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®

SHARE Orlando
February 2026
Session
SPON106

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!

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

Today's Agenda

- Who we are / what we do (Peter)
- Pivotor Latest Updates (Peter)
- Updated SMF Recording Recommendations (Scott)
- Emerging Areas of Interest (we will not get through the whole list)
  - Sites aren't fully utilizing SRB (Scott)
  - I/O Issues (Scott)
  - CPENABLE and z/OS 3.1 (Scott)
  - Implicit CPU Protection in z/OS 3.1 (Peter)
  - New enclave measurements in response to DB2 HiPerf DBATs (Peter)
  - First Reference Page Faults (Peter)
  - IXGCNFxx Keep local buffers (Peter)
- Prize drawings! (Jamie)

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, as well as contract reviews

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	Tue 13:15	Salon 18
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 14:30	Salon 19

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?
Winter / Spring 2026 Webinars

- Free z/OS Performance Educational webinars!
  - The titles for our Winter/Spring 2026 webinars are as follows:
    - *New Year's Resolutions for z/OS Performance and Capacity People*
    - *How WLM Makes Decisions*
    - *What I Learned about VSAM RLS SMF Data*
    - *z/OS Performance Spotlight: Some Top Things You May Not Know*
    - *Building a Strong Foundation When You're New to z/OS Performance*
    - *Wait…Do We Need to Re-evaluate our WLM Goals?*
    - *z15 to z16 to z17 – What has changed?*
    - *Evaluating in the Mainframe Environment*
    - *Managing Workload Manager: Multiple Sysplexes and Asymmetric Sysplexes*
    - *Introduction to z Processor Measurements*
    - *(more to be announced)*

- 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?
Spring / Summer 2025 Webinars

- Free z/OS Performance Educational webinars!
  - 30-60 minutes long
  - 2026 Schedule Being put together now:
    - New Year's Resolutions for z/OS Performance and Capacity People
    - How WLM Makes Decisions
    - What I Learned about VSAM RLS SMF Data
    - Plus more… we're finalizing!
  - 2025 Webinars included:
    - ✓ *Overseeing z/OS Performance Management With Your Outsourcer*
    - ✓ *Back to basics - Processor Consumption Analysis*
    - ✓ *WLM and CPU Critical Control*
    - ✓ *Back to Basics - Evaluating Latent Demand*
    - ✓ *Understanding SMF 98 Locking Measurements*
    - ✓ *Overseeing z/OS Performance Management with Your Outsourcer*
    - ✓ *ETR vs ITR and the Basics of IBM's LSPR*
    - ✓ *Processor Comparison Discussion*
    - ✓ *SMF 99 WLM Decision Making Traces*
    - ✓ *Understanding SMF 98 Address Space Consumption Measurements*

email: contact@epstrategies.com

- 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

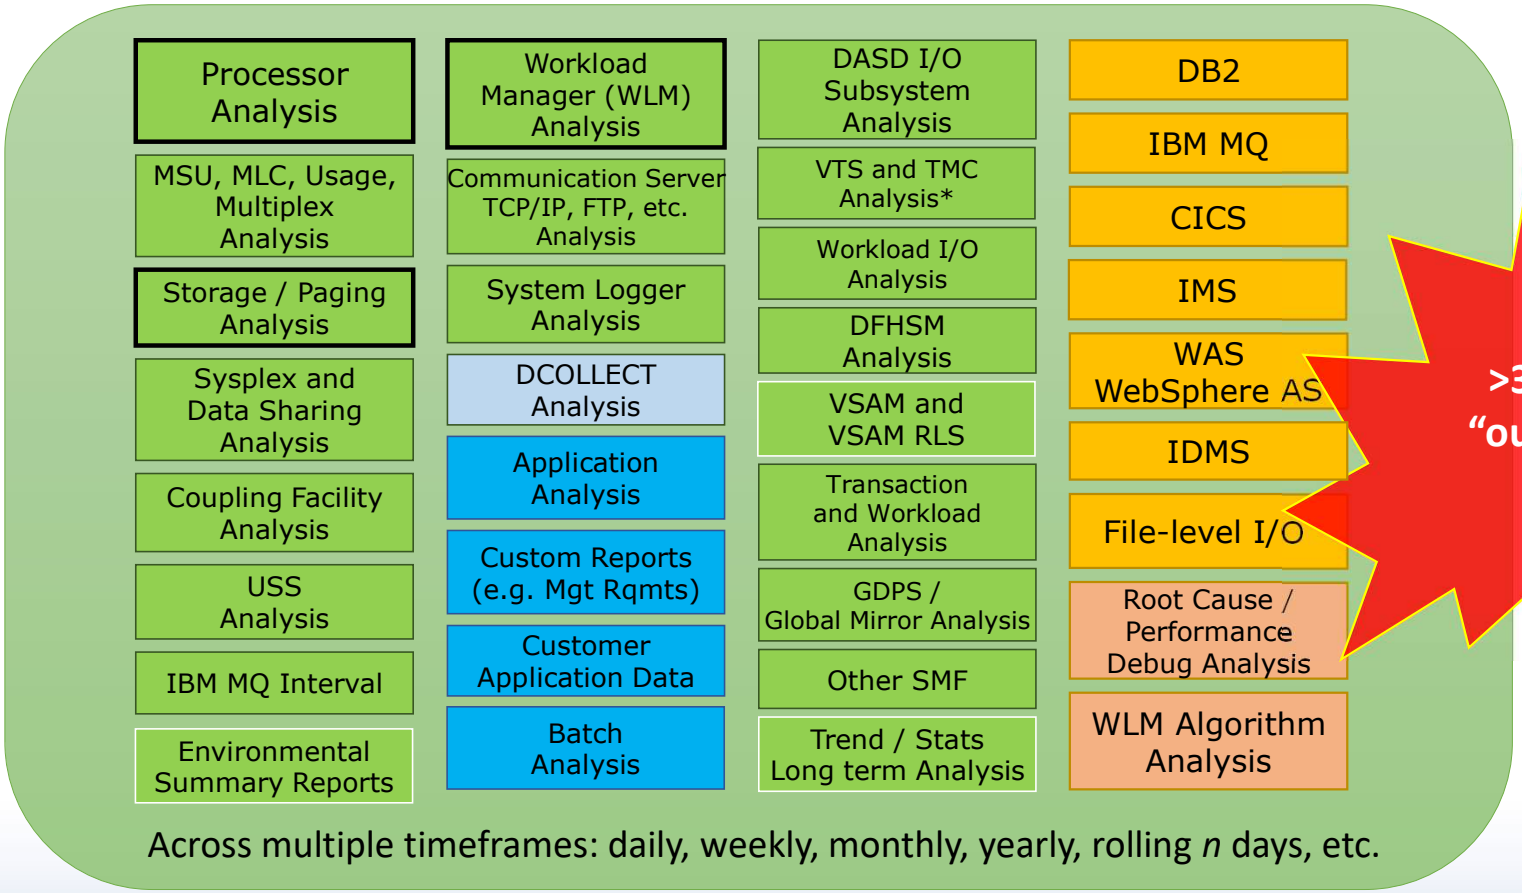

Pivotor Software as a Solution (SaaS)

- 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

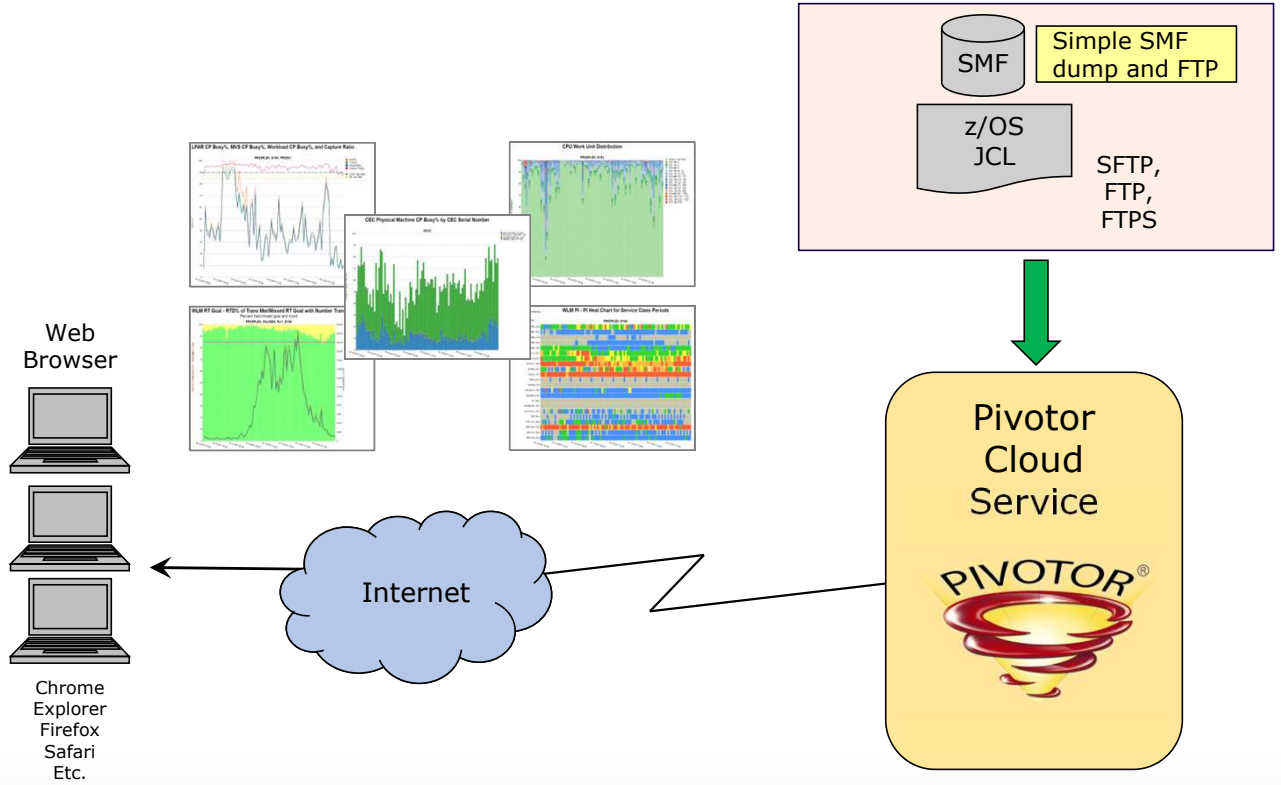

More Free Things!

- 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)

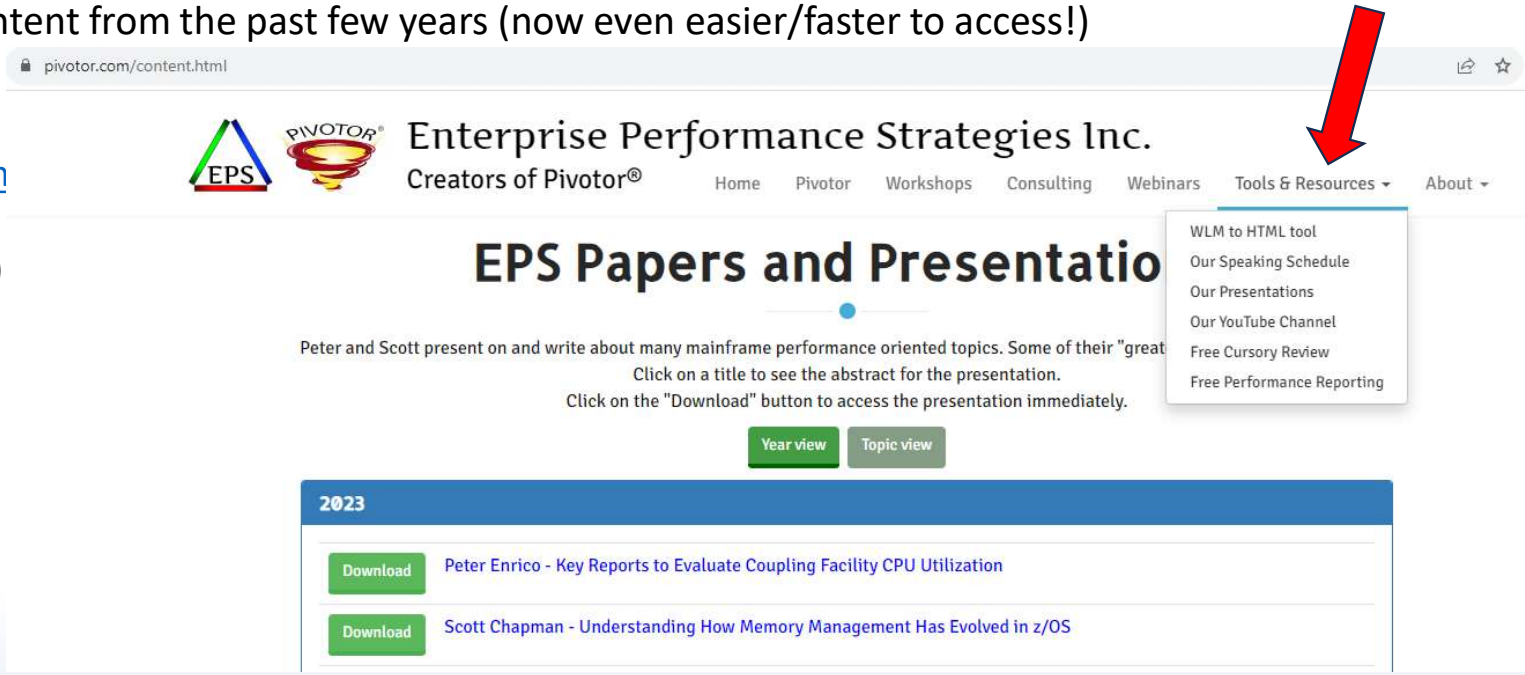

Pivotor Latest

Pivotor Batch Reporting

- As with other workload types, Pivotor is a powerful tool for Batch
  - Short term reporting
    - Job ended detail reports
    - Step ended detail reports
    - Outlier reports to zero in on anomalies
    - Top resource consumers
    - Reporting by customer application
    - New Timeline reports
  - Long term reporting
    - The ability to analyze a batch job's rolling year trending for over a year
    - Stability reports for trending
    - Top hitter reports
  - And much more…

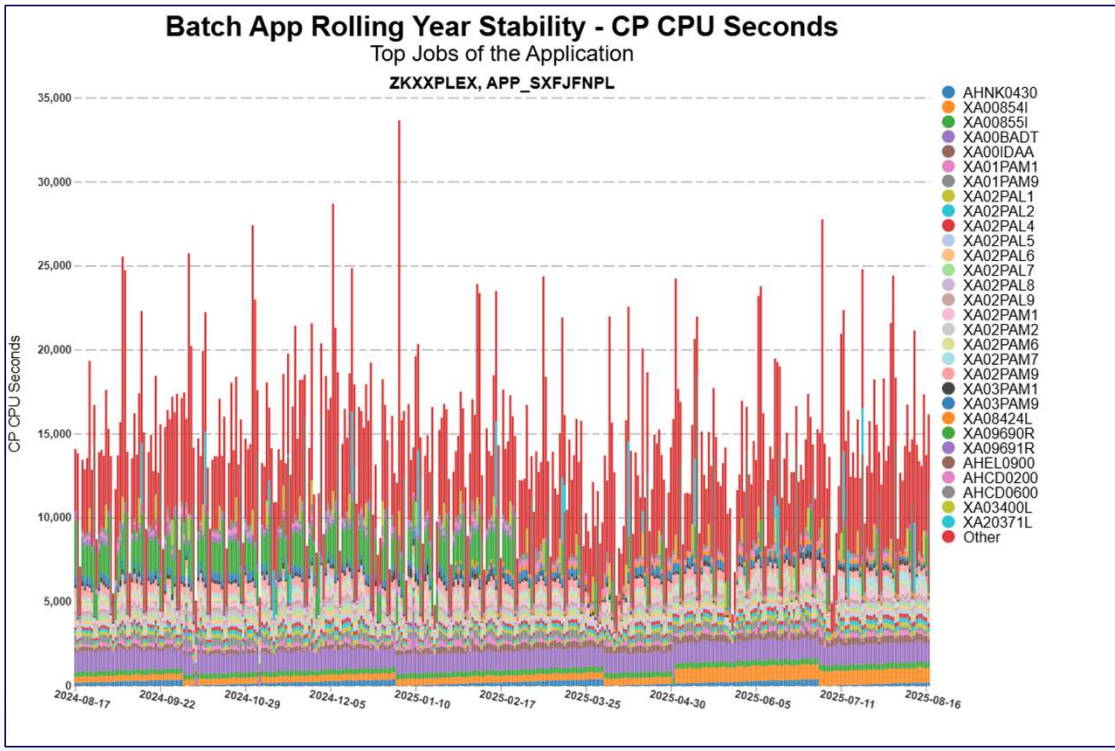

New Batch Analysis Timeline reports

- Primarily designed to help answer the question of what was running at a particular time, for a particular … something
  - Service class, job class, system, etc.
- Makes it easy to see how many jobs you have running concurrently
  - Subject to the limitation below
- Also useful to see where (job class, service class, etc.) jobs are running
  - In particular, by application (report class used as application by default)
- Limited to jobs with execution time >30 seconds <12 hours
  - Avoids cluttering up charts with excessive number of very short jobs
    - Surprising number of <30 second jobs running out there in the world
  - Avoids charting issues with CICS (e.g.) running for 24+ hours

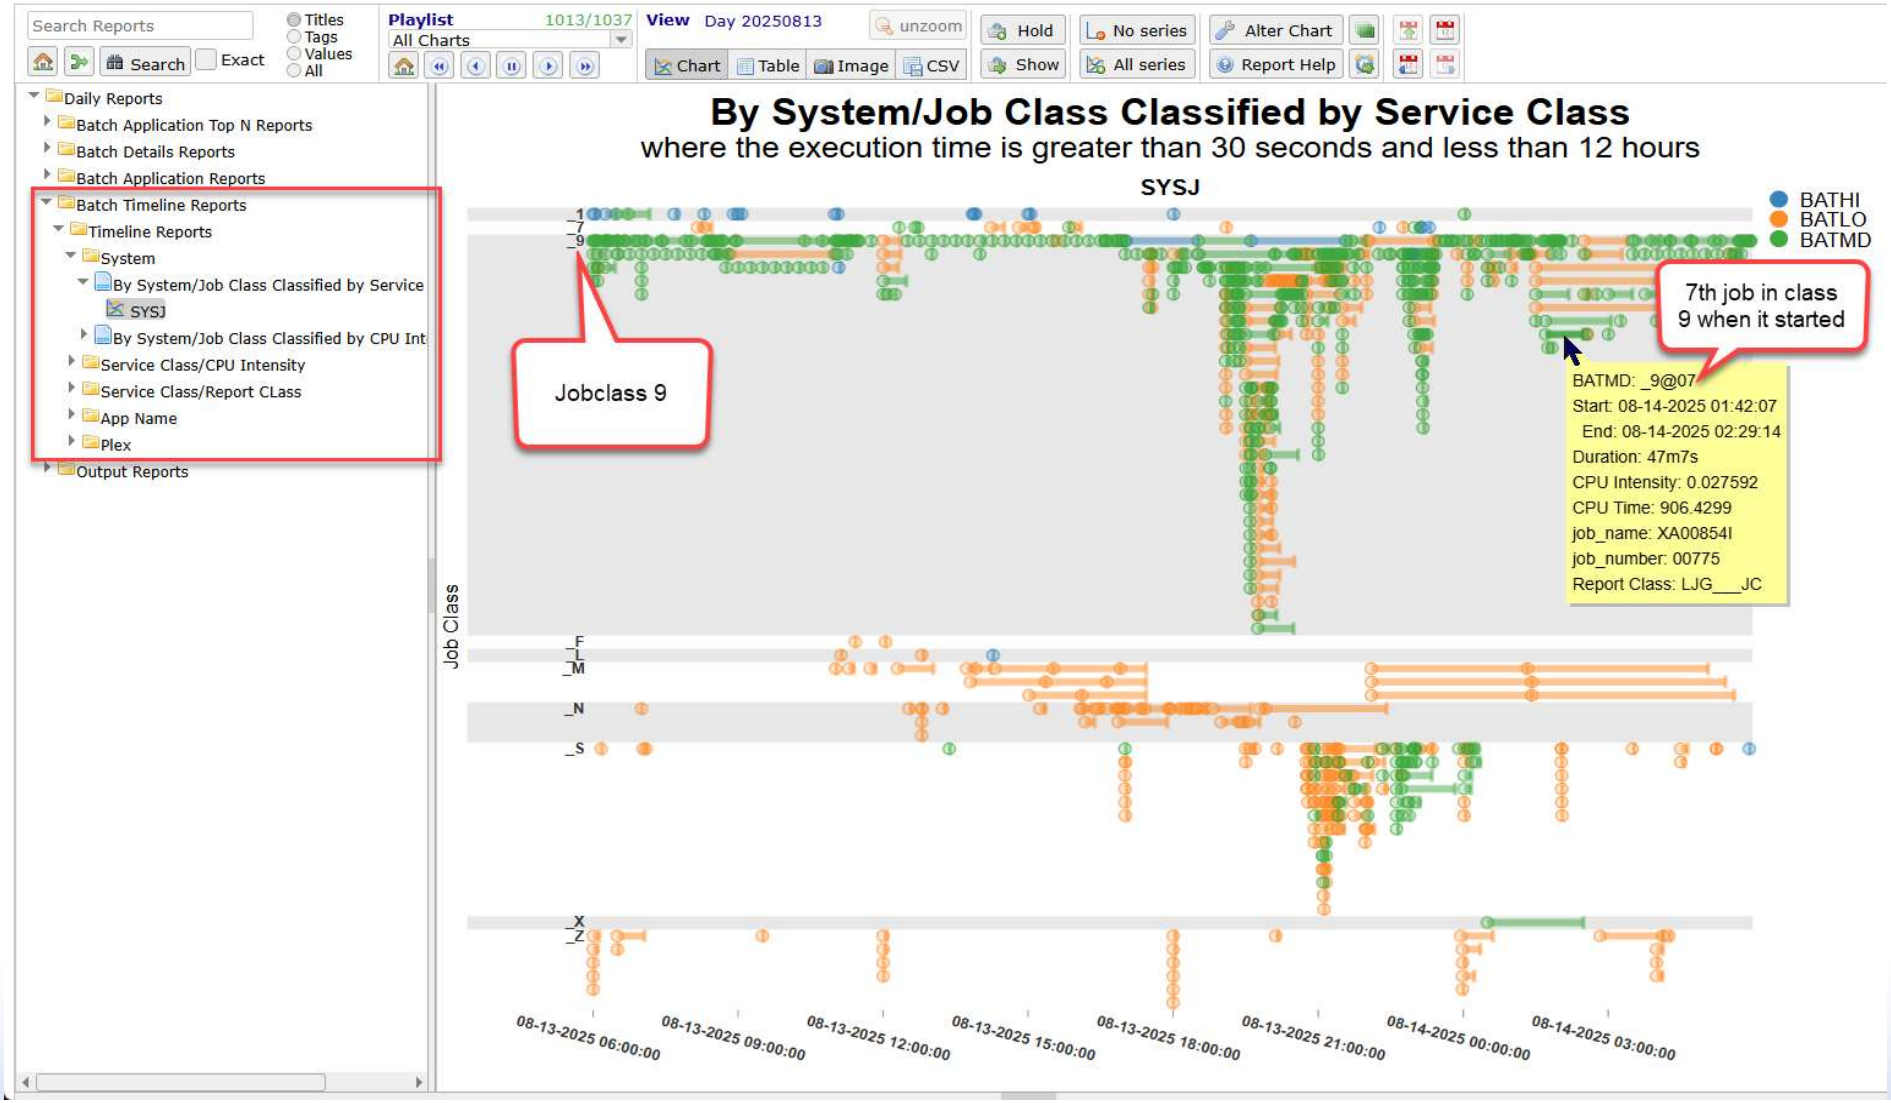
Search Reports
Titles
Tags
Values
All
Search
Exact
Playlist
1013/1037
All Charts
View
Day 20250813
unzoom
Chart
Table
Image
CSV
Hold
Show
No series
All series
Alter Chart
Report Help
Daily Reports
Batch Application Top N Reports
Batch Details Reports
Batch Application Reports
Batch Timeline Reports
Timeline Reports
System
By System/Job Class Classified by Service
SYSJ
By System/Job Class Classified by CPU Int
Service Class/CPU Intensity
Service Class/Report CLass
App Name
Plex
Output Reports
By System/Job Class Classified by Service Class
where the execution time is greater than 30 seconds and less than 12 hours
SYSJ
BATHI
BATLO
BATMD
Jobclass 9
7th job in class 9 when it started
BATMD: _9@07
Start: 08-14-2025 01:42:07
End: 08-14-2025 02:29:14
Duration: 47m7s
CPU Intensity: 0.027592
CPU Time: 906.4299
job_name: XA00854I
job_number: 00775
Report Class: LJG___JC
Job Class
_1
_7
_9
_F
_L
_M
_N
_S
_X
_Z
08-13-2025 06:00:00
08-13-2025 09:00:00
08-13-2025 12:00:00
08-13-2025 15:00:00
08-13-2025 18:00:00
08-13-2025 21:00:00
08-14-2025 00:00:00
08-14-2025 03:00:00

Pivotor has lots of daily batch reporting

Daily Batch by CPU Intensity

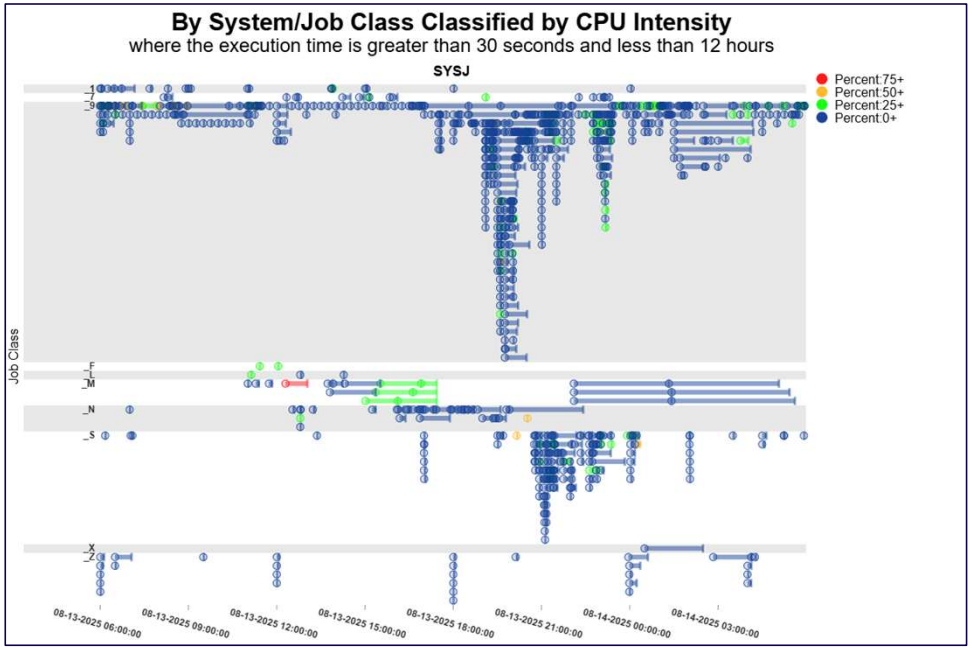

Daily Batch by Report Class

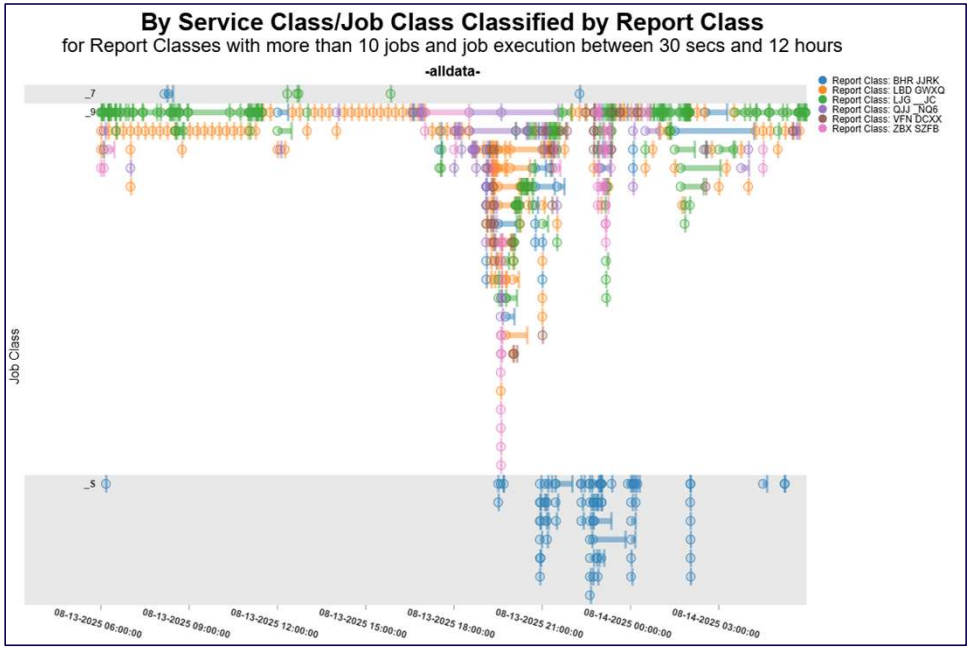

Other New/Updated Pivotor Reporting

- New system-level I/O reports
  - Including those that marry SMF 42s with DCOLLECT
- Lots of new reporting around the high frequency SMF 98s and 99s
  - Including enhanced lock reports
  - Including WLM trace and decision reports
- New “experimental” “suffering” reporting
  - Intelligent reports to help identify potential problem areas (See Scott’s session from last SHARE)
- Outlier reports
  - Sophisticated and intelligent
- Many new and unique Pivotor reports
  - Many generated from combining various record / data types
  - Many with creative formulas
  - All sorts of customer requested reports

Other New/Updated Pivotor Reporting

- Regular updates for z/OS and z17
  - Pivotor is always being updated to report on the latest and greatest measurements
- New VSAM RLS reporting
- zHyperLink Reporting has been expanded
- Updates to CICS and Db2 reports
  - Always being updated and expanded
  - Designed to help you quicky zero in on potential problems and areas of high impact
- Support for IBM IMS log records
  - This is in addition to the BMC IMS log records
  - Has both short- and long-term reporting
- Support for IDMS SMF records
  - Including transactional reporting

Something we didn’t know…

(And we’re a bit embarrassed by it)

I/O Time in SMF30 is Wrong

> **Note:** The system adjusts the connect time for FICON DASD to be 1 millisecond per request. This value differs from the channel reported connect time.

- Connect time is (historically) data transfer time
  - FICON's introduction of concurrent active I/Os per channel complicated this measurement
  - Apparently, the complication was too much for the SMF 30 records?
- In the early days of FICON 1ms might have been more reasonable than today
  - Today 0.1-0.2ms might be closer to accurate
- Makes comparing jobs by I/O time is problematic
  - Transfer times can of course vary based on size of the I/Os
    - Which may vary quite a bit for batch work
  - 1ms significantly higher than common FICON DASD total RTs today

Related: I/O is Complicated

- We came across this following up to a question from Scott's "z/OS I/O Performance: Do You Have a Problem?"
  - Thanks for your questions and feedback!
- Summary of that presentation:
  - Some of the measurements are confusing
  - I/O problems can hide in aggregated measurements
  - It's hard (essentially impossible) to know what I/O is really important to you
  - I/O reduction efforts are undervalued today
- See also "I/O, I/O It's Home to Memory We (Should) Go"
  - zHyperLink is fast, but memory is faster!

https://www.pivotor.com/content.html

Speaking of weird measurements…

How much power is your mainframe using?

- There are now (somewhat surprisingly granular) energy consumption measurements in the SMF 70 records

SMF 70 Field	Description
SMF70_CPUPower	Accumulated microwatts for all CPU resources allocated to the LPAR during the interval.
SMF70_StoragePower	Accumulated microwatts for all storage resources allocated to the LPAR during the interval.
SMF70_IOPower	Accumulated microwatts for all I/O resources allocated to the LPAR during the interval.
SMF70_CPCTotalPower	Accumulated microwatts for all electrical and mechanical components in the CPC.
SMF70_CPCUnassResPower	Accumulated microwatts for all types of resources in the standby or reserved state.
SMF70_CPCInfraPower	Accumulated microwatts for all subsystems in the CPC which do not provide CPU, storage, or I/O resources to logical partitions. These include service elements, cooling systems, power distribution, and network switches, among others.
SMF70_NumPowerSamples	Number of power samples for the interval.
SMF70_PowerPartitionName	The name of the LPAR to which the LPAR-specific power fields apply.

How much power is your mainframe using?

- Energy consumption measurements in the SMF 70 records
  - Requires z17
  - Recorded at the CEC level
  - IBM has stated that the z17 uses 19% less power than z16

- Audience poll: does anybody care about this data?
  - Other than the abstract/academic point of "that's interesting trivia"
  - Is somebody pushing you to quantify your datacenter power consumption by device?
  - Do you think this could meaningfully improve the perception of the mainframe?
  - Is the power costs of a z machine a big issue for you?
  - Do other devices, such as storage controllers, record power measurements?

  - That is… should Pivotor bother reporting on this?

Java Times

- There are a variety of new Java related measurements in the SMF 30, 70, 99
  - We are not exactly sure why these measurements were added

New SMF Fields for Java	Description
SMF30_Time_Java_On_zIIP	Time spent in Java work on zIIP including enclave time.
SMF30_ENCLAVE_Time_Java_On_zIIP	Enclave time spent in Java work on zIIP
SMF30_DEPENC_Time_Java_On_zIIP	Dependent enclave time spent in Java work on zIIP.
SMF30_Time_Java_On_CP	zIIP eligible time spent in Java work on CP including enclave time.
SMF30_ENCLAVE_Time_Java_On_CP	zIIP eligible enclave time spent in Java work on CP
SMF30_DEPENC_Time_Java_On_CP	zIIP eligible dependent enclave time spent in Java work on CP
SMF70JSU_on_zIIP	Unweighted zIIP-eligible Java service units spent on zIIP for the entire system.
SMF70JSU_on_CP	Unweighted zIIP-eligible Java service units spent on CP for the entire system.
SMF99_SUS_JAVA_ON_ZIIP	Total unweighted zIIP-eligible Java service units spent on zIIP.
SMF99_SUS_JAVA_ON_CP	Total unweighted zIIP-eligible Java service units spent on CP.

Java Measurement Examples

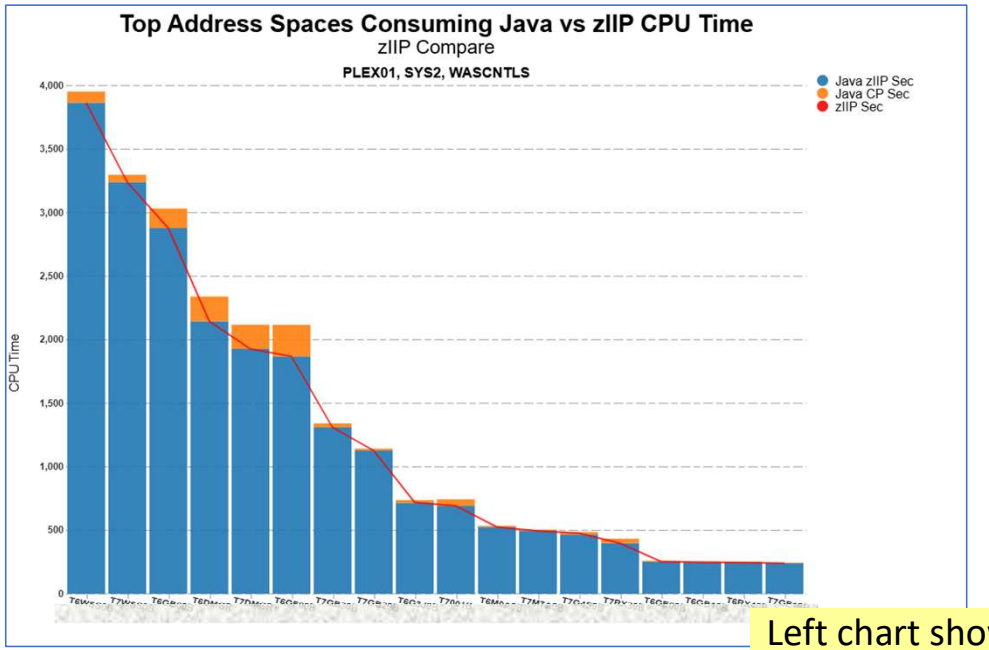

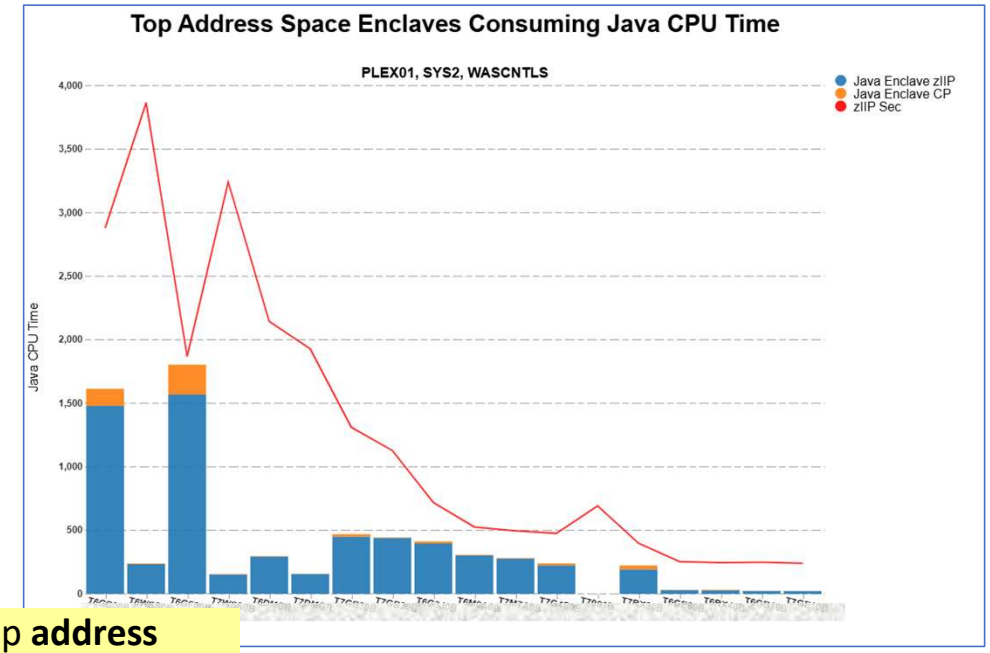

Left chart shows top **address space** Java time compared to zIIP time.

Right chart shows top **address space enclave** Java time compared to zIIP time.

Java Measurement Examples

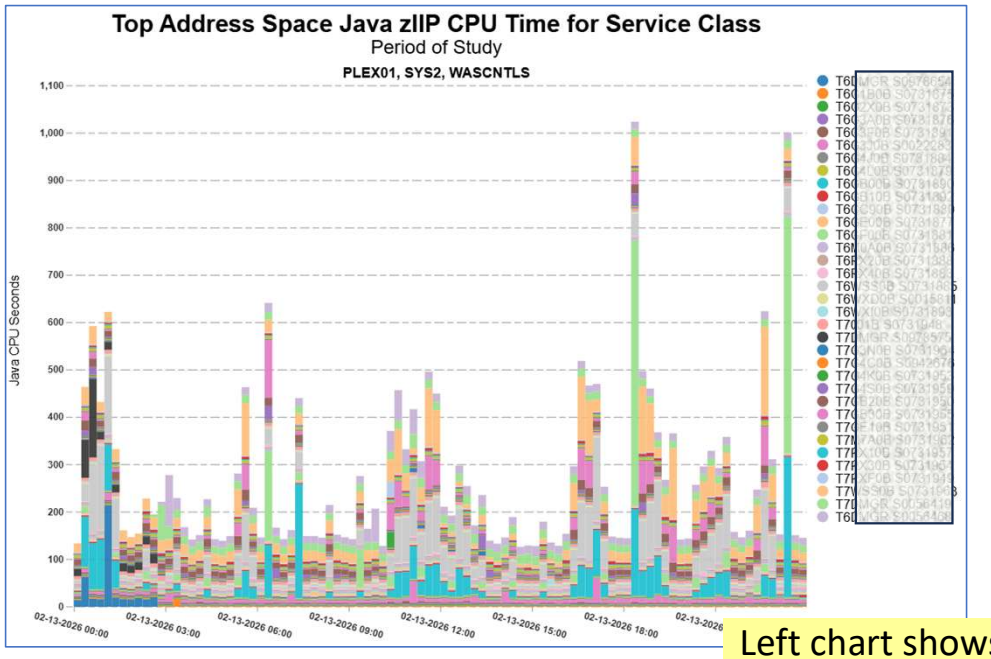

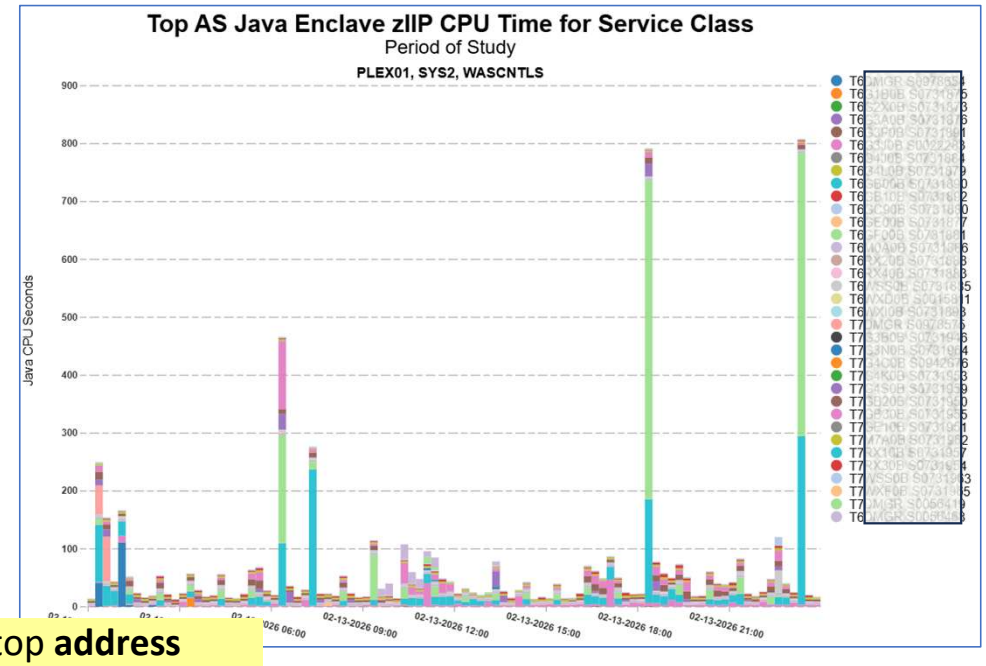

Left chart shows top **address space** Java time compared to zIIP time.

Right chart shows top **address space enclave** Java time compared to zIIP time.

Updated Recommendation

Updated SMF 98/99 Recommendations

- We've long recommended customers record the 99s and 98s
- Very useful for investigating transient performance problems (seconds to minutes)
- And the data isn't *that* large relatively speaking

Old

- Record 98s with HFTSINTVL that's greater than 5 seconds if you want
  - 15 sec interval: ~100-150 MB / system / day
- Minimally record SMF 99 subtypes 6, 10, 11, 12, and 14
  - ~100-150 MB / system / day

New

- Record the 98s with HFTSINTVL(5) (matches IBM's recommendation)
  - ~400-500MB / system / day
- Just go ahead and record all the SMF 99 subtypes
  - ~850-1500MB / system / day
  - Can save a little by turning off subtype 13

We believe the value of these records justifies the 1-2GB/system/day

Real example record sizes

Small Environment
1 System
5s SMF 98

Rec Type	Recs	Bytes
type 098	17,281	400,665,728
type 099	267,523	882,010,880
subtype 099_001	8,766	94,304,416
subtype 099_002	198,189	681,960,832
subtype 099_006	8,634	24,409,416
subtype 099_008	8,634	37,919,432
subtype 099_012	43,014	43,272,084
subtype 099_014	286	144,716
type 030	75,242	127,542,360
type 072	8,352	12,758,400
type 074	3,840	109,555,584
type 119	263,496	91,755,080

Medium Environment
1 System
15s SMF 98

Rec Type	Recs	Bytes
type 098	6,528	155,556,832
type 099	699,570	1,250,404,352
subtype 099_001	9,456	132,003,656
subtype 099_002	315,780	807,837,248
subtype 099_003	140,600	44,485,216
subtype 099_006	8,640	34,461,172
subtype 099_008	8,640	35,088,776
subtype 099_010	2	16,320
subtype 099_011	288	586,656
subtype 099_012	43,175	47,578,848
subtype 099_013	172,699	148,151,024
subtype 099_014	290	195,460
type 030	160,380	315,278,400
type 042	1,375,384	846,715,968
type 072	9,397	17,668,548
type 074	10,976	306,955,648
type 100	5,764	10,221,576
type 101	2,730,869	1,539,210,112
type 102	3,612	7,482,643
type 110	33,683	521,378,528
type 116	570,456	3,054,344,960

Large Environment
9 Systems
20s SMF 98

Rec Type	Recs	Bytes	Bytes/system
type 098	50,079	1,287,791,488	143,087,943
type 099	6,184,194	12,406,321,152	1,378,480,128
subtype 099_001	87,294	1,237,456,896	137,495,211
subtype 099_002	2,424,340	7,715,728,384	857,303,154
subtype 099_003	1,416,581	565,265,280	62,807,253
subtype 099_005	5,437	989,534	109,948
subtype 099_006	82,912	258,727,040	28,747,449
subtype 099_008	82,912	250,740,240	27,860,027
subtype 099_010	4	32,640	3,627
subtype 099_011	1,796	13,252,116	1,472,457
subtype 099_012	414,453	474,274,272	52,697,141
subtype 099_013	1,665,662	1,886,882,560	209,653,618
subtype 099_014	2,803	2,972,085	330,232
type 030	1,067,142	2,495,804,672	277,311,630
type 042	6,755,265	4,795,271,168	532,807,908
type 070	3,419	71,175,576	7,908,397
type 072	152,332	271,338,848	30,148,761
type 074	58,185	1,114,929,792	123,881,088
type 100	404,310	702,351,808	
type 101	307,838,528	172,787,793,920	
type 102	33,389,874	13,666,774,016	
type 110	53,572,512	465,821,106,176	
type 119	47,320,428	26,800,586,752	

CEC Physical Machine CP Busy% by CEC Serial Number

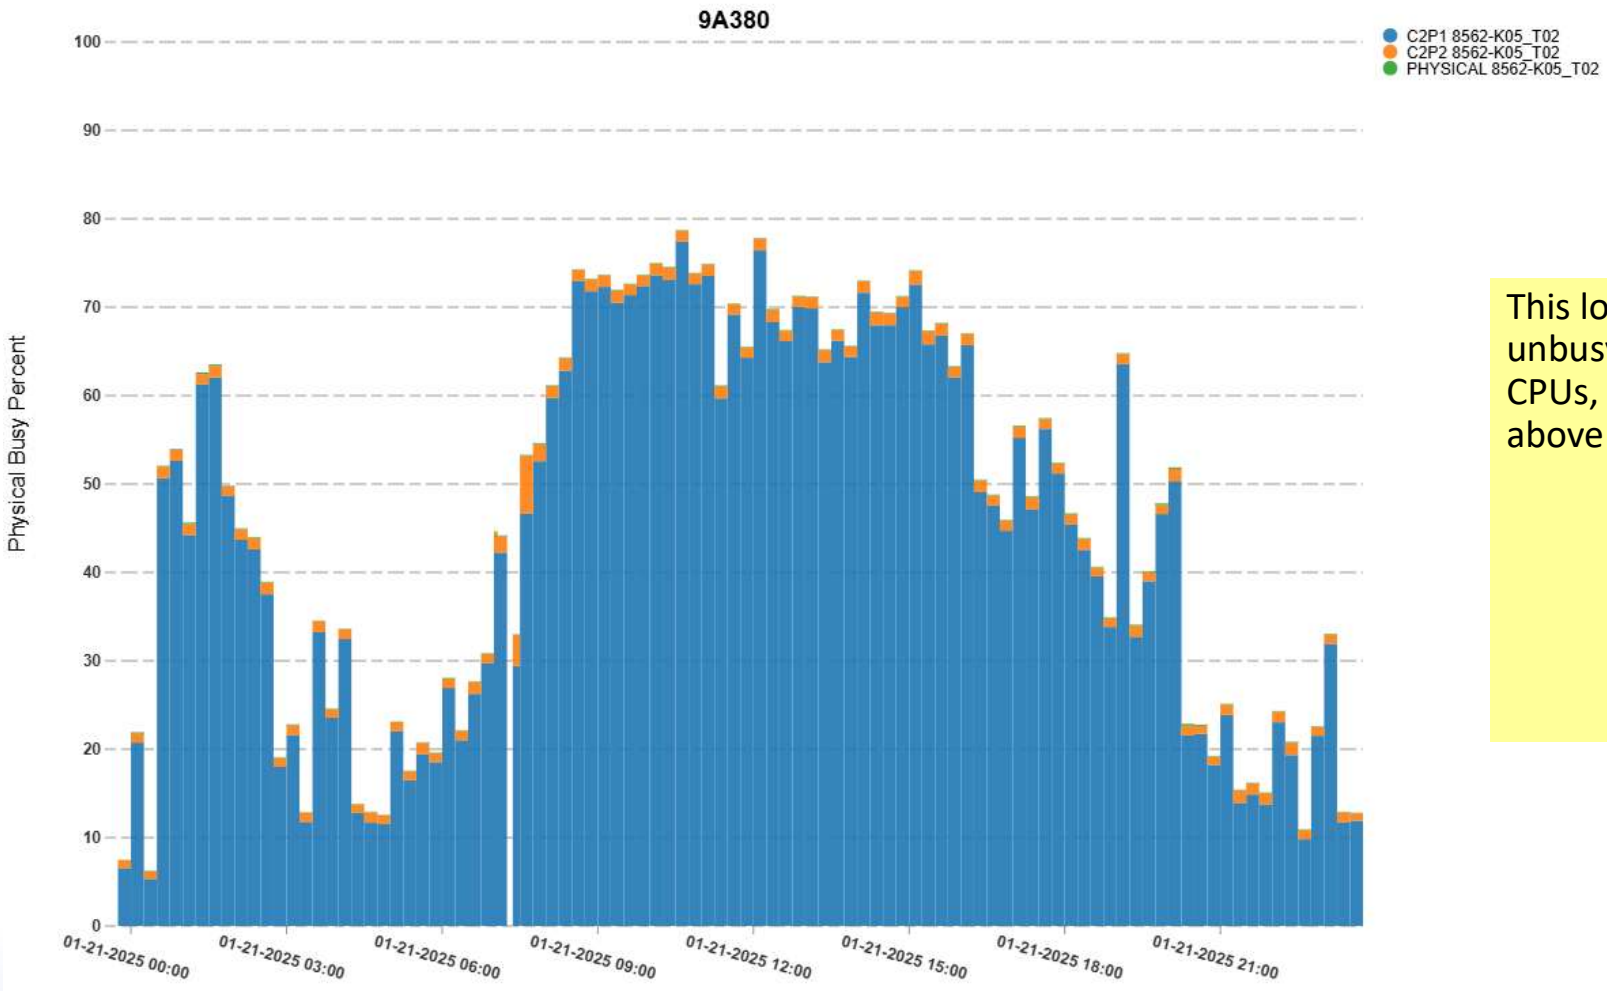

This looks like a pretty unbusy machine: 5 CPUs, never getting above 80% busy.

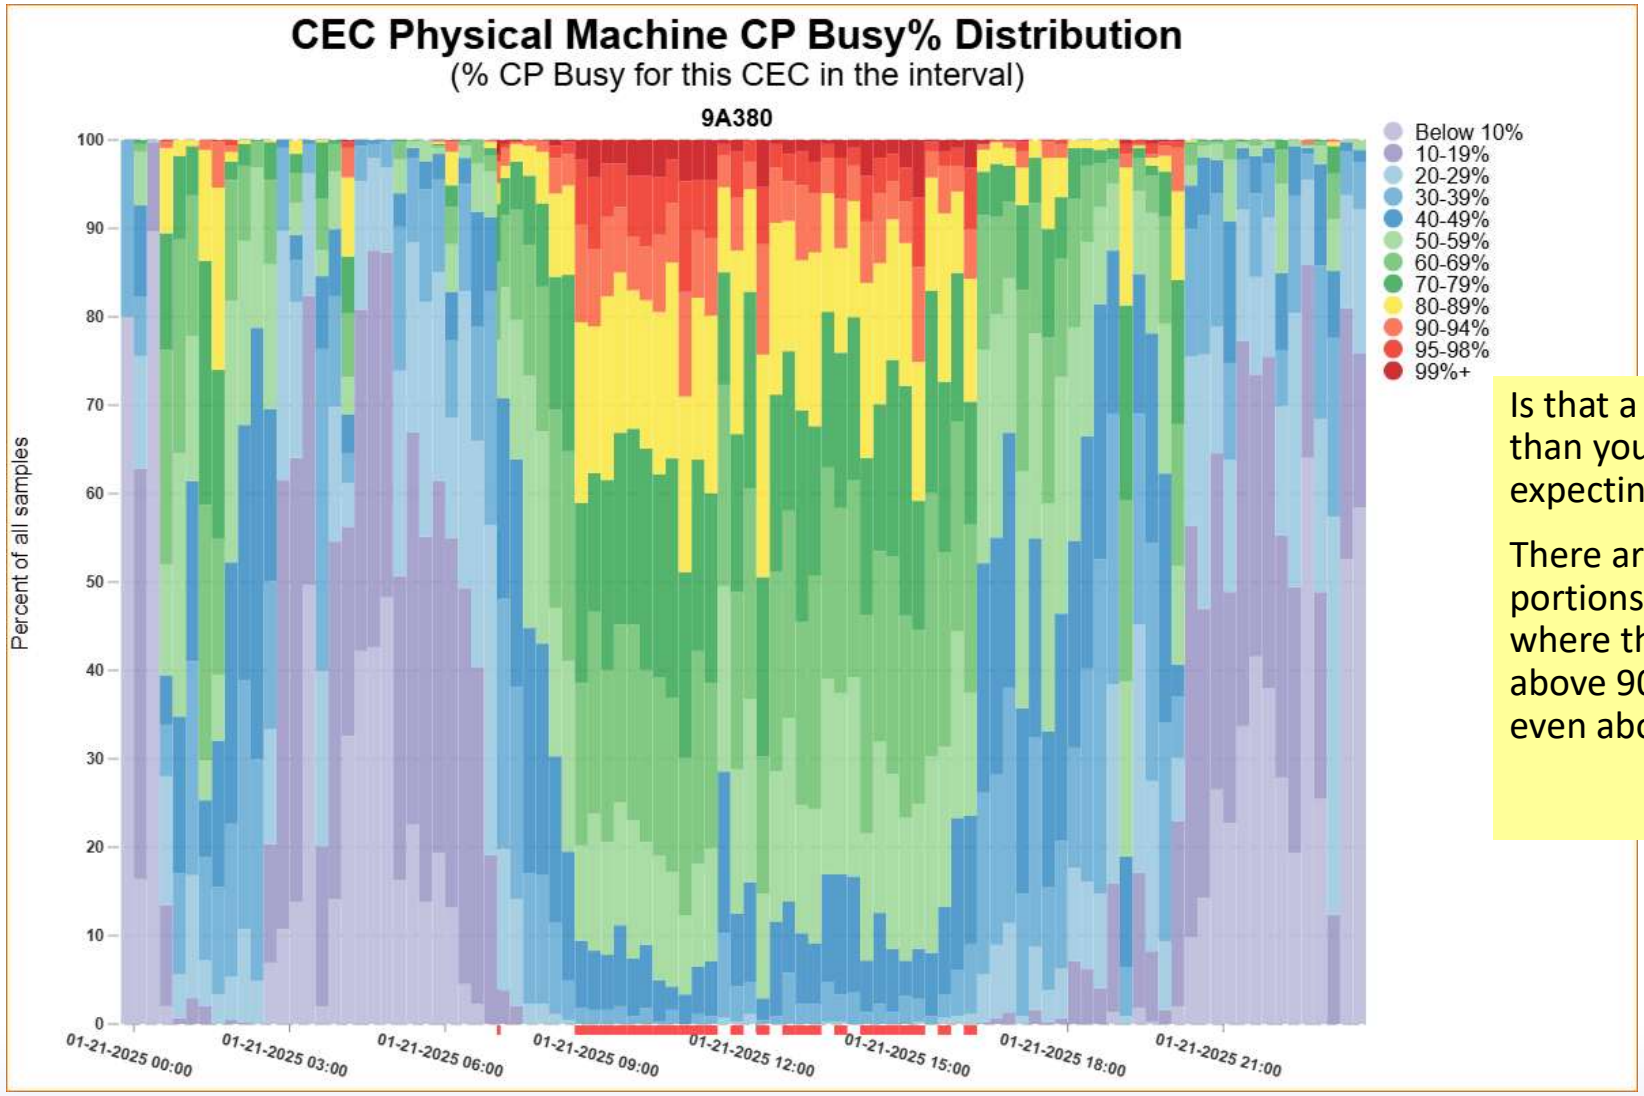

Is that a different picture than you were expecting?

There are significant portions of the intervals where the machine is above 90% busy, or is even above 99% busy!

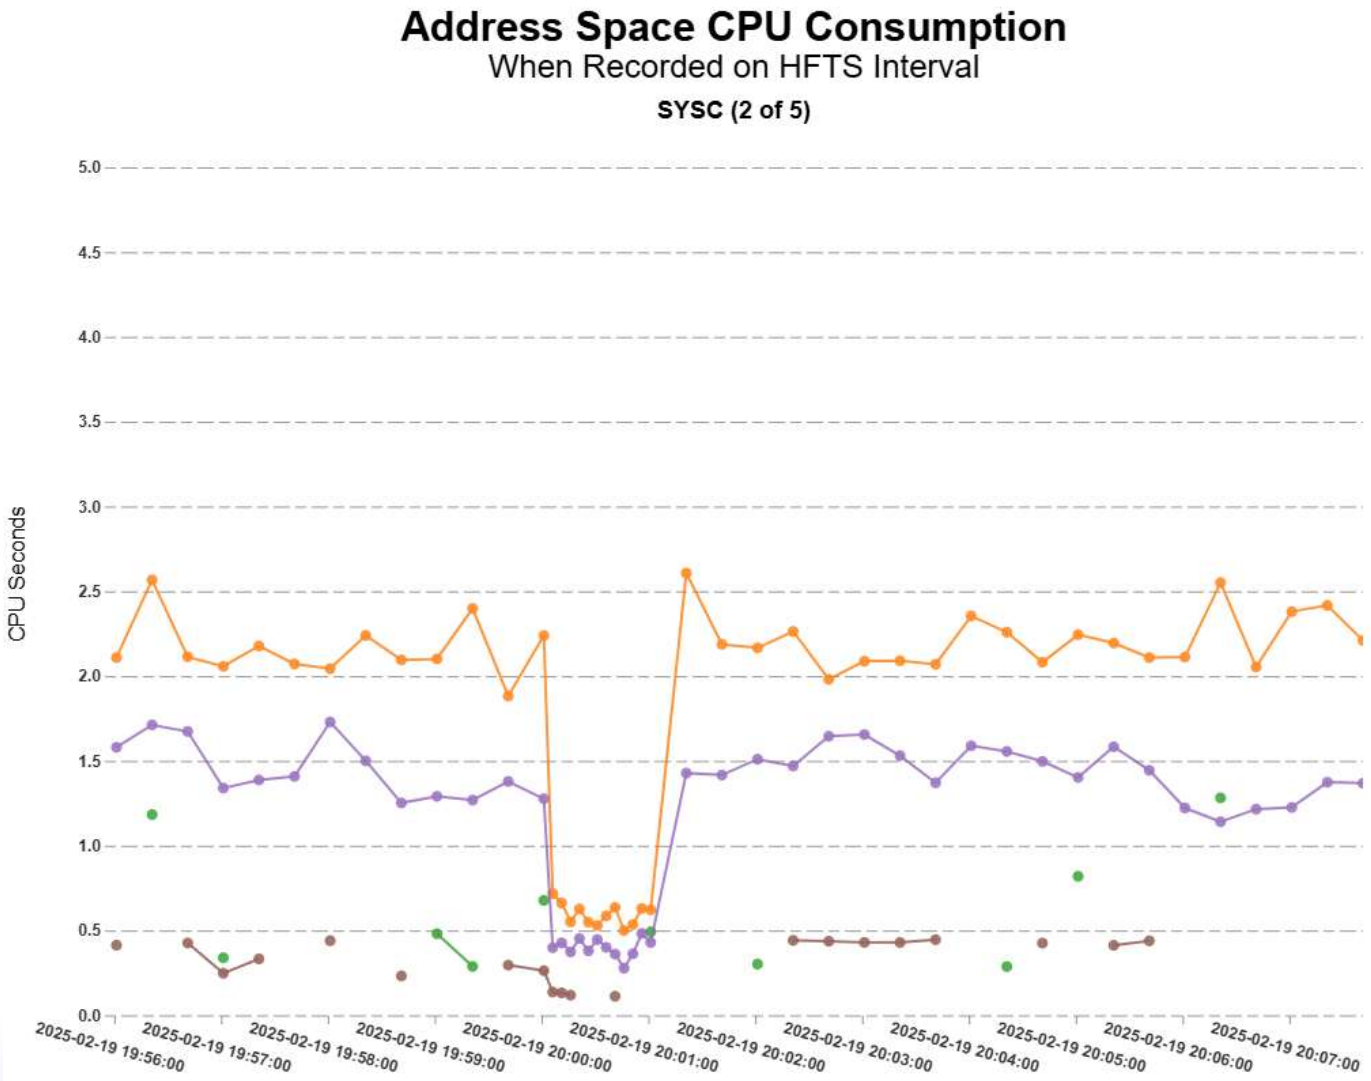

Address space CPU consumption on 20 and 5 second intervals.

That drop in in CPU used by TCP/IP and XCFAS by about 75% is due to the short intervals. (This is a good reason to use 5 second intervals.)

Also note VTAM and *MASTER* probably weren't idle when not shown… they just weren't recorded because they weren't selected to be recorded.

WLM PA - Receiver Candidate Occurrences

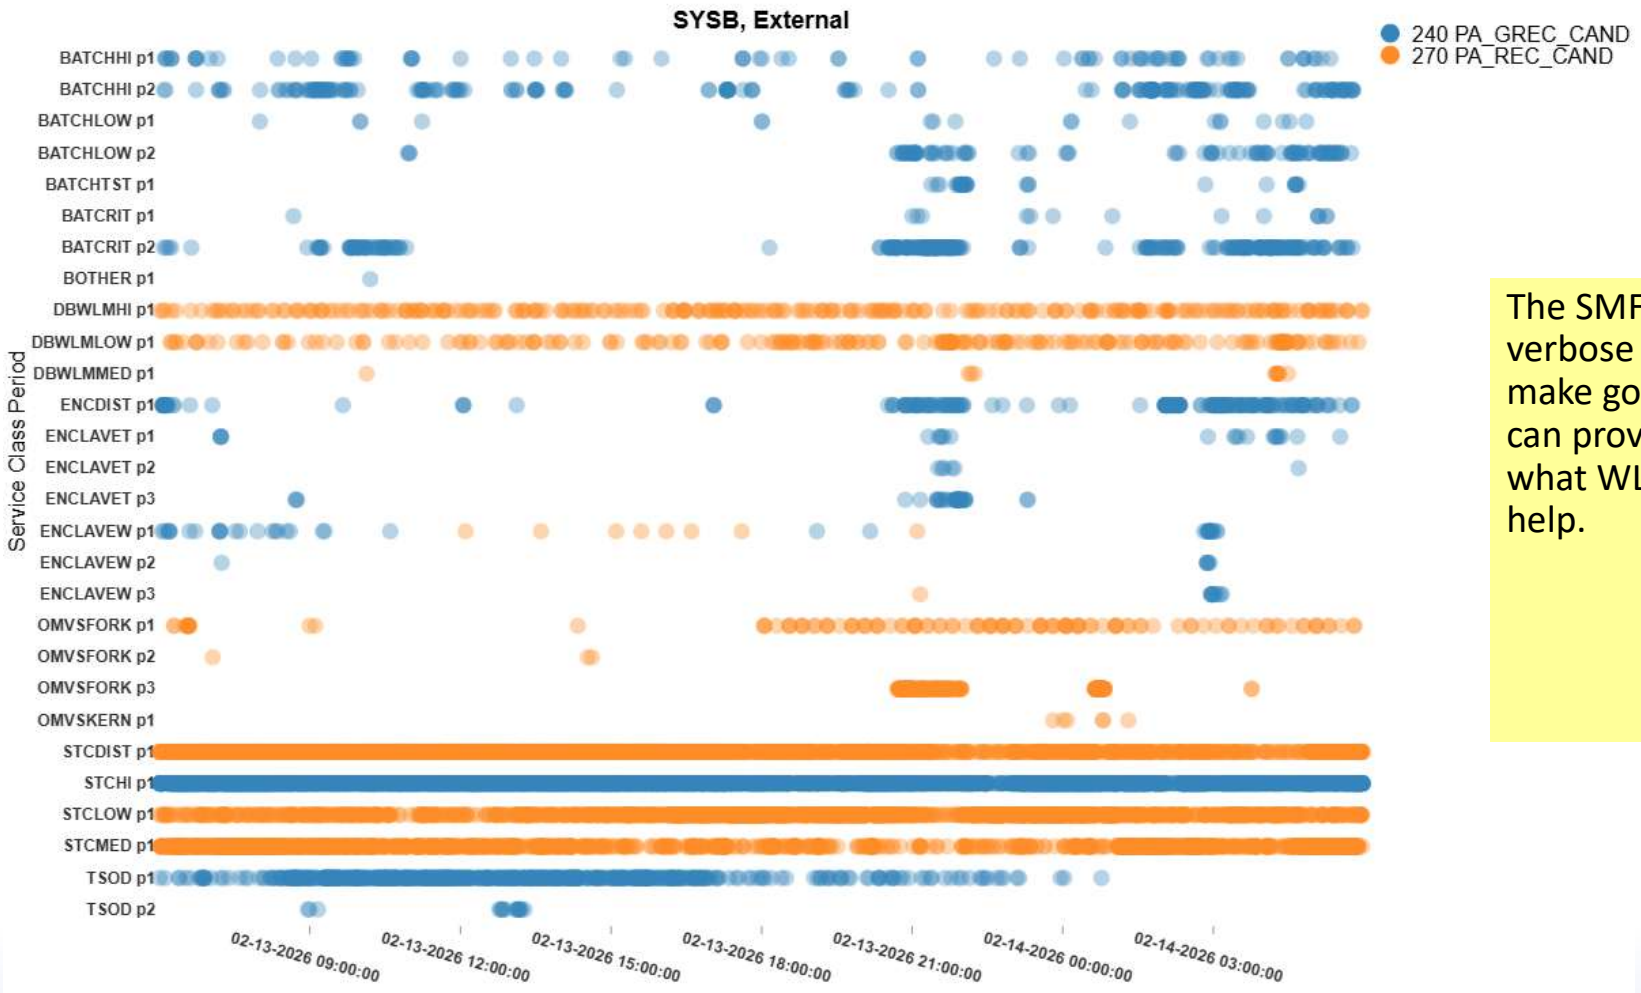

The SMF 99.1 data is verbose and difficult to make good use of, but it can provide insight into what WLM is trying to help.

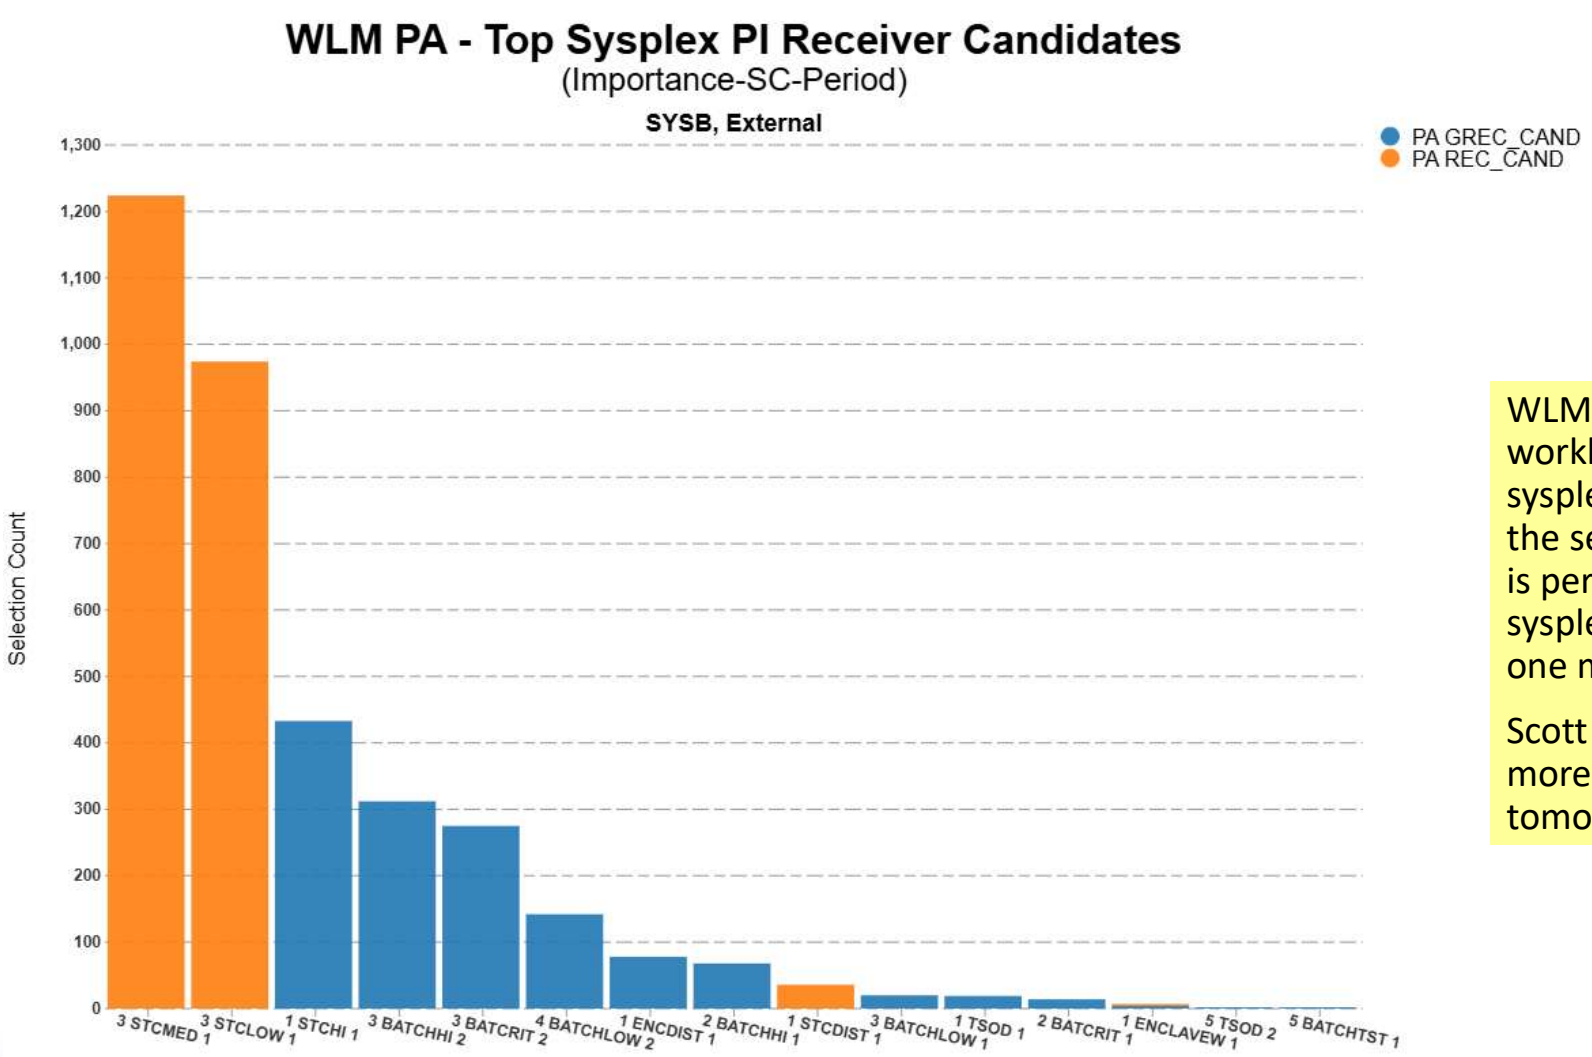

WLM might be helping workloads to help the sysplex PI (i.e. how well the service class period is performing across the sysplex) more often than one might think.

Scott will discuss this more in his presentation tomorrow.

Unsurprising recommendation

- You need a performance reporting / investigation tool at the ready
  - If all you have is RMF/CMF post processor reports, that's not great
    - And it only helps for a subset of the SMF performance data
  - Whatever you have you should use it on a semi-regular basis to be familiar with it
- Performance problems can have a real cost to the business, be prepared
- We're seeing more sites outsourcing their z/OS environment
  - There's of course a wide range of arrangements and outcomes, but…
  - Performance management in these environments can suffer
  - Retaining your own oversight of z/OS performance is a good idea
    - Have heard from multiple sites where the outsourcer surprisingly didn't provide useful performance reporting

“Emerging” Areas of Interest

Time from announcement to 100% adoption is … long

Not everybody should be first ☺

Adoption status across our customers

50% of machines are z16 or z17

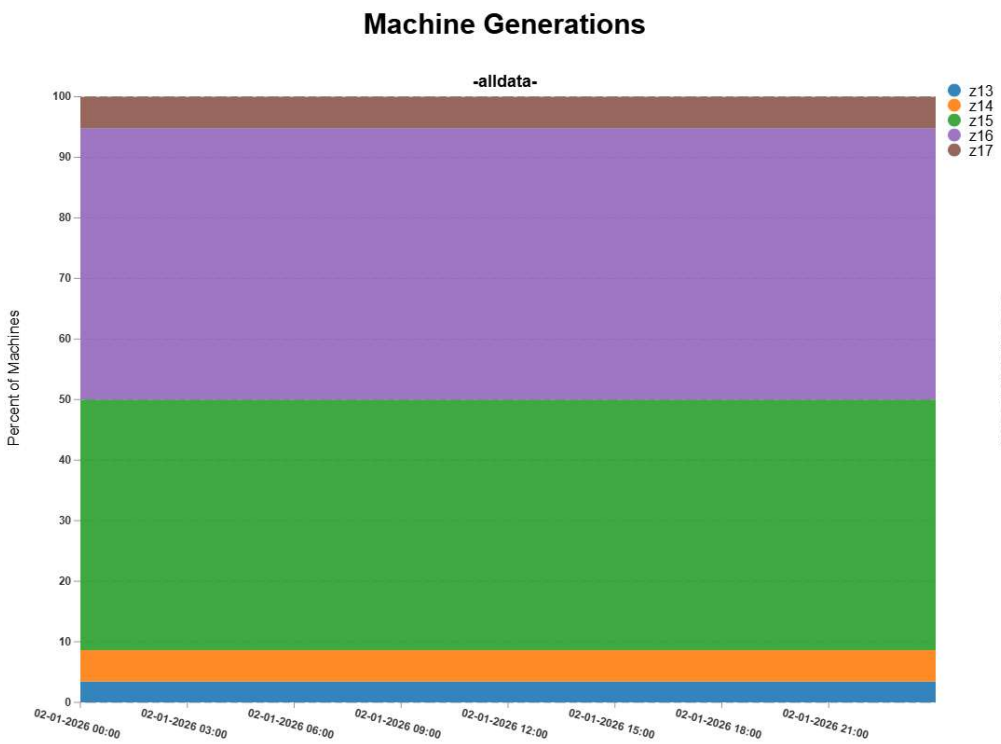

40% of systems not yet on z/OS 3.1

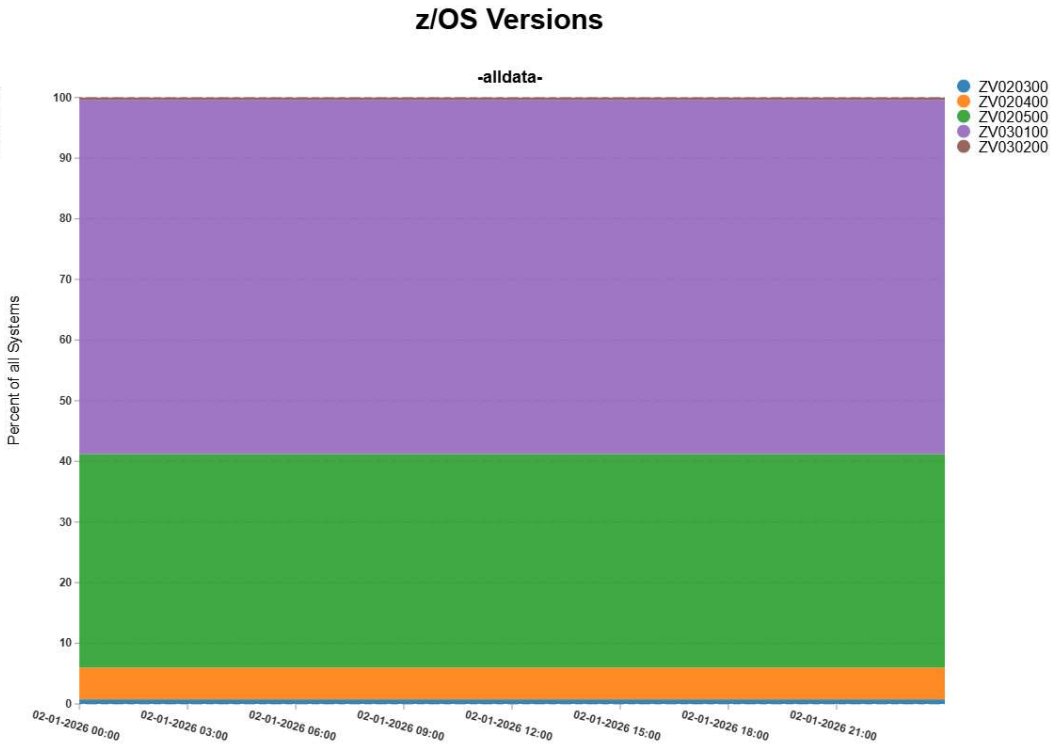

z17 Thoughts

- So far mostly “unsurprising” upgrade
  - First clock speed step-up in a while
  - Small cache size bump from z16 (significant change from z15)
  - Said to have slightly reduced cross-drawer penalty compared to z16
- Do have one report from a customer of a CICS transaction taking more CPU
- DPU is most interesting architectural difference
  - Moves certain FICON work from adapter card to processor Tellum II chip
  - DPU has 32 cores split across 4 clusters
  - Each DPU supports up to 64 PCHIDs (up to 16 adapters with up to 4 ports each)
  - Expectation for new DPU/adapter combination to more efficiently handle small I/O
  - New CPMF and SMF 73 measurements
  - For one customer doing ~100,000 IOPS: IOP busy down, IO RT essentially flat, Percent I/O requests retried up (although tiny numbers)

System Recover Boost – Are you using it??

A lot of people aren't using it as much as they could…

System Recover Boost

- Overall goal: provide additional capacity and performance to better recover from certain planned or unplanned events
  - Shutdown faster
  - Startup faster
  - Faster DR site switch
  - Faster sysplex recovery
  - Faster subsystem restart
- Requires z15 or later machine
  - + z/OS 2.3 and up on z15, z/OS 2.4 and up on z16
- No additional charge for basic System Recovery Boost
  - Optional extra cost item discussed in just a moment

Types of boost

- Speed Boost: *GP engines run as full speed engines*
  - Applicable only to sub-capacity systems
  - I.E. this is of no use to 8561-7xx and 8562-Z0x systems
  - Only boosted LPARs run as full speed!

- zIIP Boost: *zIIPs will run work not eligible to run on zIIPs*
  - I.E. ***any*** workload may be dispatched to a zIIP
  - Reserved (but physically available) zIIPs brought online to the boosted LPAR
  - Note that zIIPs always run full speed
  - What if you want more zIIPs for zIIP boost?

Boost classes = triggering events

Class	Where	Duration	WLM Work Routing
Shutdown	Single system	30 minutes	Routes work away from boosted system
IPL	Single system	60 minutes	Routes work to boosted system
Recovery Process	Multiple systems	<=5 minutes	No change

- Recovery Processes Boost
  - CF data sharing member recovery
    - Triggered by disconnection from lock structure while locks were held
  - Sysplex partitioning
    - System removed from a sysplex
  - CF structure recovery
    - Structure rebuild, duplex failover, or reduplexing
  - HyperSwap Recovery
    - Recovery from storage controller failover

New with z16:
- SVC Dump
- HyperSwap Configuration Load
- Middleware Region Startup
- Dynamic I/O Activation

More people should be using more SRB

- Boosts you get by default, without really taking action:
  - IPL
  - Recovery process for things like sysplex partitioning
- You must take action to get boosts for:
  - Shutdown
  - Middleware Recovery
  - SVC Dump boosts
- We're not seeing much of the latter—why??
  - If you're on TFP, all the more reason to think about this
- The extra capacity available to a boosted LPAR can be significant!

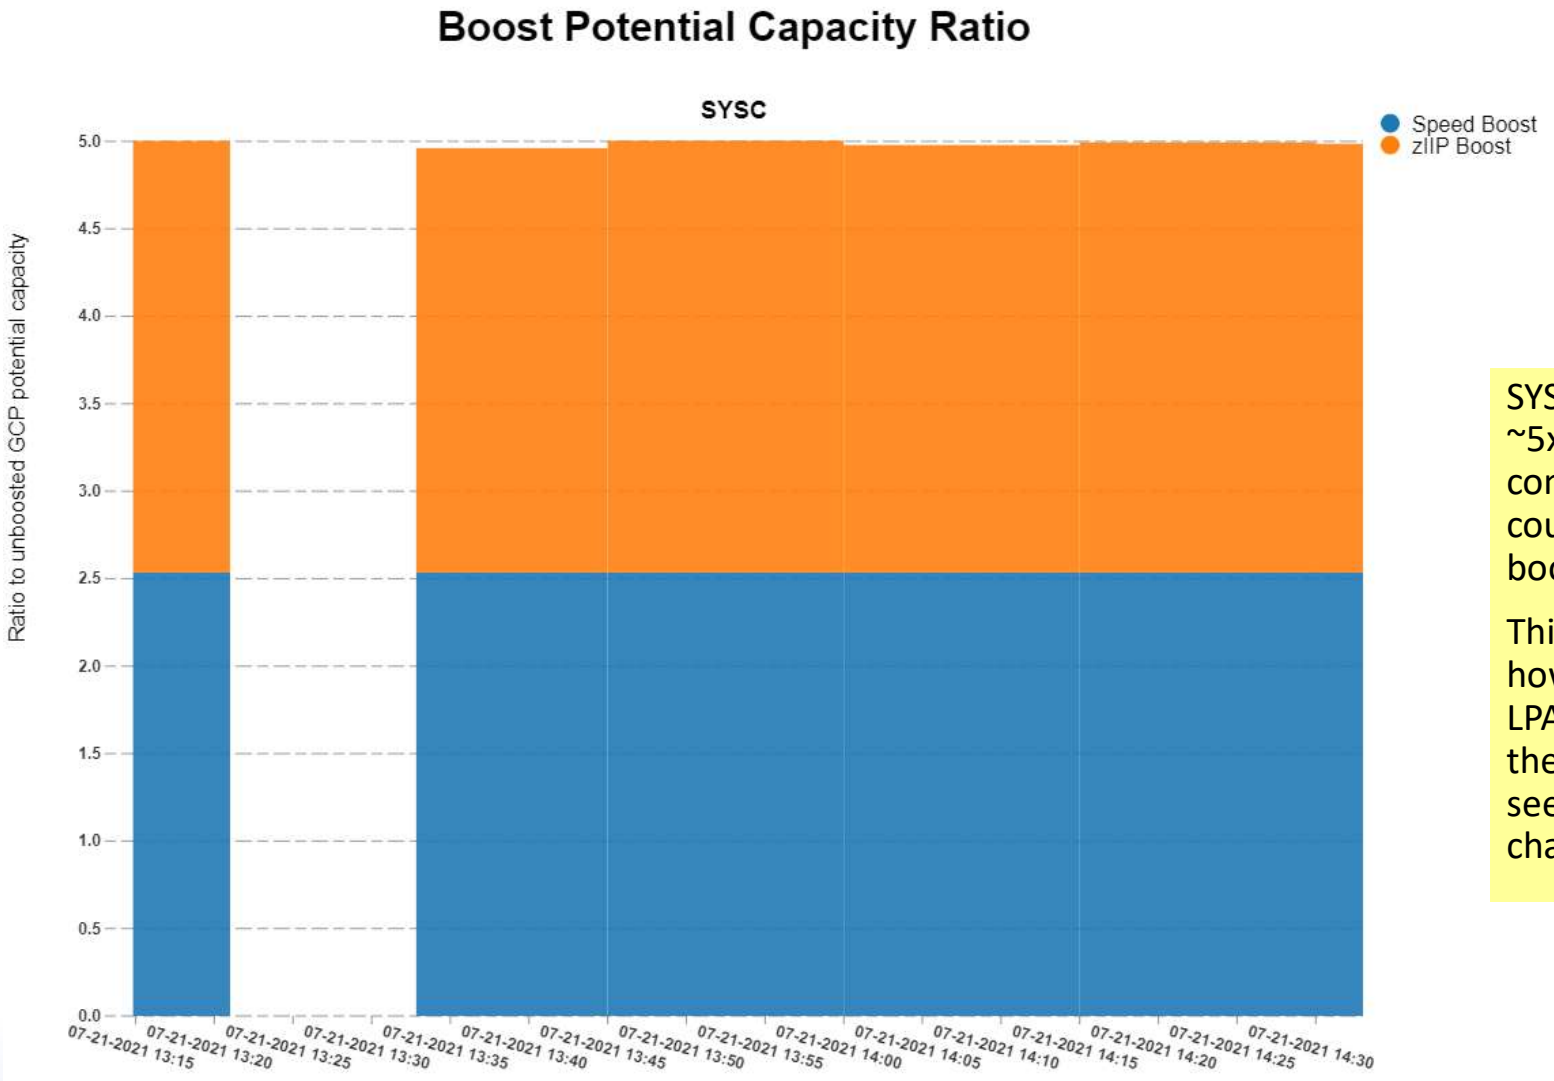

SYSC could have done ~5x the GP work compared to what it could have done un-boosted.

This report considers how much the other LPARs are using hence the slight variations you see in the ratios on each chart.

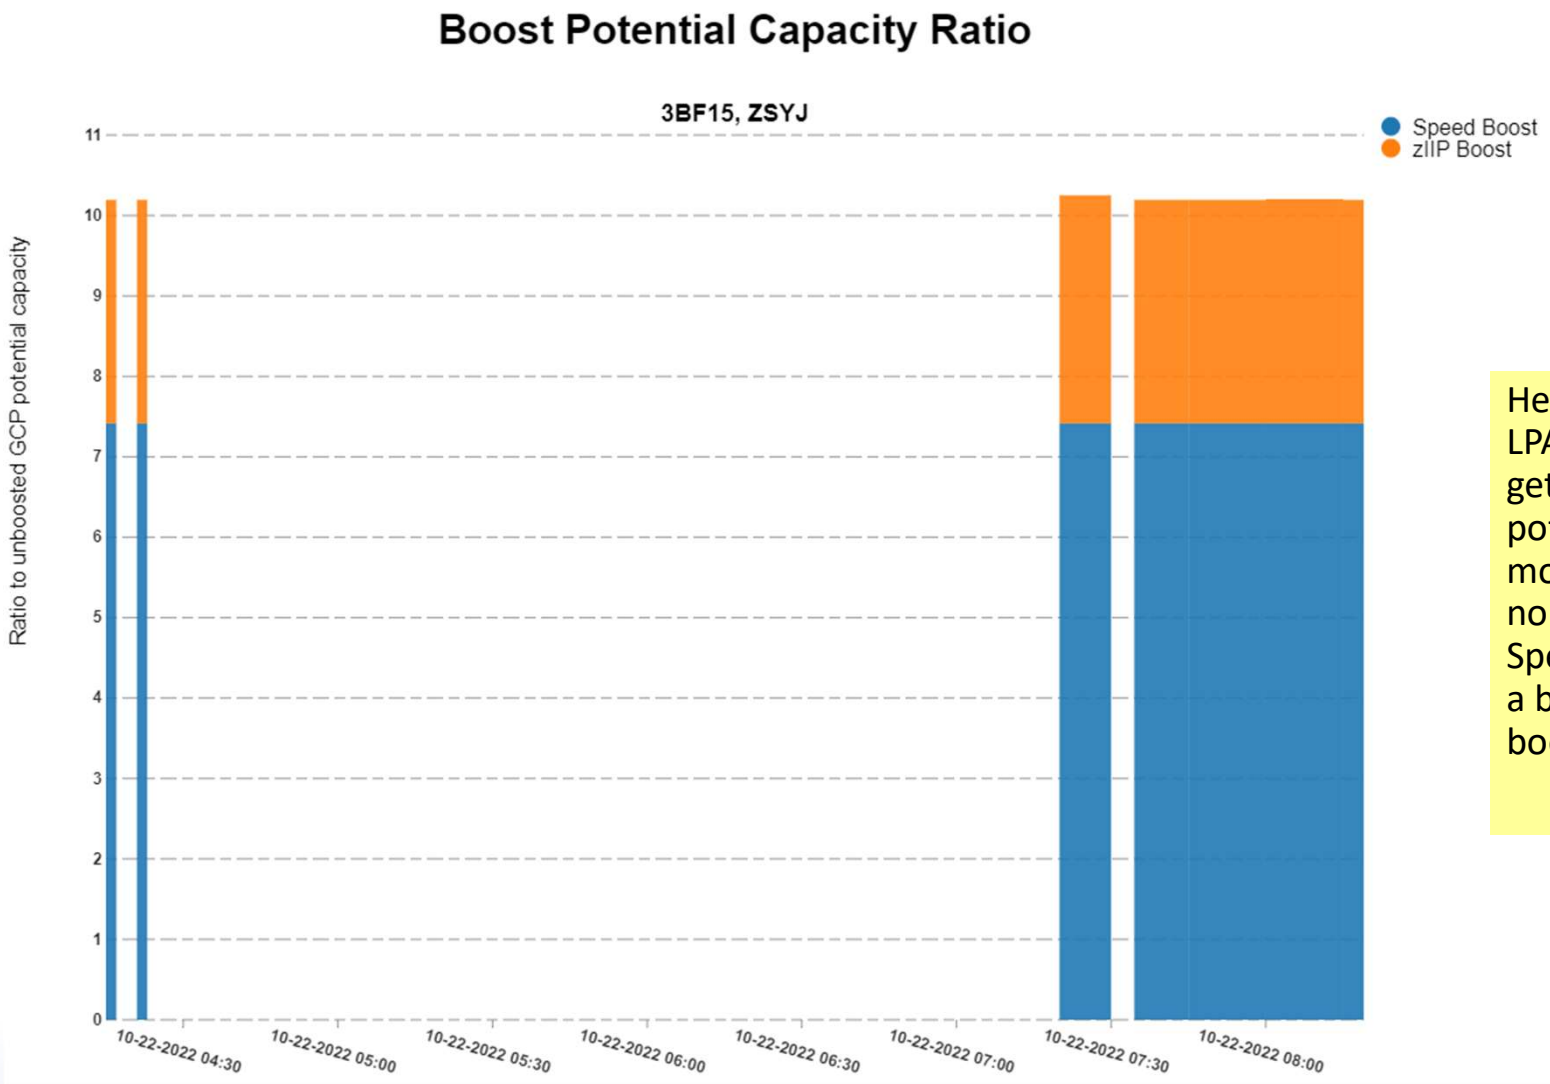

Here's an example of an LPAR on a z15-411 getting access to potentially over 10x more capacity than it normally has access to. Speed boost is obviously a big part of this, but zIIP boost is helping too.

LPARs *not* being boosted

- These LPARs *could* be negatively impacted
- CPU cache effectiveness may be impacted by boosted LPARs (at least theoretically)
  - Speed boosted LPARs do more work per unit of time on the GPs
  - zIIP boosted LPARs may drive more work to zIIPs
  - zIIP caps ignored during boost periods
- Higher physical zIIP utilization may impact non-boosted LPAR's ability to get work dispatched on zIIPs
  - Could potentially lead to more crossover
  - zIIP caps ignored but weight enforcement still applies
  - May need to change relative weights during boost periods, if trying to protect unboosted LPARs
    - Or (in theory) if trying to give more capacity to boosted LPAR
- Resource Group caps with sysplex scope don't count work running on boosted LPARs
  - May allow more than expected work to run on unboosted systems
- Conversely:
  - If boosted systems consume less of GP capacity, might help non-boosted systems

Weighty issues

- Boosts happen on a per-LPAR basis
- PR/SM dispatches physical CPs to LPARs as normal
  - Speed boosted LPARs simply get more useful work done in their dispatch interval
- Absent weight changes:
  - zIIP boosted LPARs with low weights may end up using low pool zIIPs
    - May not be able to access that capacity if other LPARs busy
  - Other LPARs using more than their zIIP weight may be limited if they are borrowing weight from the boosted LPAR
- So in some cases, you *may* want to consider setting up automation to make weight changes (particularly for zIIPs) during boost intervals

See also https://pivotor.com/content.html

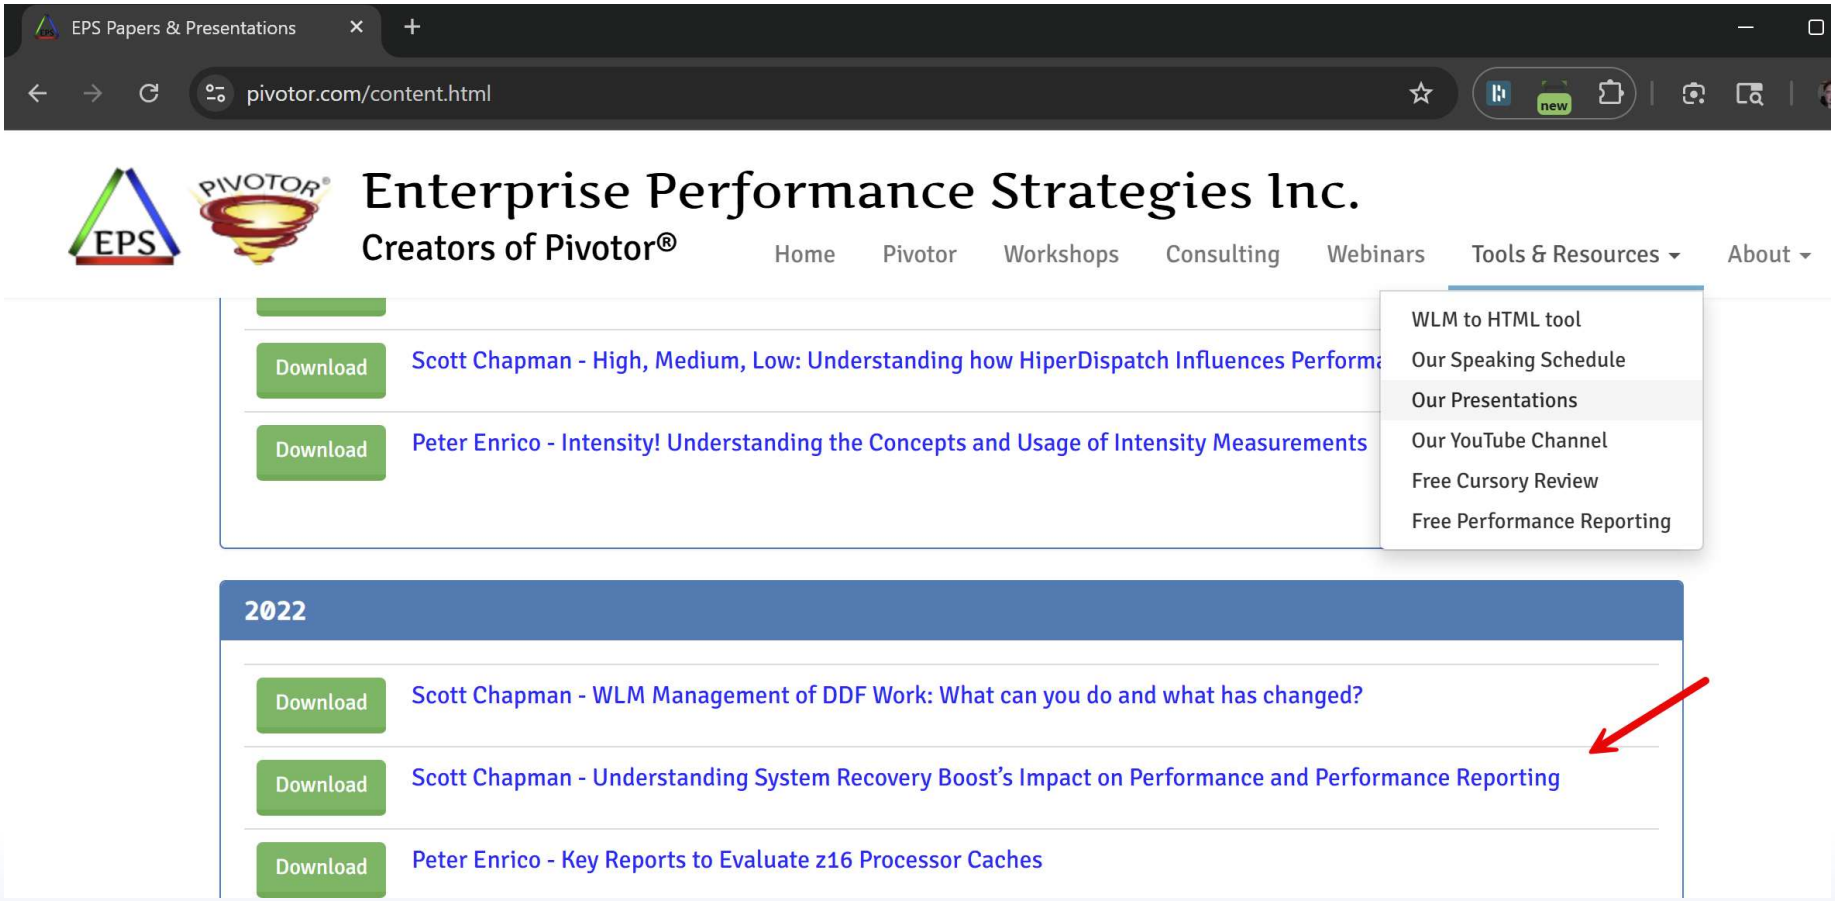

Implicit CPU Protection in z/OS 3.1

But not in 3.2 (they fixed it!)

CPU Critical aka Long-term CPU Protection

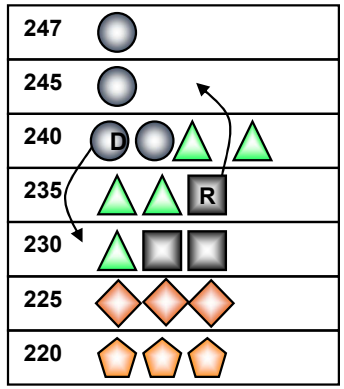

Note: To make the point, just a few priorities between DP 203 and DP247 are shown.

- 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 CPU DP 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.

New IBM Defaults in z/OS 3.1

- New option for “Implicit” Long-Term CPU Protection in IEAOPTxx
  - Default is “On” for importance 1 service classes and optional, but “Off” for importance 2 service classes

- CCImp=0|1|2
  - 0 : CPU protection not assigned to work unless it has been explicitly assigned for the SC in the WLM policy.
  - 1|2 :Specifies the importance level up to which CPU protection is implicitly assigned to the work

- 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
  - In 3.2 they reverted back to the historical default (yay IBM!)

- 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

WLM PA - CP Critical Movements

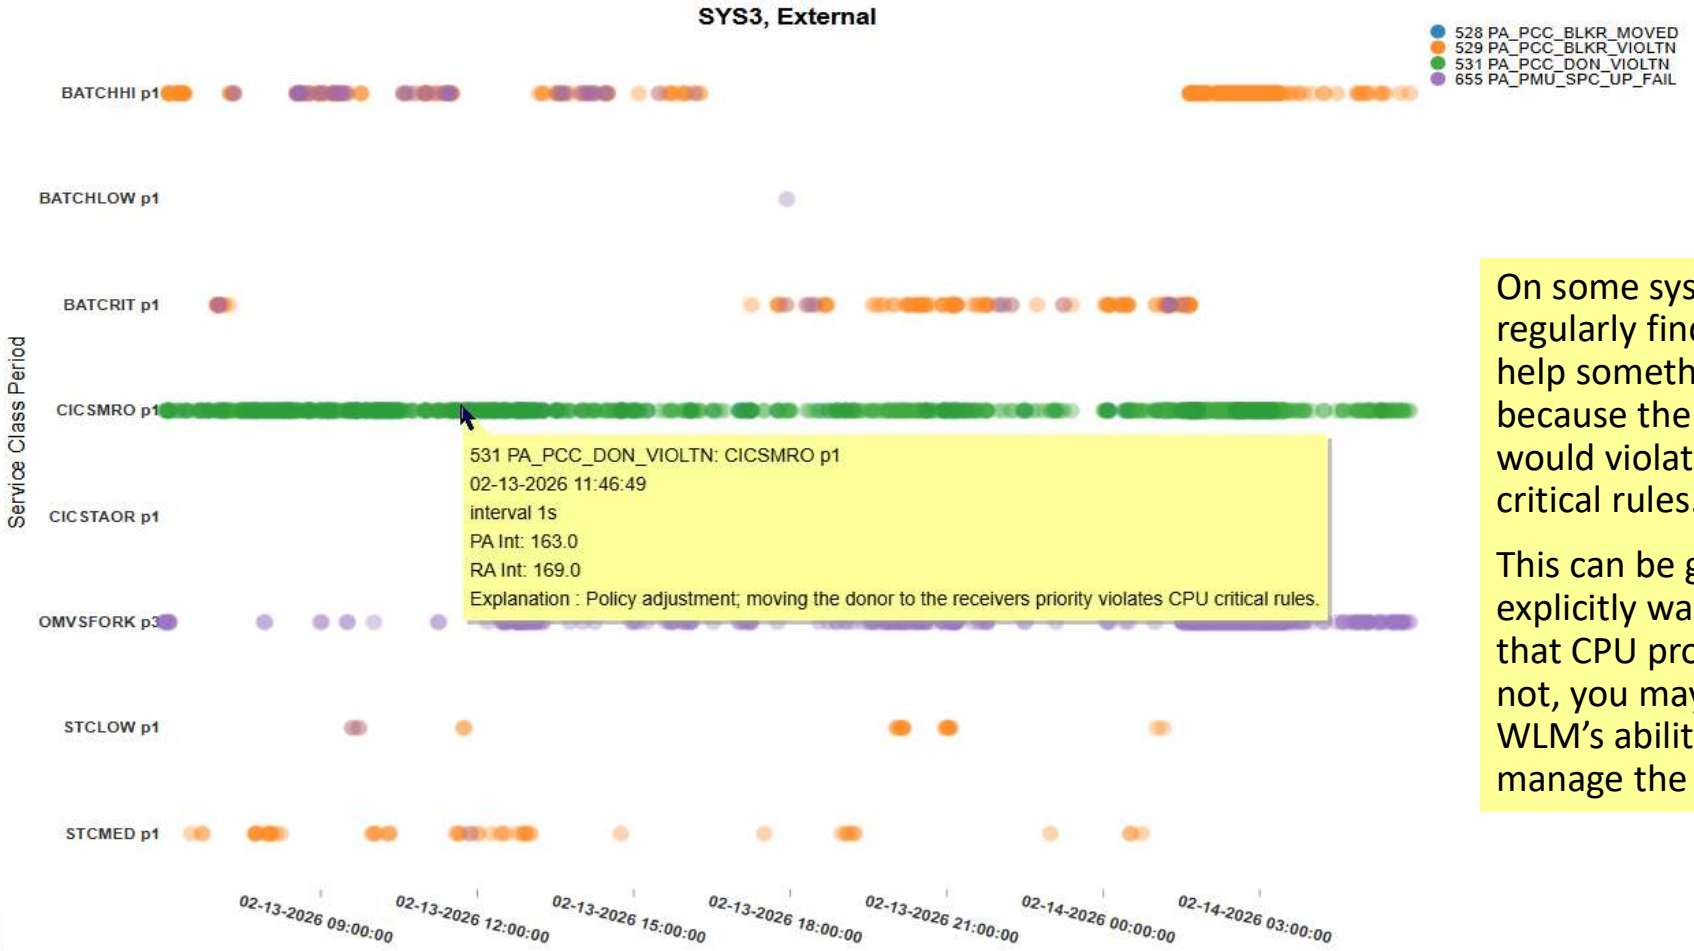

On some systems WLM regularly finds itself trying to help something and can't because the desired change would violate the CPU critical rules.

This can be good if you explicitly wanted/needed that CPU protection. But if not, you may be limiting WLM's ability to ideally manage the work.

Our thoughts

- 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
- Emerging area of study, we might refine our recommendations over time

Session feedback

- Submit your feedback using the Whoova app or at the conference website - https://conferences.gse.org.uk/2025/feedback/HI

- Make sure you are signed into Whoova app or MyGSE if using the website

- The session code is HI

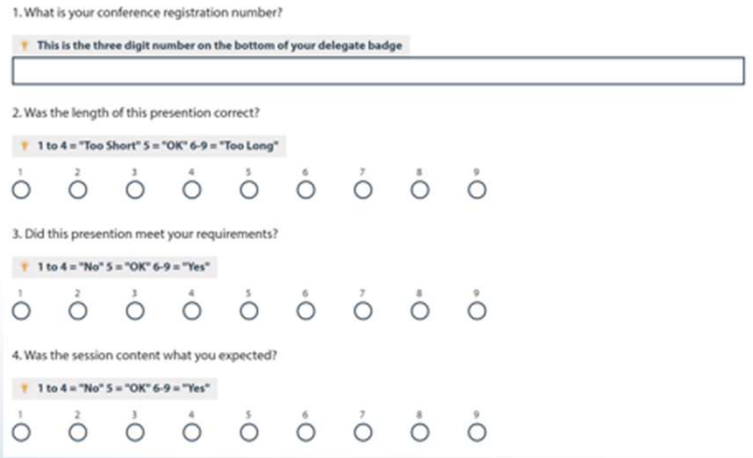

Wrap-up

- 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