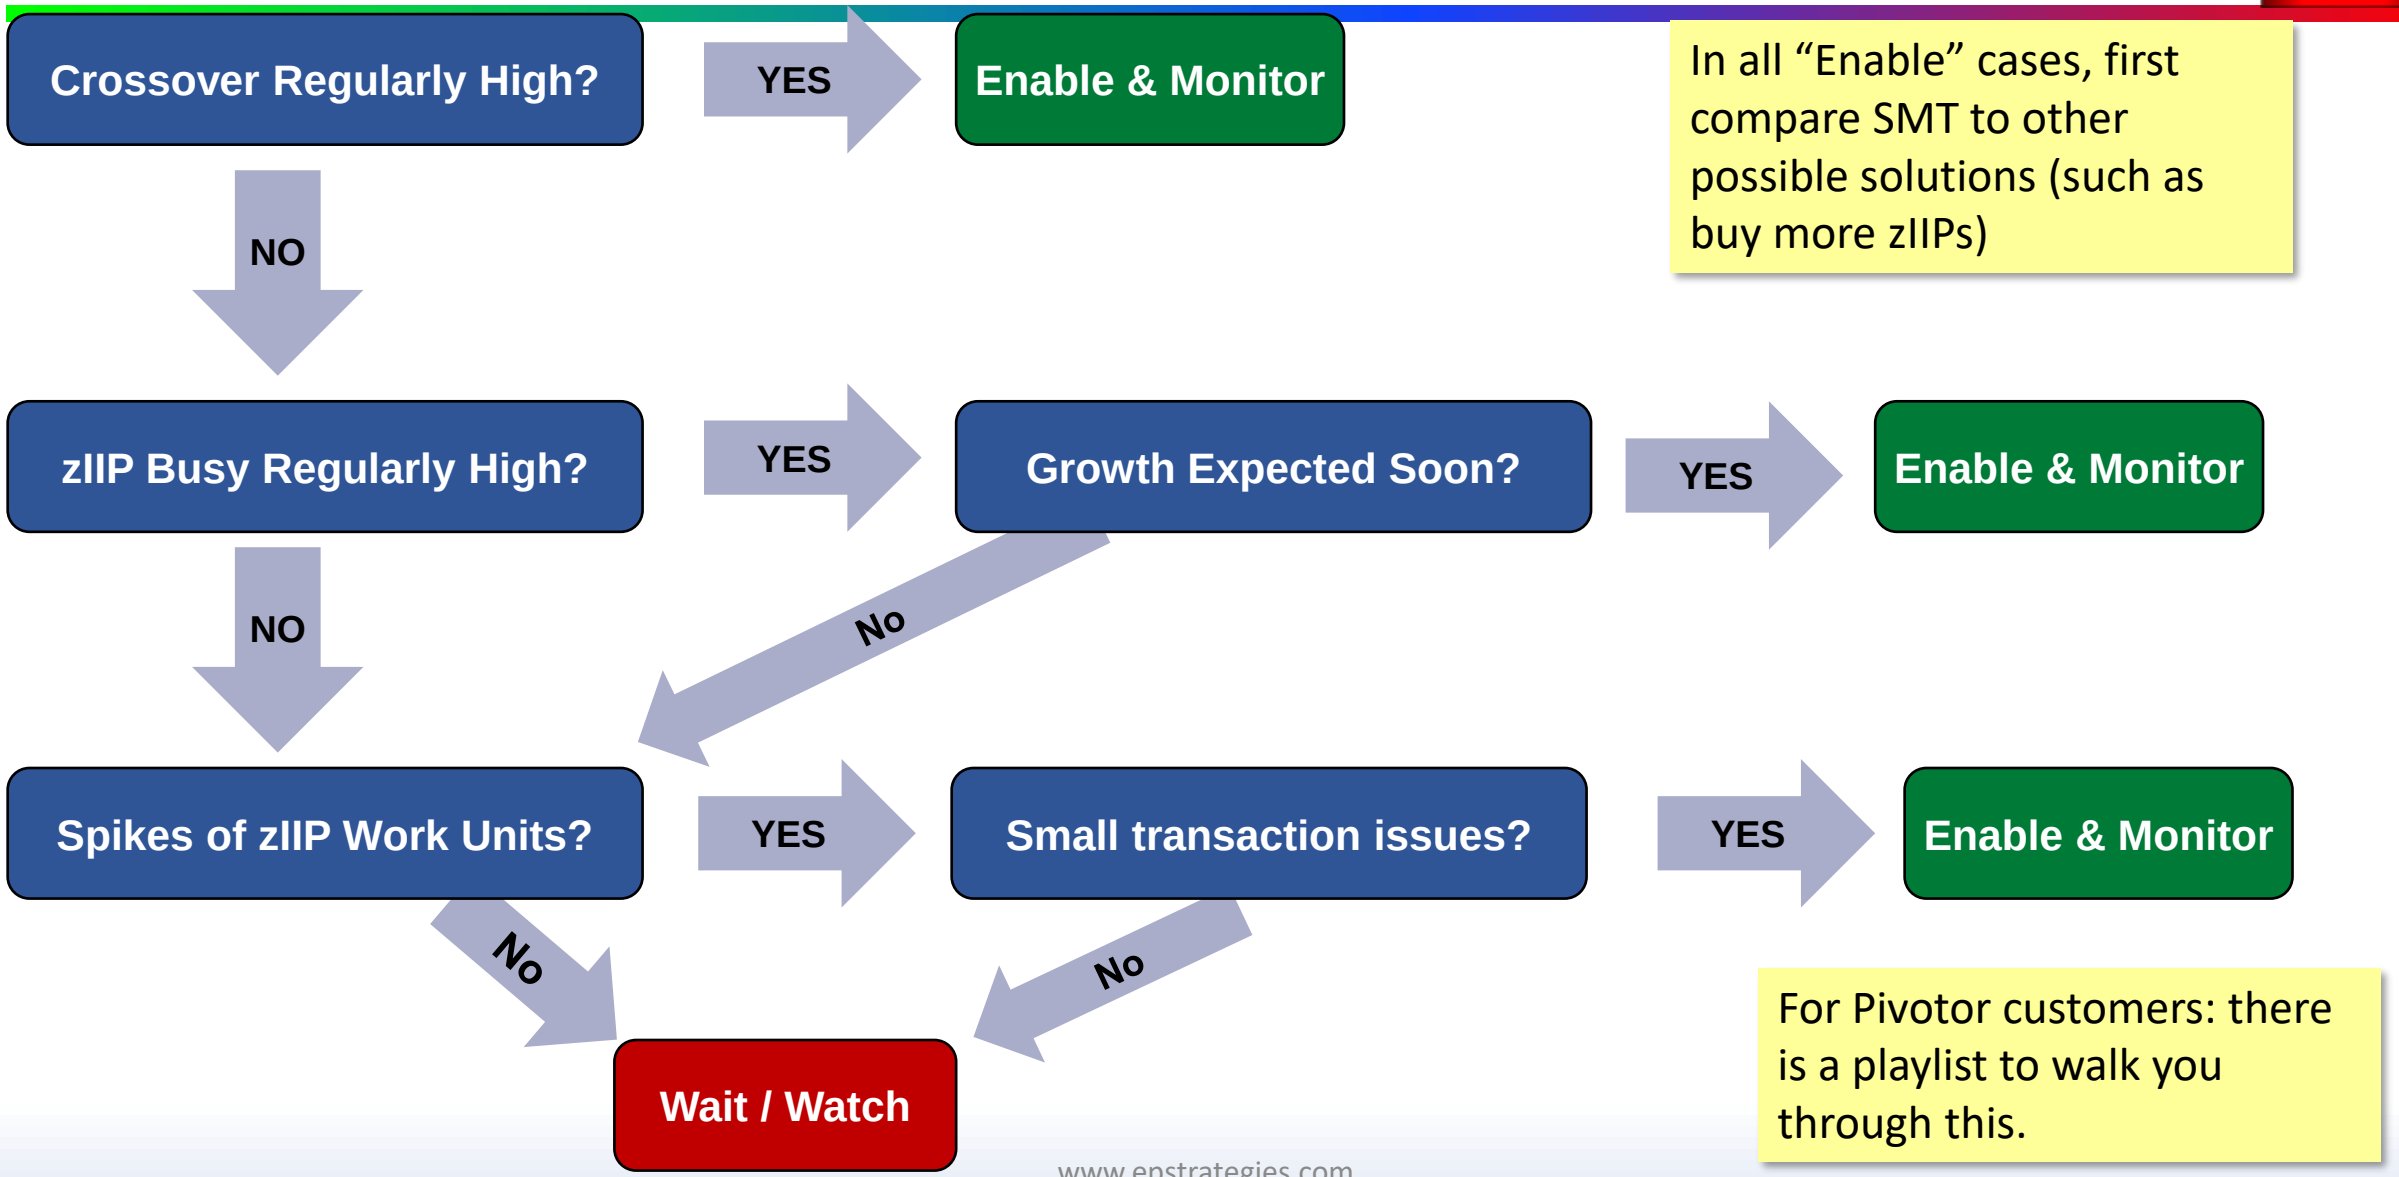

In all "Enable" cases, first compare SMT to other possible solutions (such as buy more zIIPs)

For Pivotor customers: there is a playlist to walk you through this.

Service Definition Coefficients Updates

- Recommended values by EPS since about 2018 (maybe earlier)
  - CPU=1, SRB=1, IOC=0. MSO=0
  - Summary of reasoning: Aging a transaction based on I/O no longer made much sense since I/O priority management mattered much less due to advent of PAVs, and most I/O processing is also outside the z/OS operating system. So why age a workload based on its I/O characteristics. It is CPU that matters.
- z/OS 2.5 the SDCs go away, and the values will default as follows
  - CPU=1, SRB=1, IOC=0. MSO=0
  - Basically, it is durations are now based on CPU and SRB service units, and not longer based on the concept of 'service'.
- Most customers are using 1,1,0,0
  - If you haven't made the transition yet, read next slides…

IBM's z/OS 2.5 Migration Step

The following is an excerpt from SHARE presentation:
PERFORMANCE INFRASTRUCTURE IMPROVEMENTS IN Z/OS V2.5 WLM
Presenter:
ANDREAS HENICKE (IBM WLM)

Presentation discusses the z/OS 2.5 migration steps suggested to migrate your period durations prior to migrating to z/OS 2.5.

Basically, IBM is suggesting to take CPU and SRB 'service', divide by your current SDCs to convert to 'service units'. Then take the sum of those values and multiple them by the ratio of current duration to service consumed.

Or put a little simpler…
Blah, blah, blah…

Feel free to take this approach, but a bit to complicated for me.

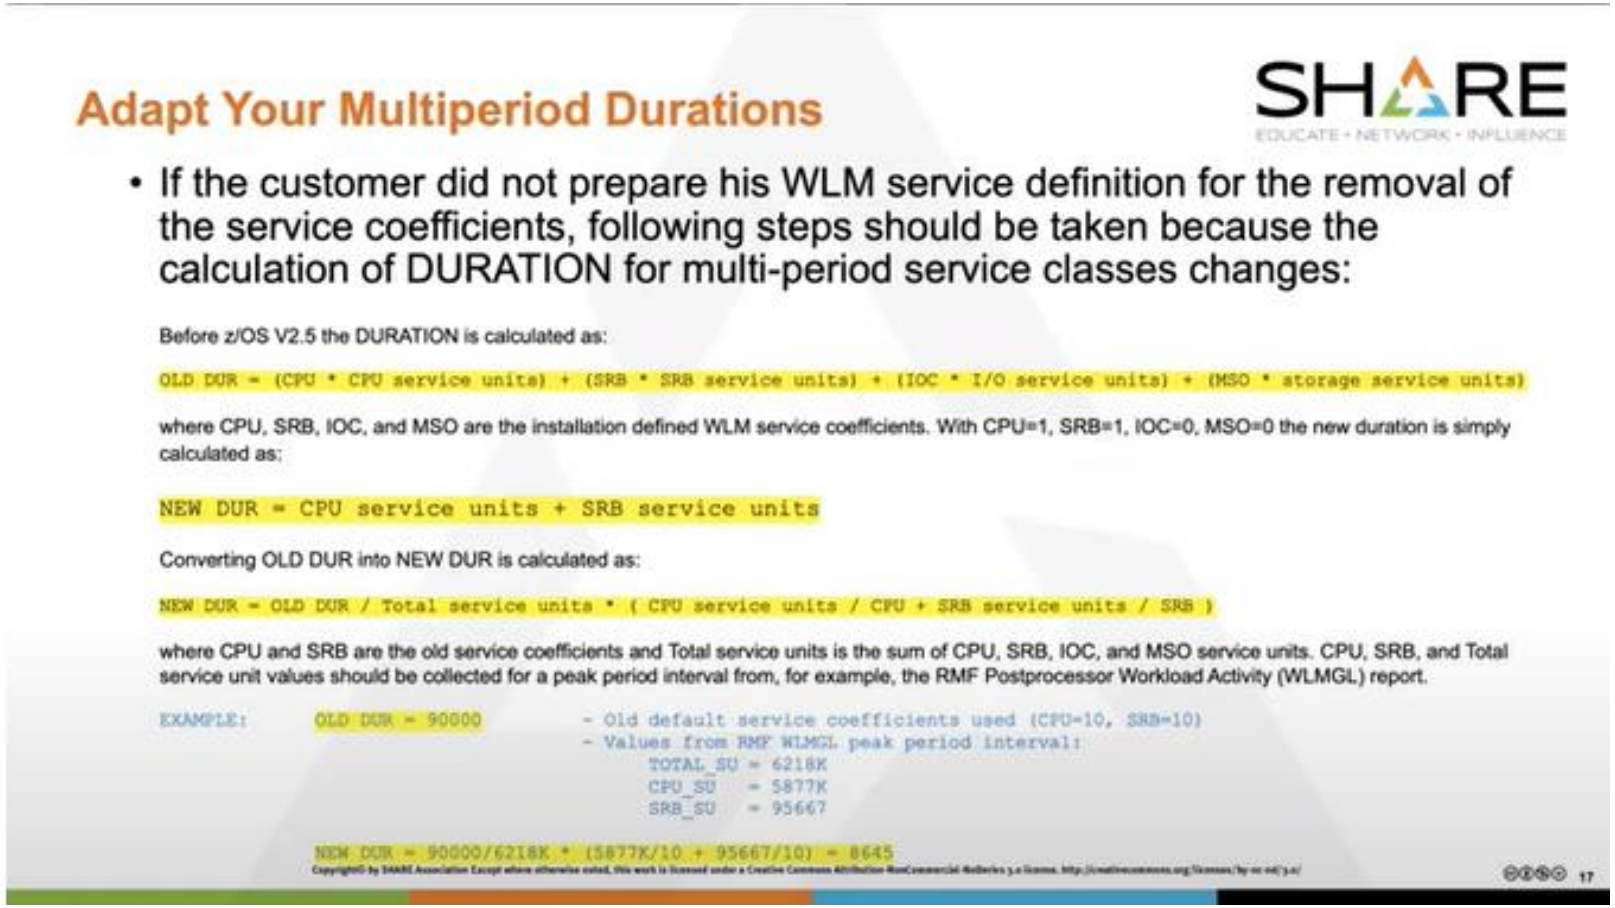

Adapt Your Multiperiod Durations

SHARE
EDUCATE • NETWORK • INFLUENCE

- If the customer did not prepare his WLM service definition for the removal of the service coefficients, following steps should be taken because the calculation of DURATION for multi-period service classes changes:

Before z/OS V2.5 the DURATION is calculated as:

```
OLD DUR = (CPU * CPU service units) + (SRB * SRB service units) + (IOC * I/O service units) + (MSO * storage service units)
```

where CPU, SRB, IOC, and MSO are the installation defined WLM service coefficients. With CPU=1, SRB=1, IOC=0, MSO=0 the new duration is simply calculated as:

```
NEW DUR = CPU service units + SRB service units
```

Converting OLD DUR into NEW DUR is calculated as:

```
NEW DUR = OLD DUR / Total service units * ( CPU service units / CPU + SRB service units / SRB )
```

where CPU and SRB are the old service coefficients and Total service units is the sum of CPU, SRB, IOC, and MSO service units. CPU, SRB, and Total service unit values should be collected for a peak period interval from, for example, the RMF Postprocessor Workload Activity (WLMGL) report.

```
EXAMPLE:    OLD DUR = 90000        - Old default service coefficients used (CPU=10, SRB=10)
                                   - Values from RMF WLMGL peak period interval:
                                        TOTAL_SU = 6218K
                                        CPU_SU   = 5877K
                                        SRB_SU   = 95667

            NEW DUR = 90000/6218K * (5877K/10 + 95667/10) = 8645
```

Peter's Approach to Migrating SDCs to New z/OS 2.5

- Understand that most durations for multiple periods are usually wrong to begin with.
  - If you feel yours are correct, then do this exercise

- My general approach is a follows:
  1. Determine your current SDCs
  2. Determine your current multiple period service classes
     - Most likely multiple periods are only being used for the following interactive workloads or certain batch
       - TSO, Interactive OMVS, DDF, WAS CB, Batch (sometimes)
  3. Determine which multiple period service classes are consuming I/O service and how much
  4. Then ignore any sort of duration migration exercise for the following enclave workload types since these enclave workloads do not consider I/O service
     - DDF
     - WAS CB
     - So will be left with workloads such as eft with only TSO, interactive OMVS, and Batch,
  5. Revisit duration
     - Either start fresh (which should be done for many periods regardless of this change)
     - Ignore and accept
     - Tweak

I/O Priority Management

- A few (several?) years ago we made the recommendation that most customers should disable I/O Priority Management in WLM
  - Recommendation had been for ~20 years to enable it
  - Changing reality of I/O meant that having it enabled inflated velocities
- At the time we said probably 90% of sites shouldn't have it enabled
- Having seen even more data over the years, that's probably now >99%
  - It makes WLM focus on just CPU using and delays
  - May have to revisit/reset your velocity goals when you do this though
  - "Worst" case is that turning it off makes no difference
- IBM is also now recommending to turn off I/O Priority Management

Wrap-up

- We hope you enjoyed this and that you've learned something
- Let us know if you like this potpourri of topics format
- We'll be around now and all week for questions

- Questions?

- Please visit our website: www.epstrategies.com
  - Past presentations
  - WLM to HTML tool
  - More information about Pivotor
  - Future educational webinars