/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)

Like what you see?

- Free z/OS Performance Educational webinars!
  - The titles for our Fall 2022-2023 webinars are as follows:
    - ✓ *Key Reports to Evaluate z16 Processor Caches*
    - ✓ *Understanding System Recovery Boost's Impact on Performance and Performance Reporting*
    - ✓ *WLM Management of DDF Work: What can you do and what has changed?*
    - ✓ *Intensity! Understanding the Concepts and Usage of Intensity Measurements*
    - ✓ *High, Medium, Low: Understanding how HiperDispatch influences performance in z/OS*
    - ✓ *How and why Pivotor is different than other performance management reporters*
    - *Putting a lid on XCF*
    - *Key Reports to Evaluate Usage of Parallel Access Volumes*
    - *Key Reports to Evaluate Coupling Facility CPU Utilization*
    - *Understanding how memory management has evolved in z/OS*
  - Let me 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

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

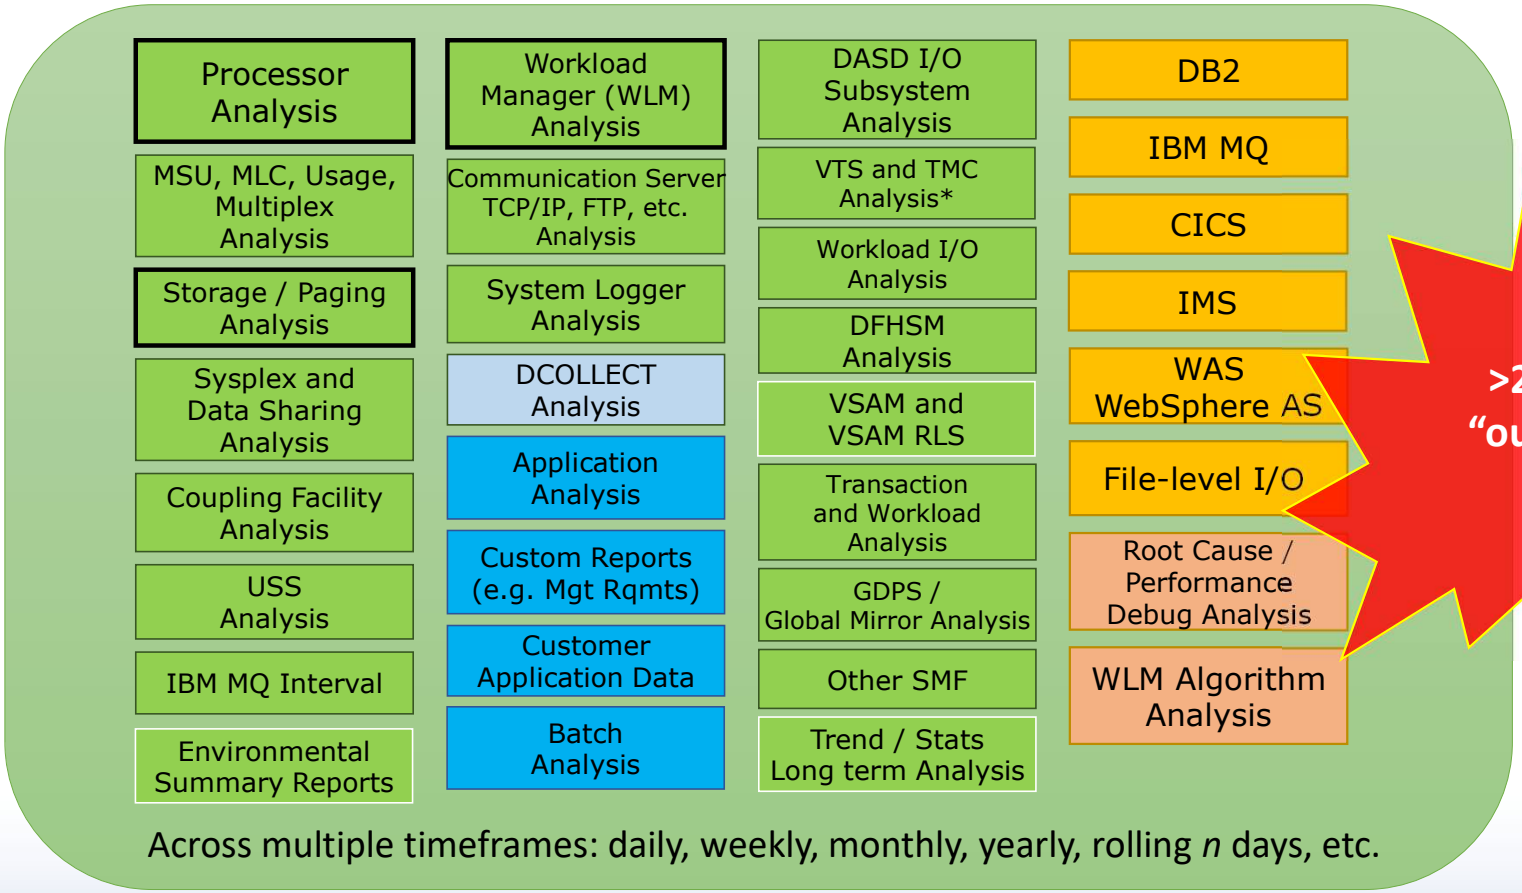

Pivotor Software as a Solution (SaaS)

- 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

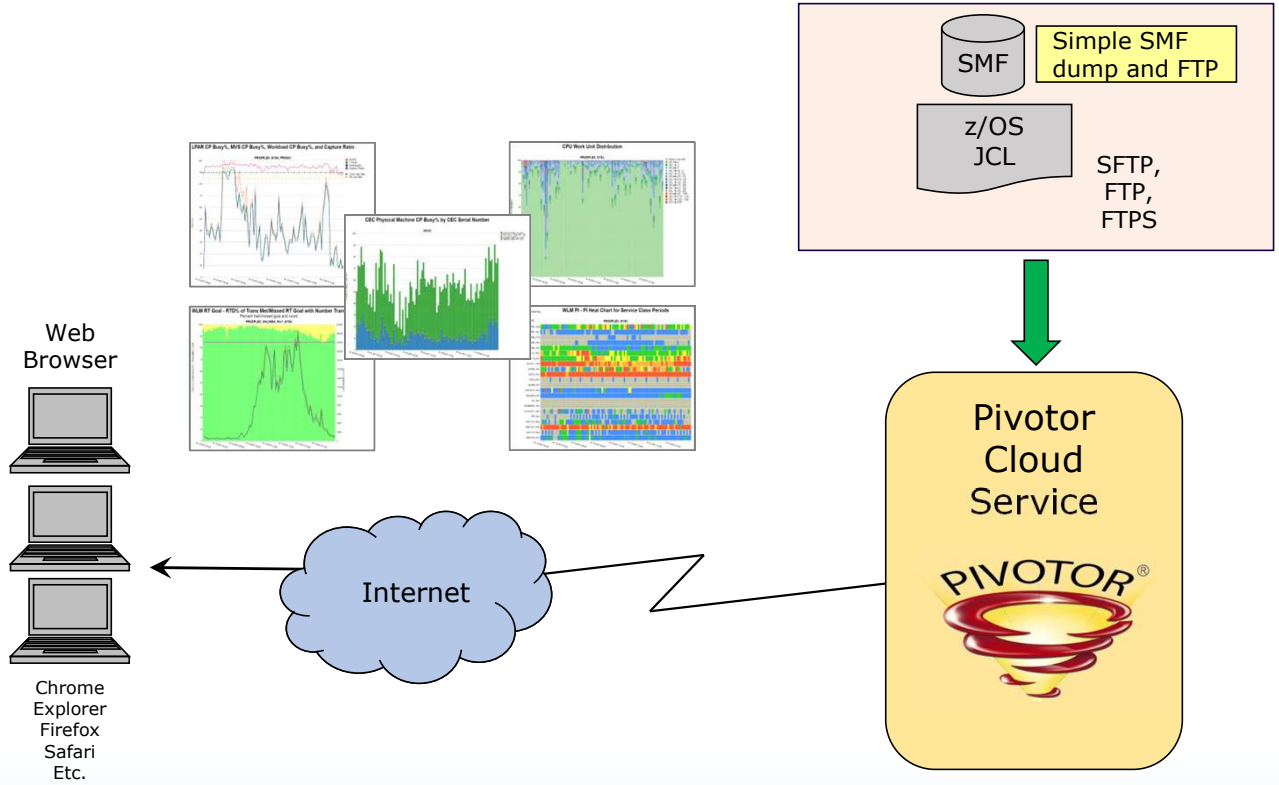

- Pivotor pricing is clear and affordable

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!

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

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

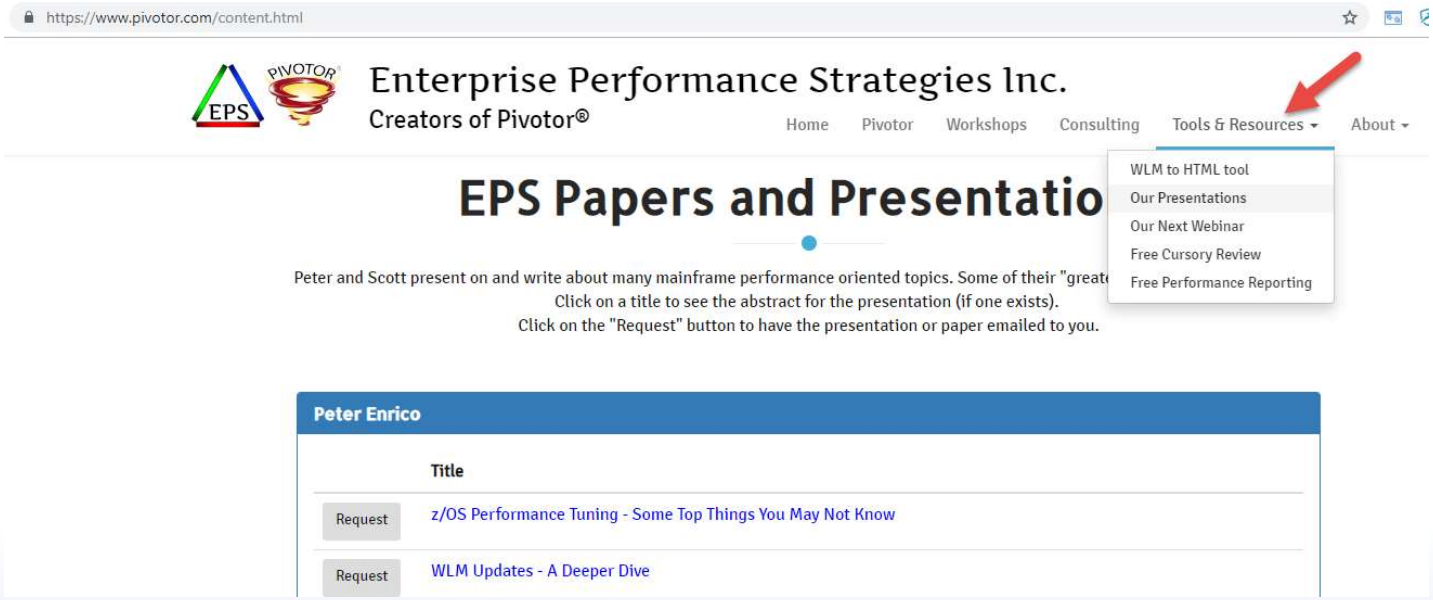

EPS presentations this week

What	Who	When	Where
Evolution of z/OS Memory Management: Large Memory, Large Pages, and How to Use Them	Scott Chapman	Mon 1:15	Hanover D
Back to Basics – Introduction to Parallel Sysplex and Data Sharing	Peter Enrico	Wed 9:15	Hanover D
z/OS WLM - Revisiting Goals Over Time	Peter Enrico	Wed 10:30	Hanover D
PSP: z/OS Performance Tuning – Some Top Things You May Not Know	Peter Enrico Scott Chapman	Wed 1:15	Hanover D

Must use Velocity Goals when using DB2 High Performance DBATs

DBAT

- Database Access Thread
- How DDF traffic connects to DB2
  - Vs. CICS/Batch/etc. which connect directly from those address spaces
- Consists of the thread (large) and connection (tiny)
- As a thread is used, it tends to grow and become “fat”
  - So can’t keep them around forever
- Threads have to have a place to run, for DDF that’s an enclave
  - DB2 creates the enclave to run the DBAT in

What is a transaction?

- From DB2 perspective:
  - Transactions start at first SQL read/write
  - Transactions end at COMMIT, ROLLBACK, or end of application/connection
- From WLM perspective, a unit of work:
  - Batch: life of job
  - STC: life of started task
  - DDF: life of the *enclave*

Traditional / Typical DB2 DDF Enclave

- DDF transactions are managed by WLM as independent enclaves
  - The life of the enclave is the life of the transaction

- In this case, 200 transactions come in, and for each transaction an enclave is created
- When each transaction ends, the enclave for that transaction is deleted, and an ended transaction and response time are posted to the WLM Service Class period
- WLM response time goals are best since each transaction is measured and considered by WLM

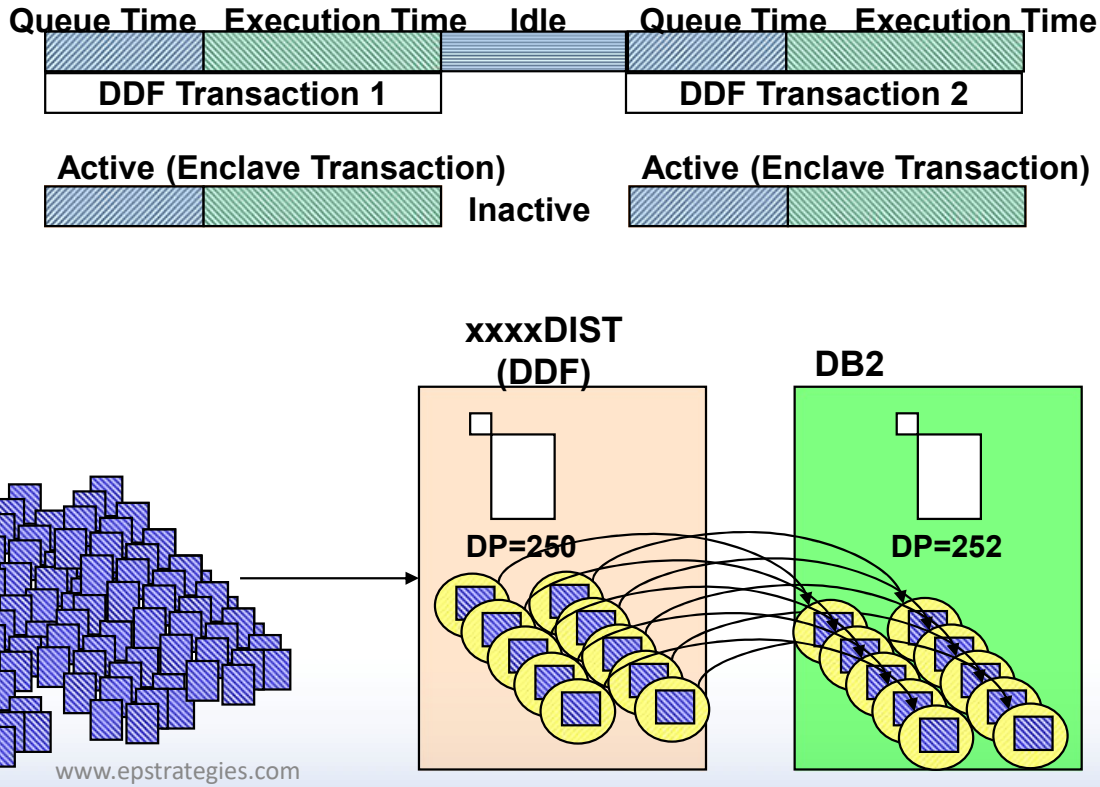

High-Perf DBATs w/ PH34378

- The DDF transactions still run in independent enclaves
  - But now the life of the enclave is for as long as up to 200* transactions in the enclave
  - Thus, you may have 200 ended transactions, but from a WLM point-of-view, the 200 transactions are part of a longer running enclave

* 500 in DB2 v13

- In this case, up to 200 transactions come in, and all are associated with a single enclave.
- After the last of the up to 200 transactions end, enclave is deleted, and one ended transaction is posted to the WLM service class period. Response time is from start of first transaction to end of last ended transaction
- Must use a velocity goal to effectively manage the work

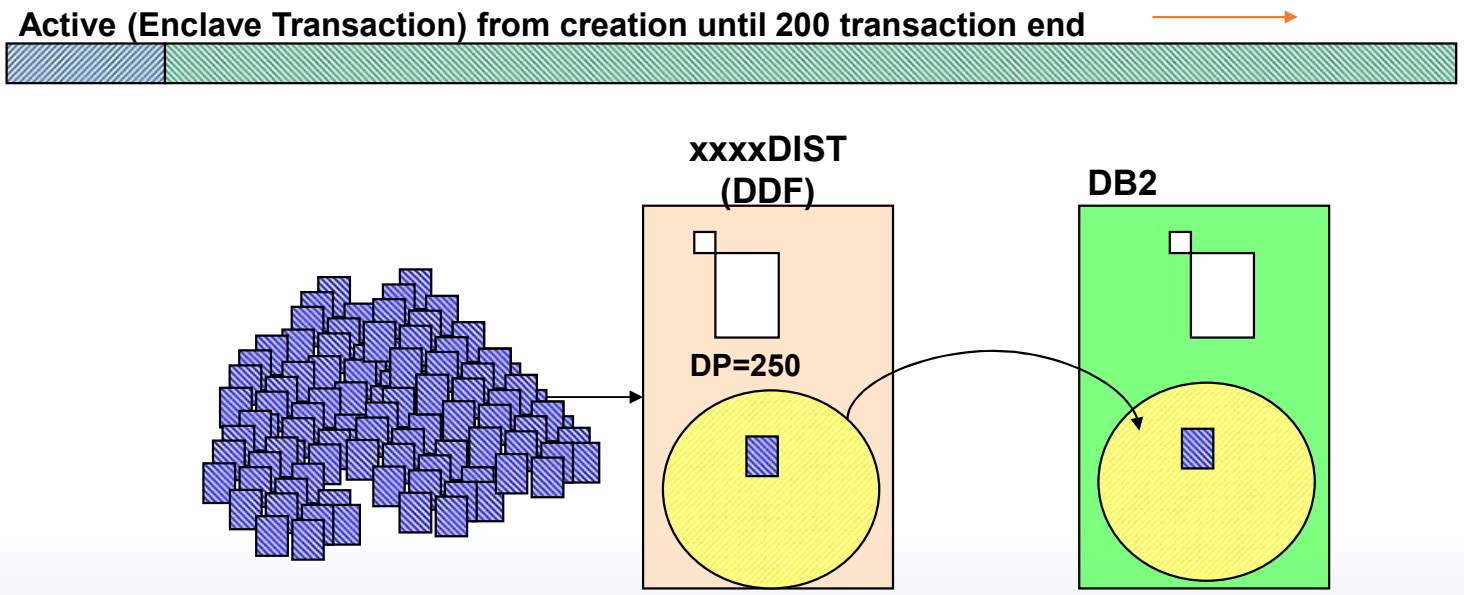

So why the change and why do we care?

- Apparently in some specific high-volume instances, retaining the enclave and not allocating a new one for each transaction provides performance benefit
  - Seems like a small benefit, likely only visible in very high-volume situations
- Would have been nice if DB2 would have added a flag/option for this behavior but instead they just changed it for everybody!
- So now WLM administrators need to know what DDF work is potentially using high-perf DBATs and change those to single-period velocity goals
  - Which limits the flexibility of managing that DDF work
  - If you fail to do this for your high-perf DBATs, that work will not be managed the way you expect
    - This is a risk point for the future unless everybody involved remembers this!

Help to report on actual transactions

- PH41024 / UI77711 Db2 support for WLM OA61811
- In combination, adds fields to smf072 to report on actual DDF transactions from the DB2 perspective instead of number of enclaves
- But does not change how the work can be managed!

Recommendations for DB2 DBAT Workloads

- Find out from your DB2 group if they are using High Performance DBATs
- If so, consider segregating these DBAT trans to separate WLM service classes
  - Assign those service classes single-period velocity goals
- If your DDF work is not using single-period service class with a velocity goal, and it starts using high-perf DBATs (after PH34378) your DDF work will no longer be managed as you expect
  - Seems like a risk sitting out there waiting to bite people

SRB – System Recovery Boost

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
- Requires z15 or higher machine
- Requires z16 or higher machine for middleware boost
  - Trigger a boost on starting site-specified STCs
- No additional charge for basic System Recovery Boost

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
  - 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 from storage controller failover
  - Note: z/OS 2.5 added additional boosts such as Middleware boost

LPARs being boosted

- Seems obvious: they get more work done!
- But may need changes to prepare them to get more work done
  - Maybe: define reserved (offline) zIIPs to be brought online during boost period
  - May need to add zIIP weight to be able to access newly online zIIPs
- Increased zIIP usage could cause some zIIP work to cross over to the GPs
  - Work with HONORPRIORITY=NO prevented from crossing over
  - So *may* want to consider changing *from* HONORPRIORITY=NO during IPL boost
    - Probably unnecessary in most cases
    - Recovery process boosts likely short enough to not be a major concern
    - Work should be routed away from the LPAR during shutdown boost anyways
- Boosted LPARs over-achieving work won't be capped to help discretionary
- New SMF interval started when boost starts/ends

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
- 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

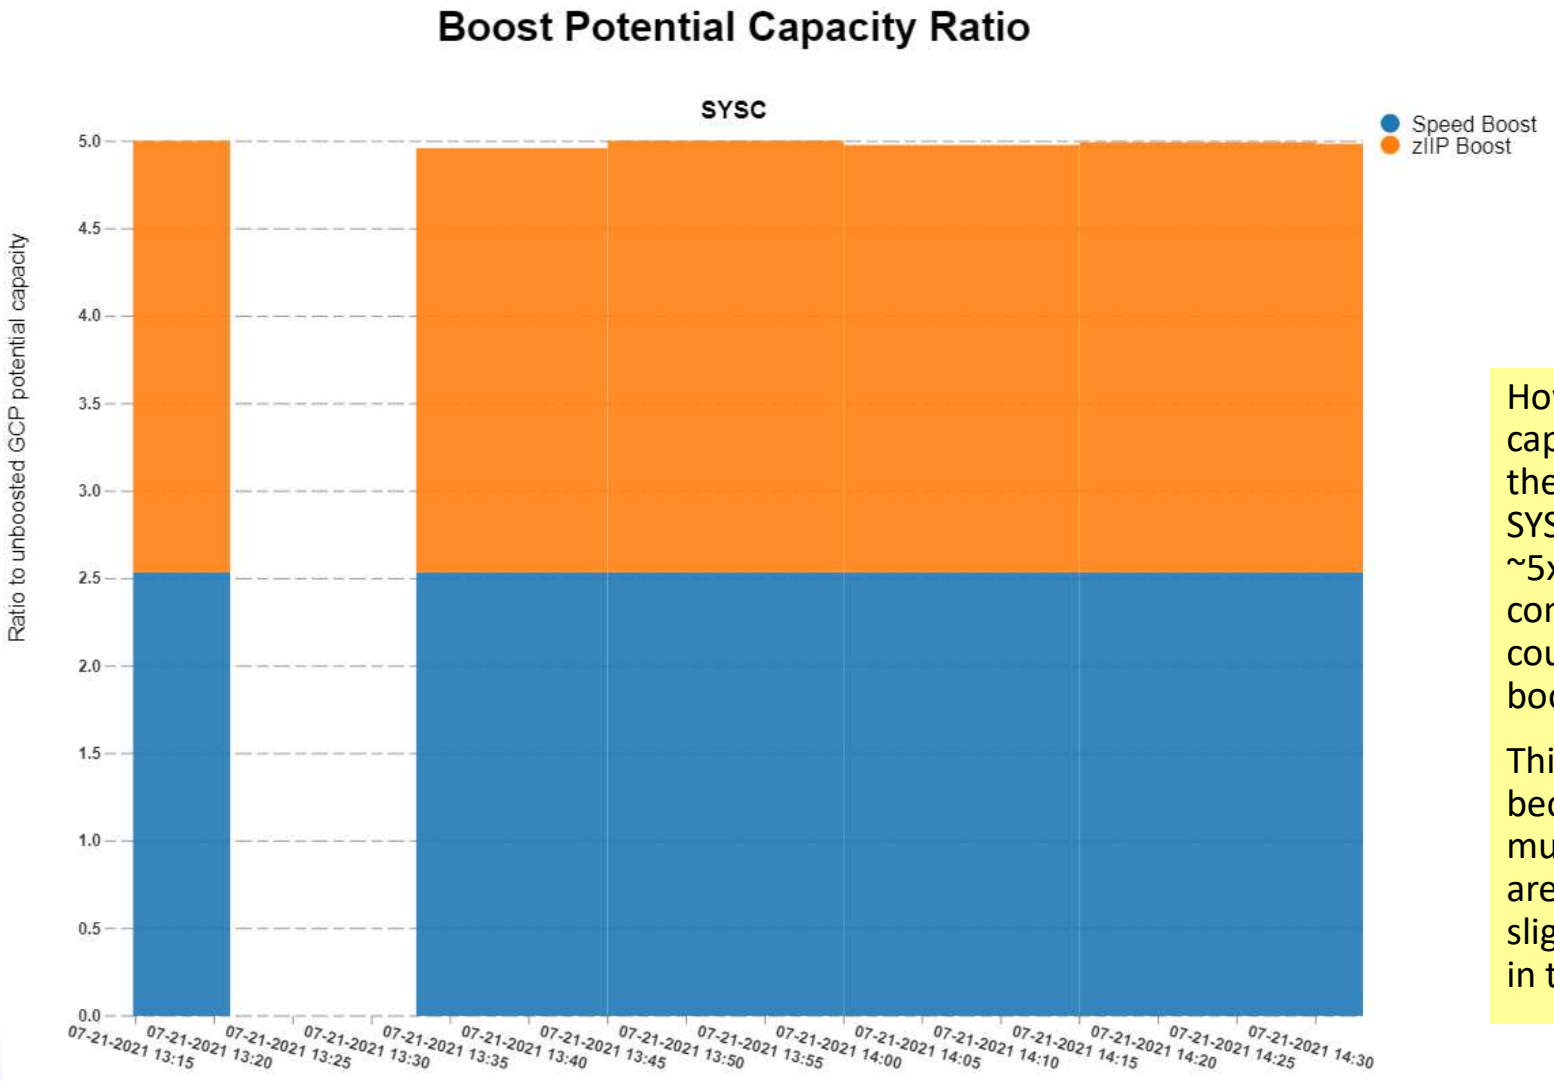

How much extra capacity did we get from the boost? Potentially SYSC could have done ~5x the GP work compared to what it could have done un-boosted!

This report was tricky because it considers how much the other LPARs are using hence the slight variations you see in the ratios.

Things that limit SRB's effectiveness

- Full speed GCPs eliminates speed boost
- Busy zIIPs
  - Less available capacity to the boosted LPAR
- Single physical zIIP shared among several LPARs
  - PR/SM will still move the zIIP between LPARs limiting the time the boosted LPAR might be able to use it
- Few or no zIIPs available to the boosted LPAR

Scott's Opinion: In the 2020s, almost every machine should have at least 2 zIIPs!

Some interesting SRB considerations

- If you don't have your SMF intervals synced, you won't get new SMF 30 interval records when boost changes
  - Meaning you'll have records that cross boost/unboost time, making those records particularly problematic for performance analysis
- If boost makes the zIIPs busier, zIIP-eligible work may start crossing over
  - Not a big deal in the boosted LPAR, but could impact non-boosted LPARs, especially if those non-boosted LPARs had been "borrowing" zIIP weight
- Importance 1 & 2 work treated as CPU critical during boost periods
  - This probably isn't a big deal in most cases, but we're struggling a bit to come up with a good justification for why they did it
  - If you have importance inversions at non-boosted time this could (maybe) impact those workloads during boost

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

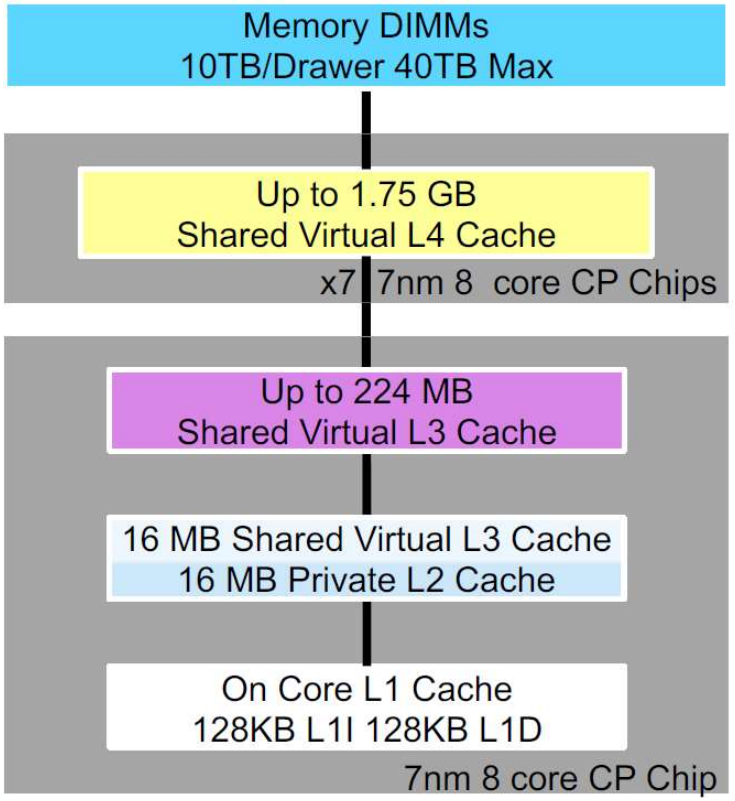

z16 – Virtual L3 and L4 Caches (slide source: IBM)

- Virtual Cache Layers built from L2 Cache blocks
  - Mirrors physical hierarchy of prior designs
- Base Virtual Cache design behavior
  - Private L2 Cache when processor is active (dynamic, 16MB)
  - Virtual L3 Cache when processor is inactive (victim)
  - Virtual L4 Cache when processor chip is inactive (victim)
- Scales as additional processors are brought online
  - L2 Caches switch from Victim L3 to Private L2 behavior
  - Similar to Cache Inclusivity tax of prior designs
- Logical Hierarchy remains
  - 1.5x more cache per core at vL3, vL4
  - More efficient use of cache array space
  - Overcomes limits of traditional architecture
  - Extendable for future generations

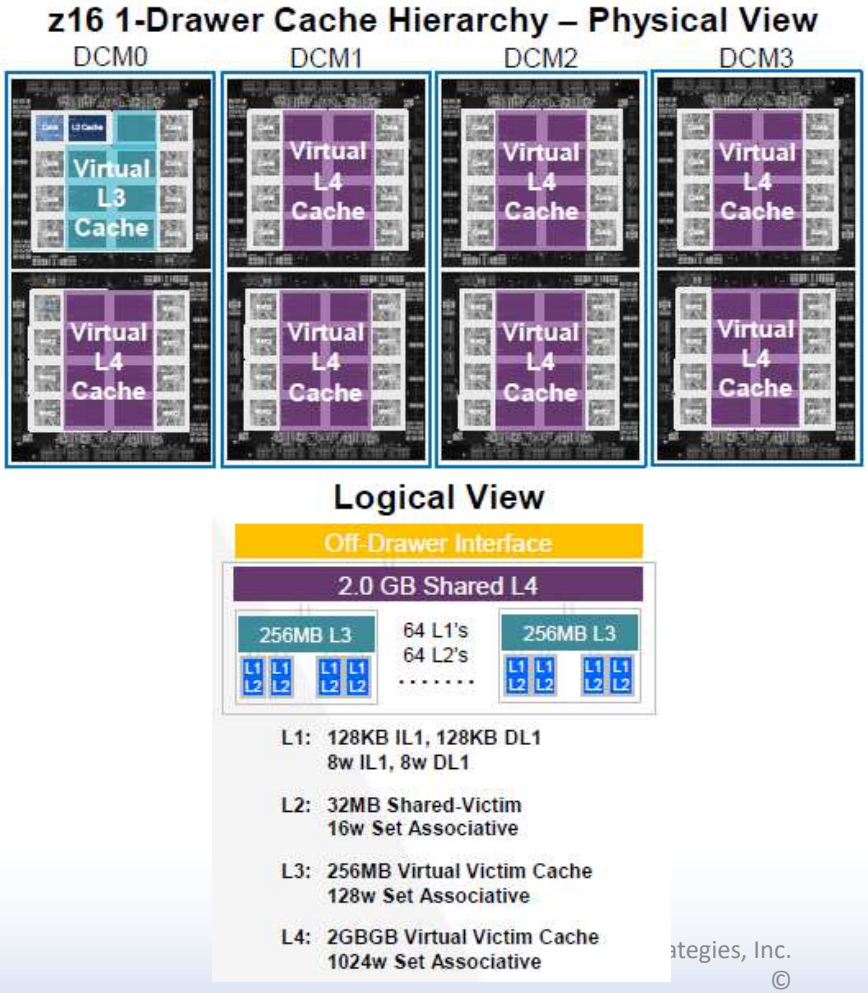

z16 Virtual Cache Provisioning

- One chip example (just to make the point)

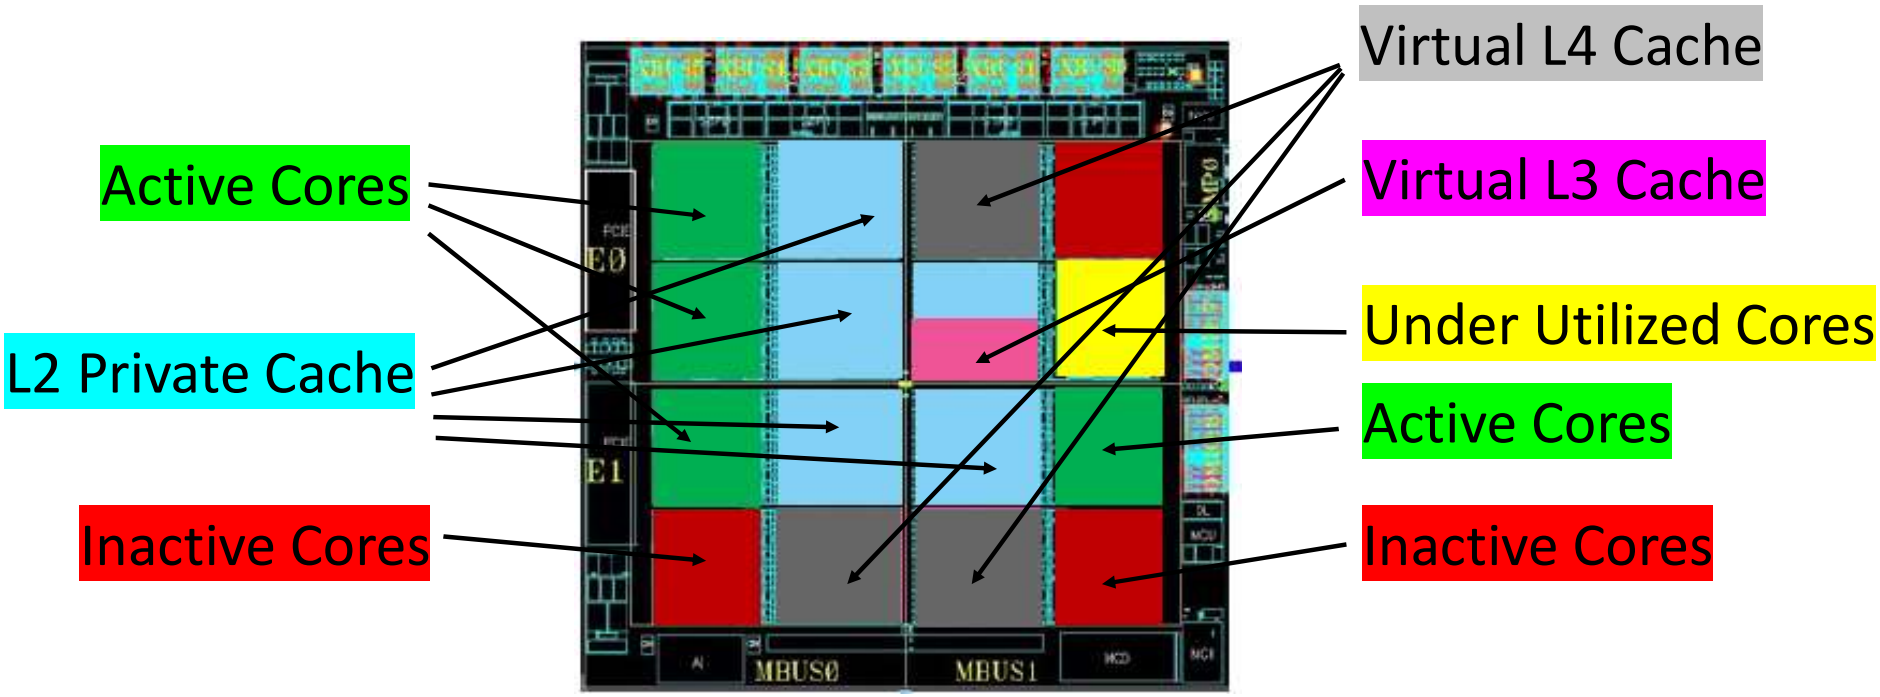

Cache Demo

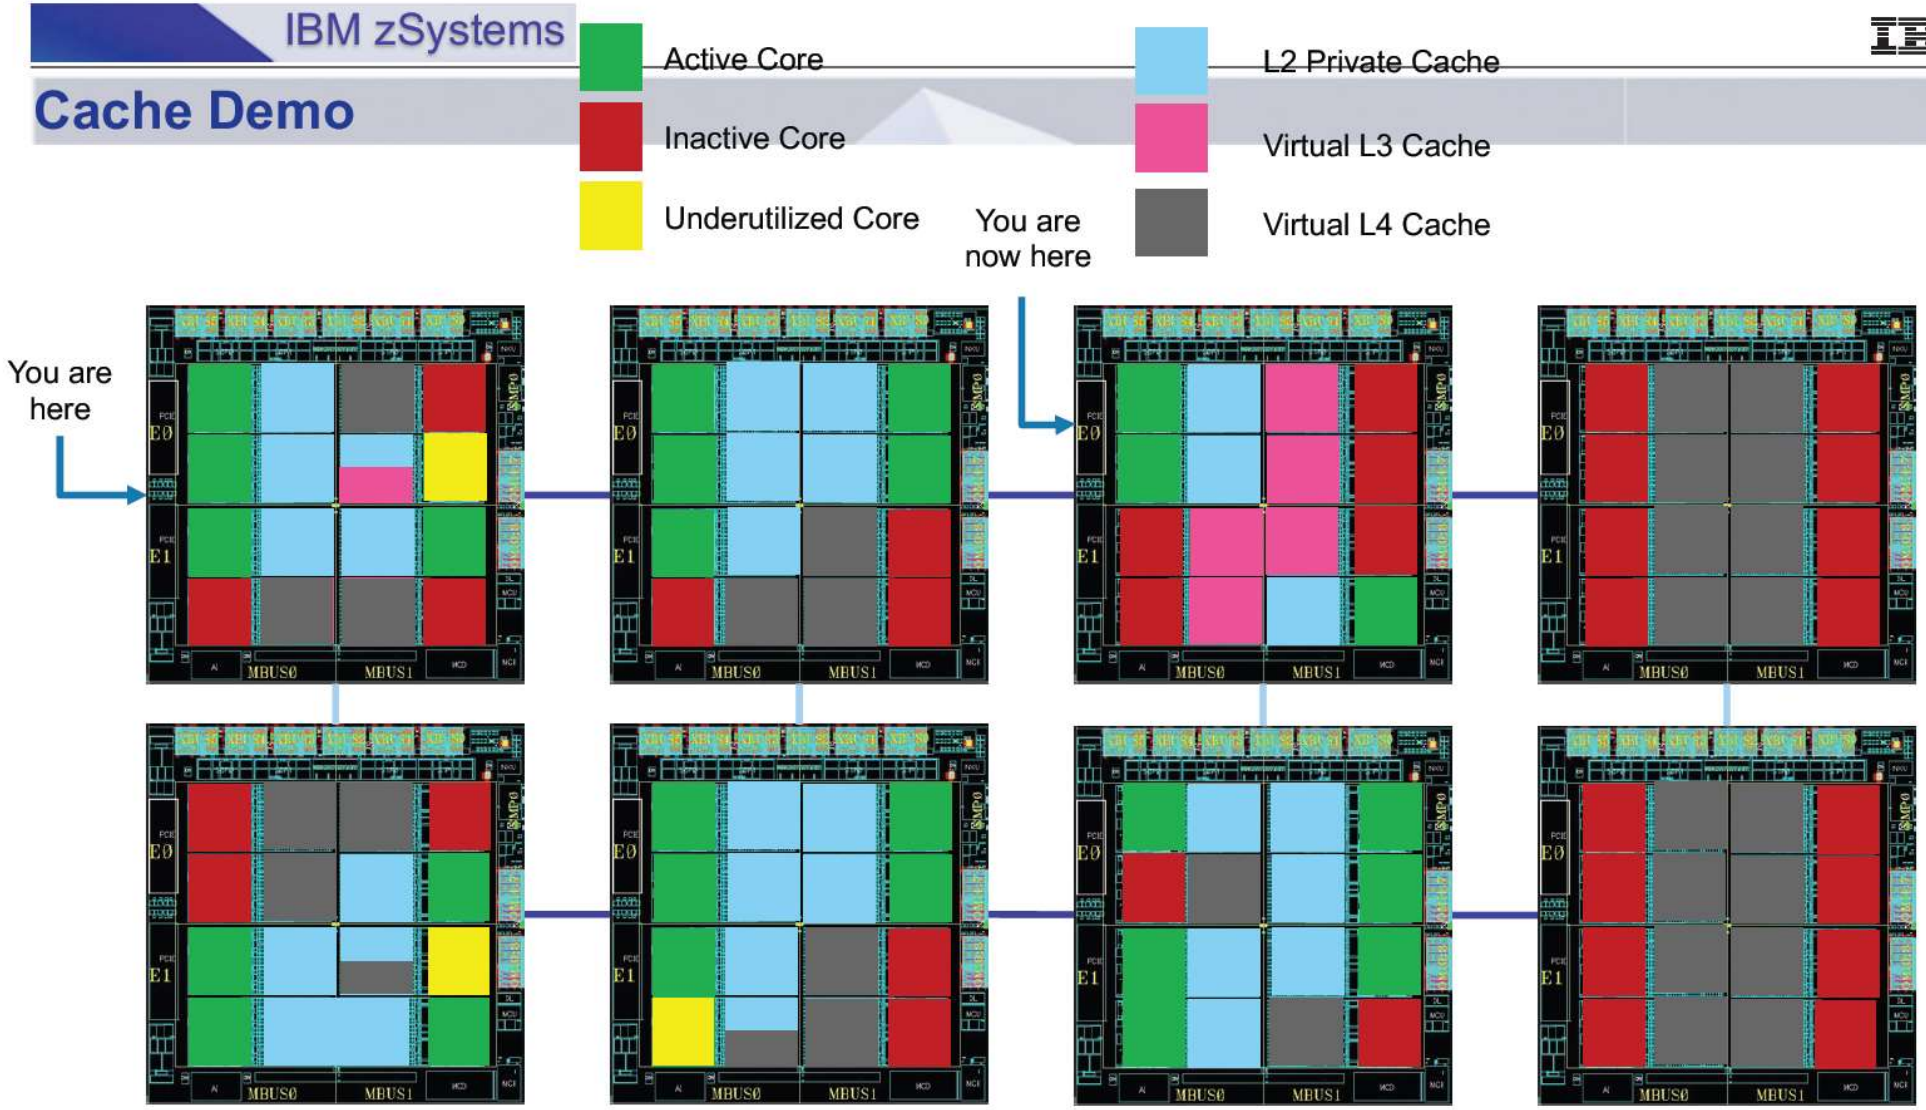

Case Study CEC LSPRs: z14 vs z16

- 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

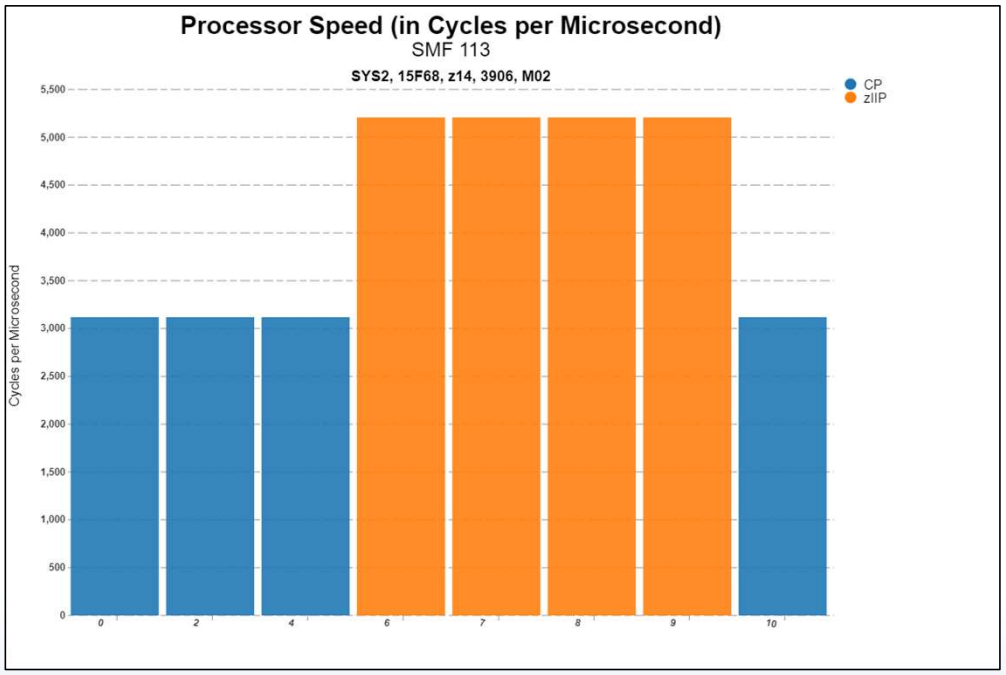

- z16 (3931-606 A01)
  - 6CPs, 4 zIIPs
  - SYS2: 4 CPs, 4 zIIPs

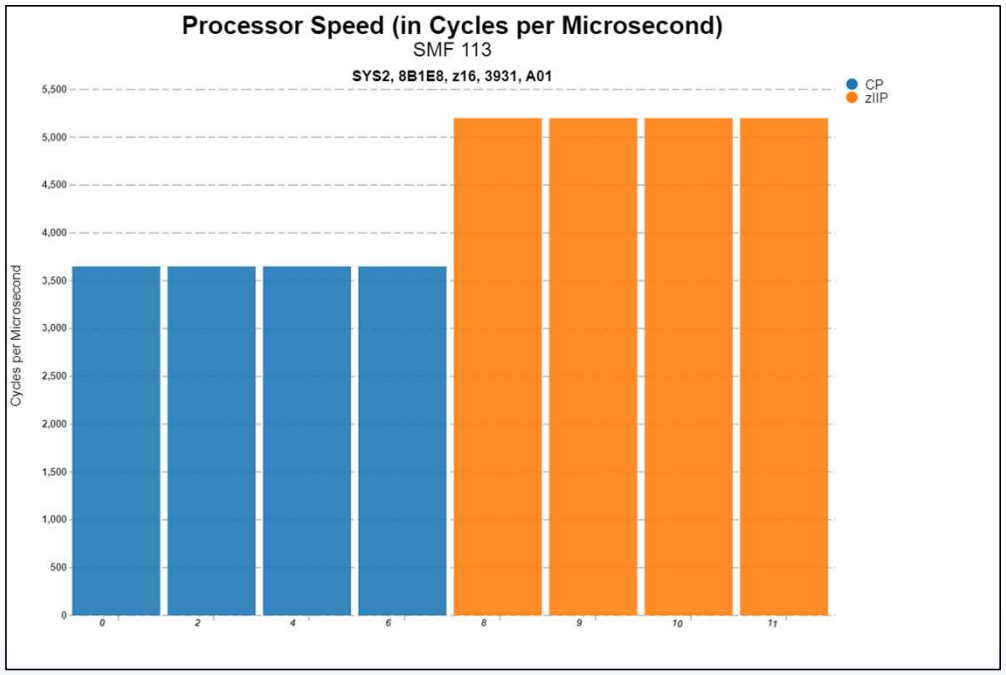

z14 vs z16 – Cache Sourcing

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

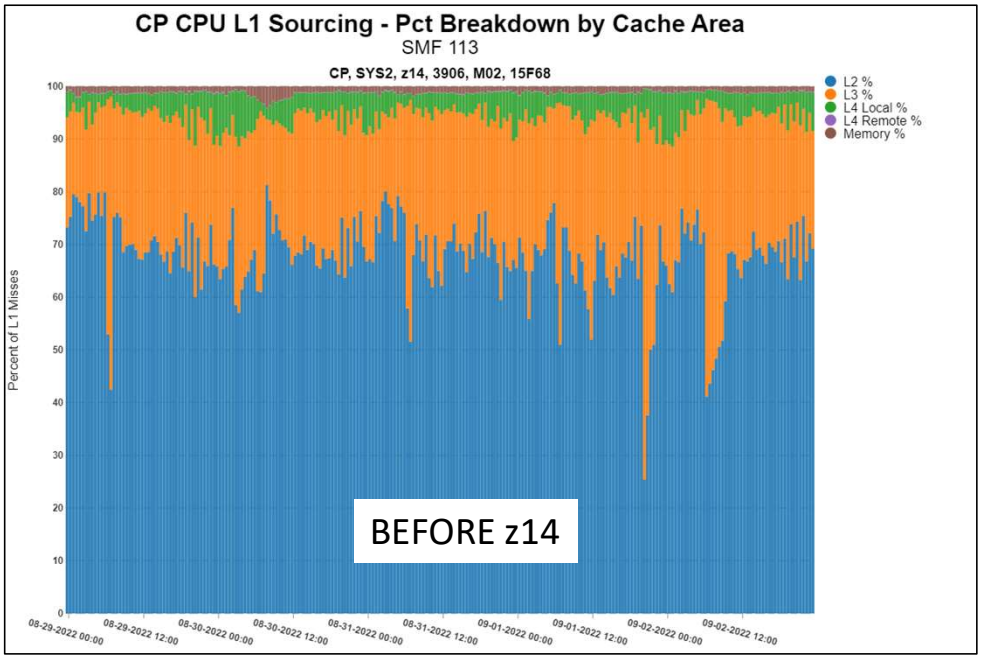

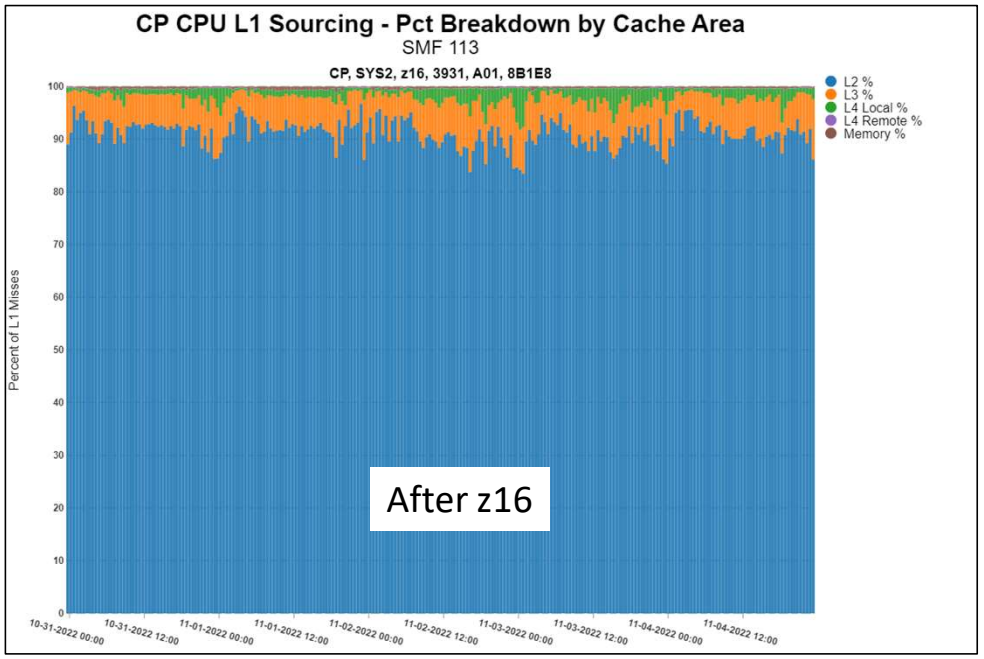

RNI – Breakdown by Cache

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.

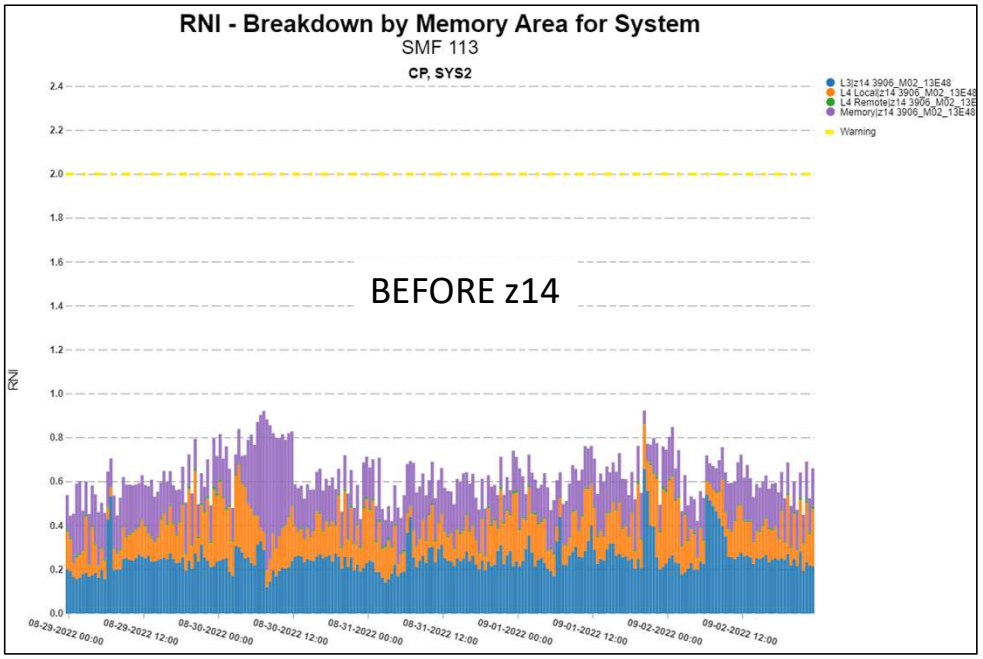

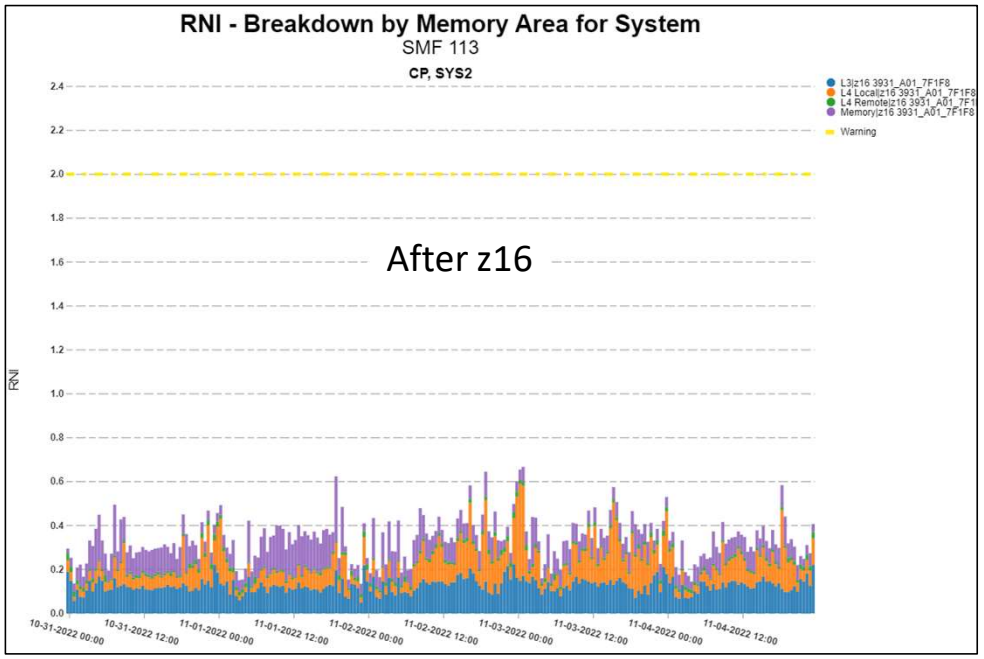

Service Definition Coefficient Changes (Migrating from CPU=1,SRB=1,IOC=xx,MSO=0)

Multiple Periods and Durations

- As transactions are processed, they consume system resources
  - CPU and SRB service units (i.e. processor)
  - I/O service units
  - MSO service units (i.e. storage)
- Traditionally, durations are in terms of 'service' and not 'service units'
  - Service is service units weighted by service definition coefficients (SDCs)
  - When duration is set for a period, the service consumed determines period switch
  - Why weight? Historical…

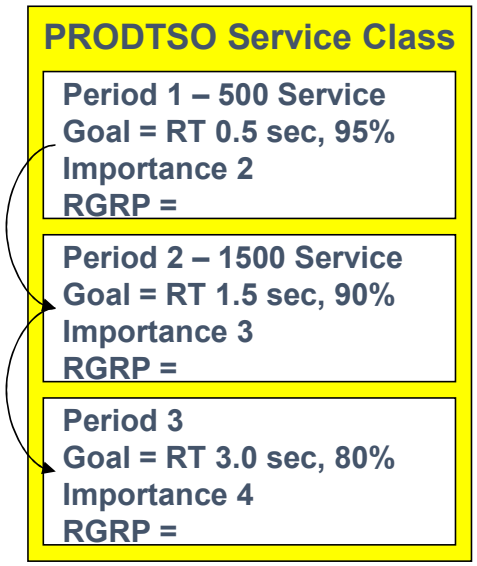

$$\text{Service} = \begin{pmatrix} (\text{CPU SDC} * \text{CPU Service Units}) \\ + (\text{SRB SDC} * \text{SRB Service Units}) \\ + (\text{IOC SDC} * \text{IOC Service Units}) \\ + (\text{MSO SDC} * \text{MSO Service Units}) \end{pmatrix}$$

Recommendations for WLM Service Definition Coefficients

- 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'.

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.

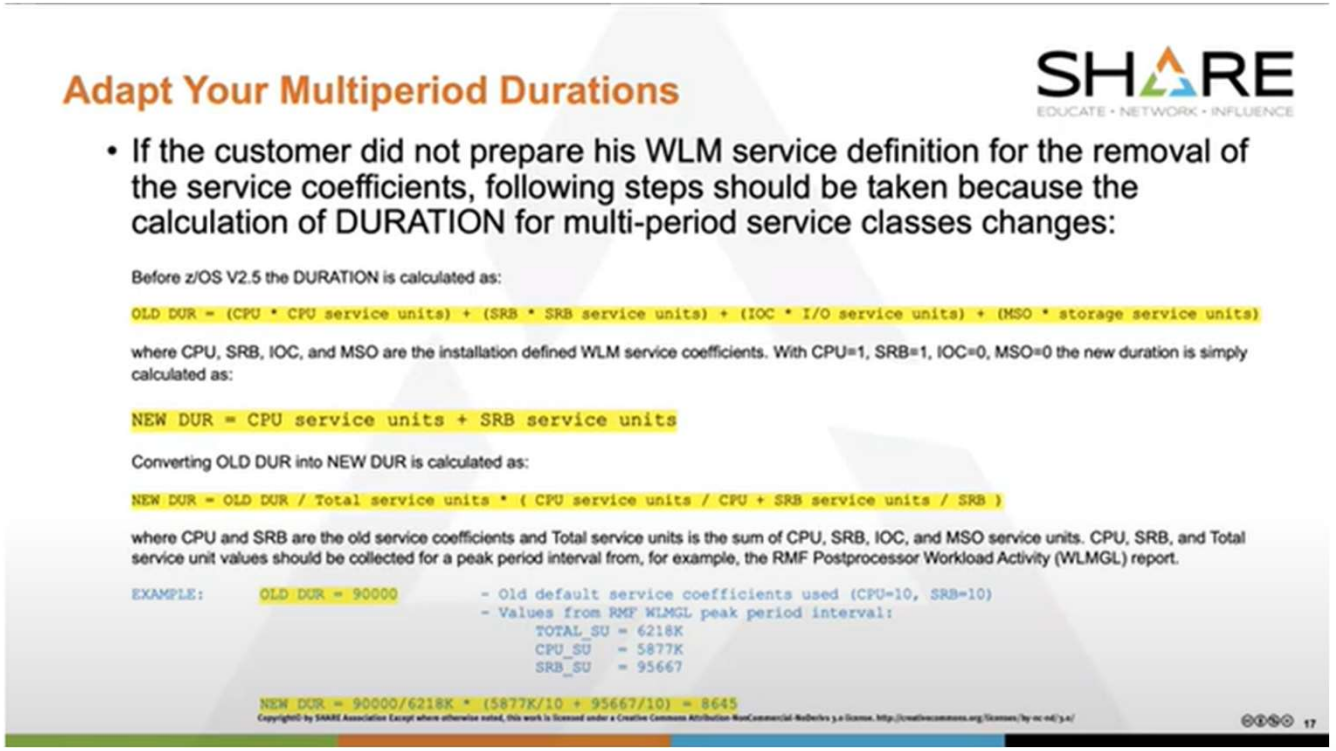

Adapt Your Multiperiod Durations

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

Steps 1 and 2

1. Determine your current Service Definition Coefficients
2. Determine your current multiple period service classes

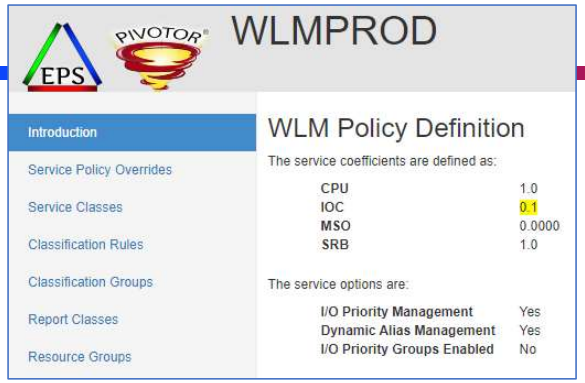

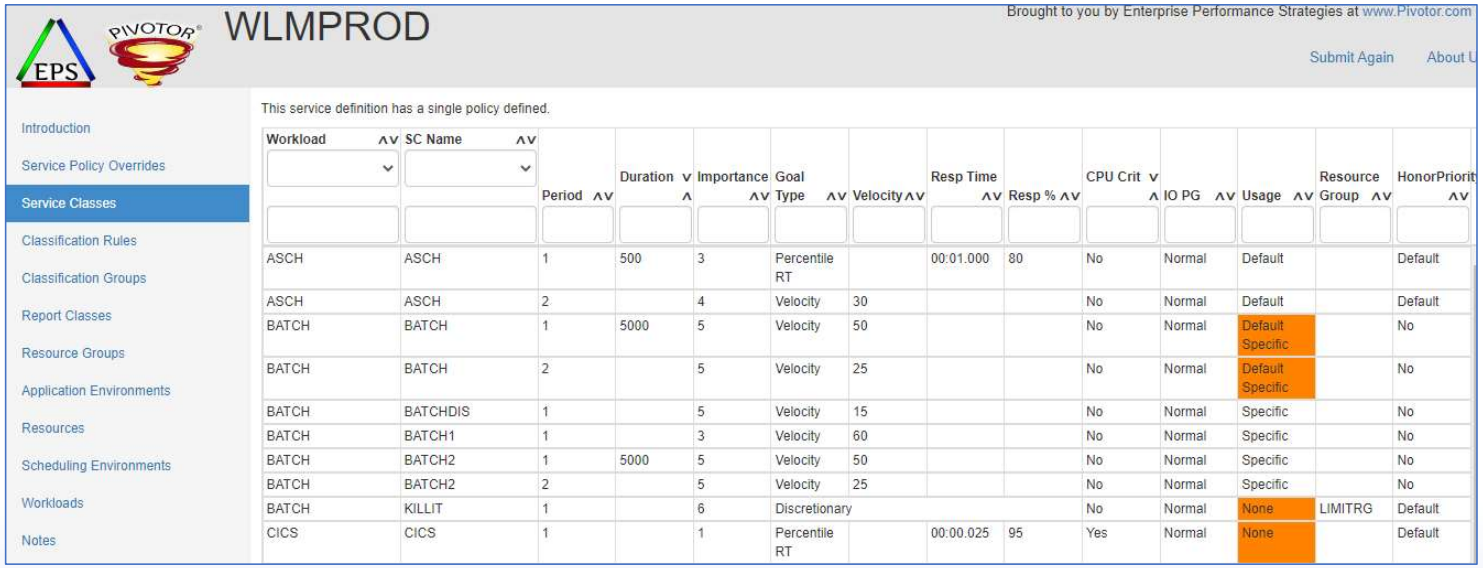

Workload	SC Name	Period	Duration	Importance	Goal Type	Velocity	Resp Time	Resp %	CPU Crit	IO PG	Usage	Resource Group	HonorPriorit
ASCH	ASCH	1	500	3	Percentile RT		00:01.000	80	No	Normal	Default		Default
ASCH	ASCH	2		4	Velocity	30			No	Normal	Default		Default
BATCH	BATCH	1	5000	5	Velocity	50			No	Normal	Default Specific		No
BATCH	BATCH	2		5	Velocity	25			No	Normal	Default Specific		No
BATCH	BATCHDIS	1		5	Velocity	15			No	Normal	Specific		No
BATCH	BATCH1	1		3	Velocity	60			No	Normal	Specific		No
BATCH	BATCH2	1	5000	5	Velocity	50			No	Normal	Specific		No
BATCH	BATCH2	2		5	Velocity	25			No	Normal	Specific		No
BATCH	KILLIT	1		6	Discretionary				No	Normal	None	LIMITRG	Default
CICS	CICS	1		1	Percentile RT		00:00.025	95	Yes	Normal	None		Default

Determine with multiple period service classes are consuming I/O service

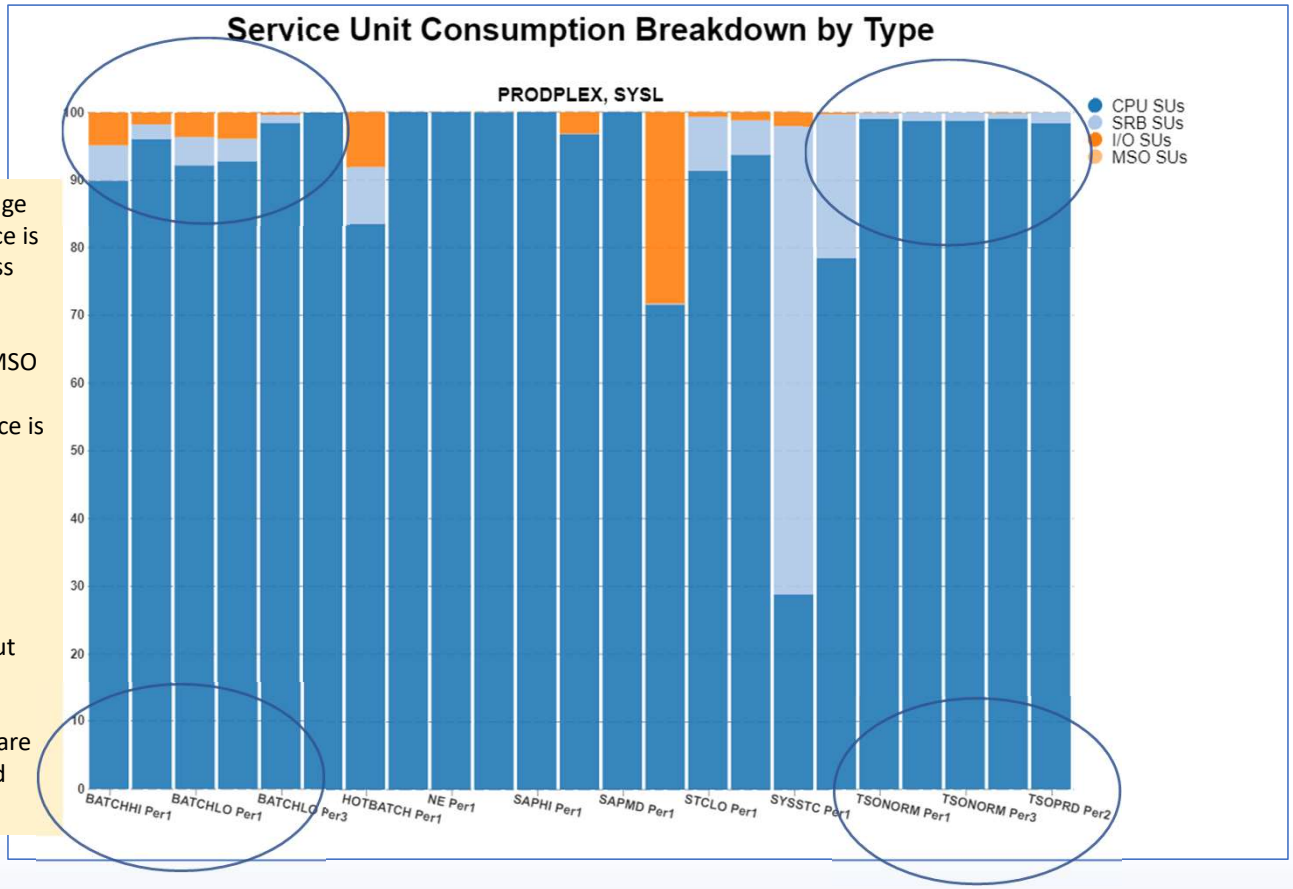

This report shows the percentage breakdown of how much service is consumed by every service class period with activity.

The orange is I/O service and MSO service.

- Chances are your MSO service is zero since MSO=0 has been recommended for 20+ years
- I/O service is what we care about
- Even then we only care about the TSO, DDF, OMVS, Batch periods.
  - Even then, we only care all but the last period

Determine with multiple period service classes are consuming I/O service

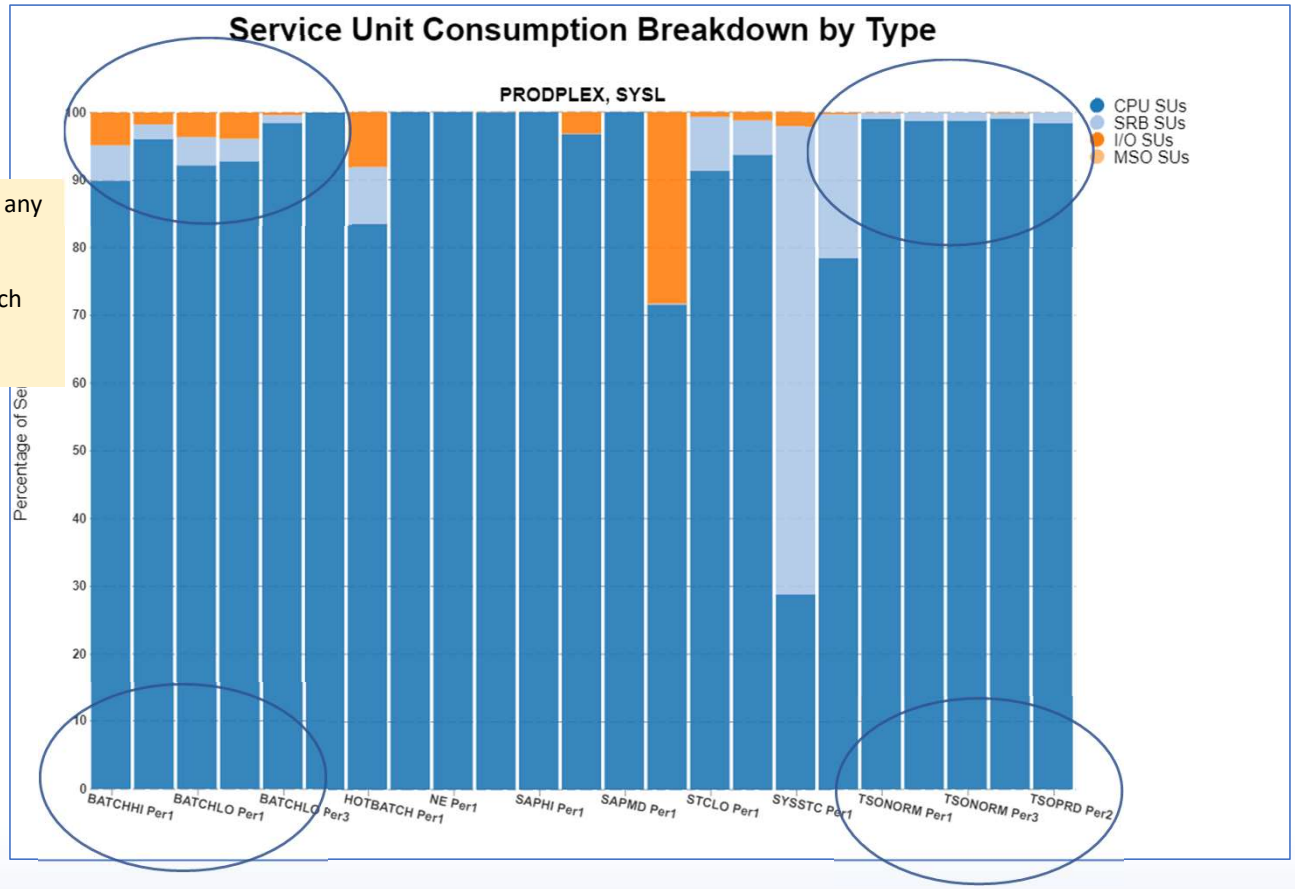

In this example, I would ignore any migration exercise.

Not enough I/O to warrant much of a migration.

Beware High Velocities
Turn off WLM option I/O priority management

WLM Velocity Goal - Achieved Velocity Across Sysplex

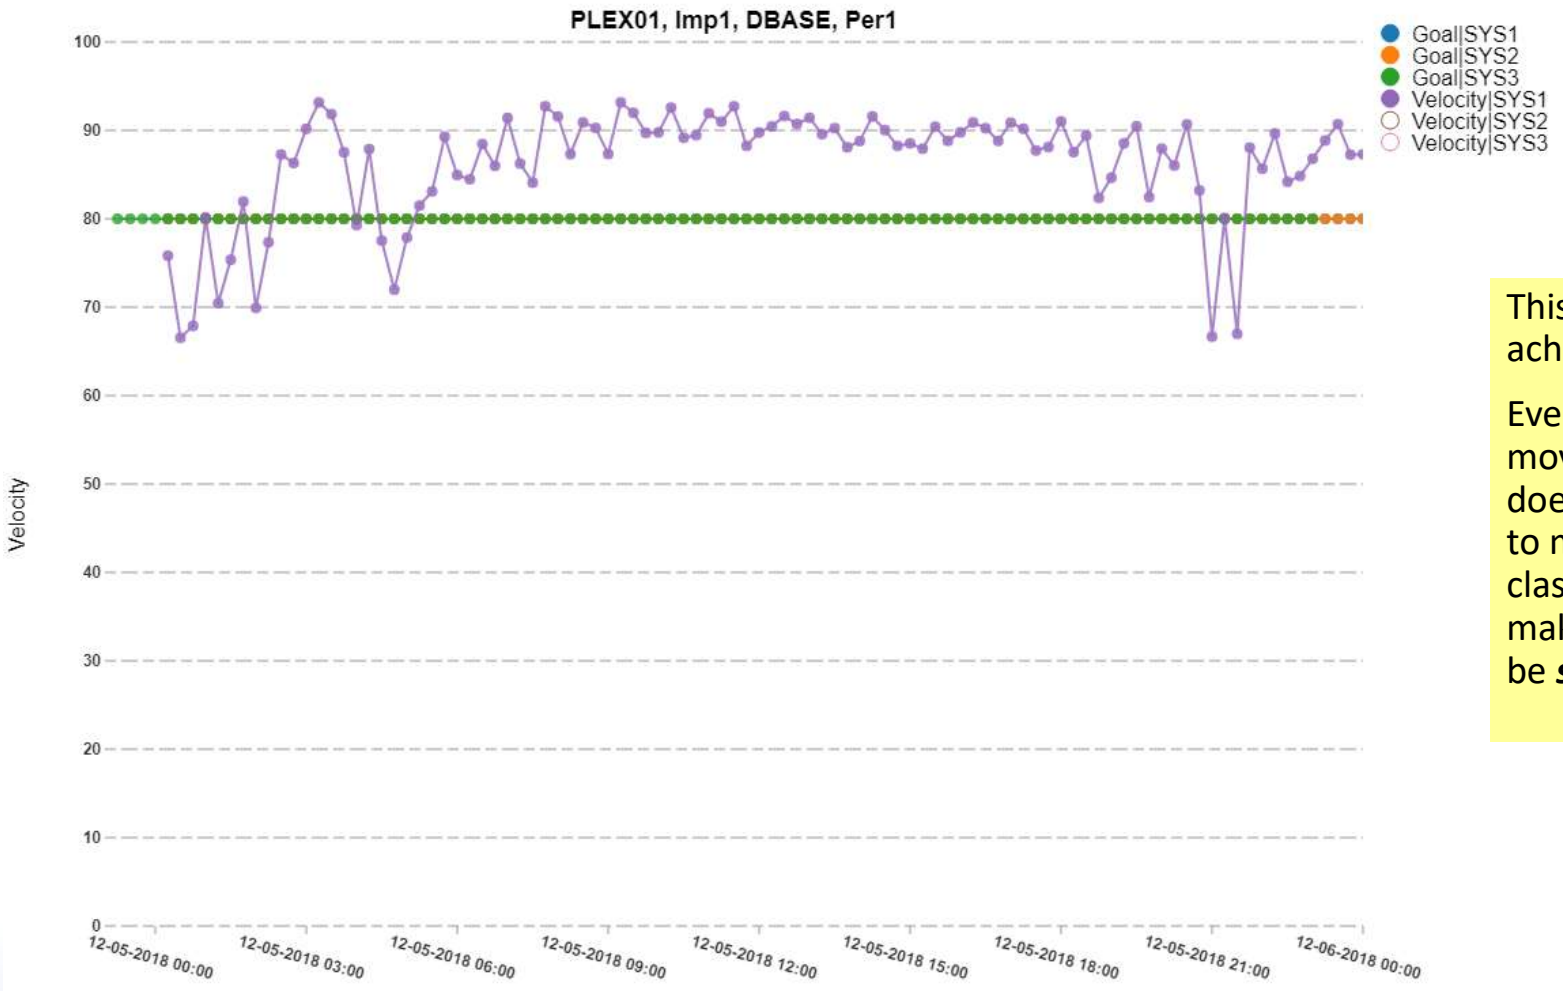

This SCP is easily over-achieving its goal.

Even if the goal was moved to 90 or 95, WLM doesn't really have room to manage this service class because it looks to make changes that will be ***significant***.

Velocity Formulas and I/O Samples

- WLM I/O priority management option 'Yes' allows WLM to manage I/O dispatching priorities separate from CPU dispatching priorities
  - . When enabled, velocity formula includes I/O samples

$$\left(\frac{\text{CPU Using}}{\text{CPU Using} + \text{CPU Delay} + \text{Storage Delays}}\right) \text{ versus } \left(\frac{\text{CPU Using} + \text{I/O Using}}{\text{CPU Using} + \text{CPU Delay} + \text{I/O Using} + \text{I/O Delay} + \text{Storage Delays}}\right)$$

$$\left(\frac{257}{257 + 126 + 14}\right) = \text{vel } 65\% \qquad \left(\frac{257 + 289}{257 + 126 + 289 + 5 + 14}\right) = \text{vel } 79\%$$

CPU & I/O Contribution to Velocity

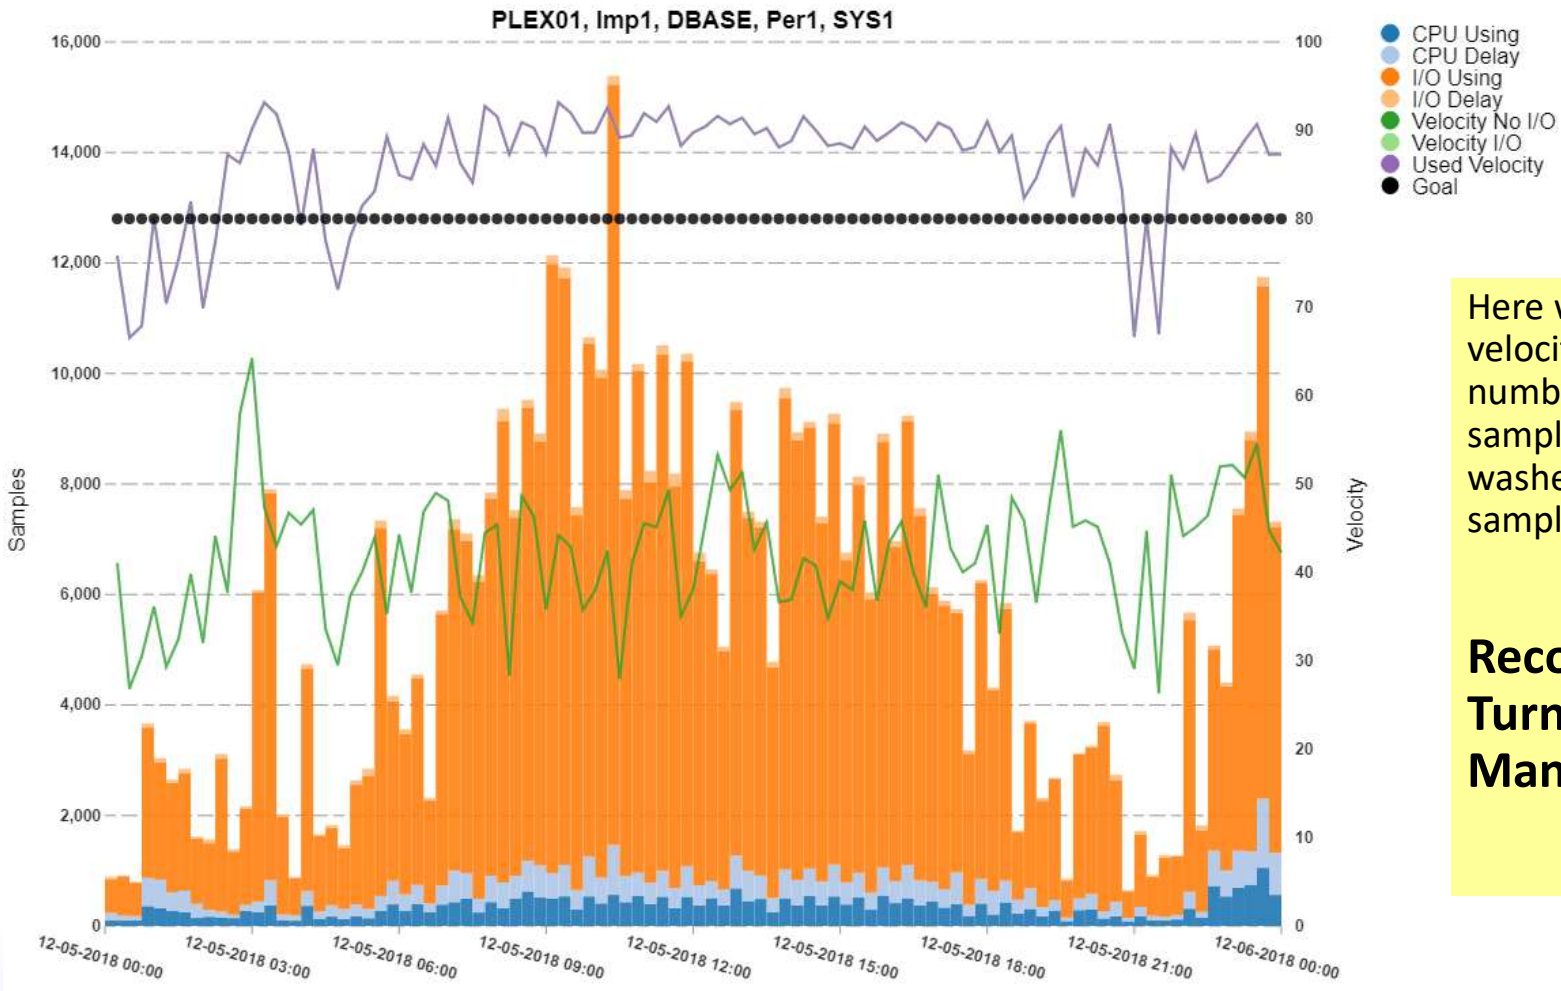

Here we see why the velocity is so high—the number of I/O using samples completely washes out the CPU samples.

Recommendation: Turn off I/O Priority Management

Enable SMF 98 and 99 Record Recording

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

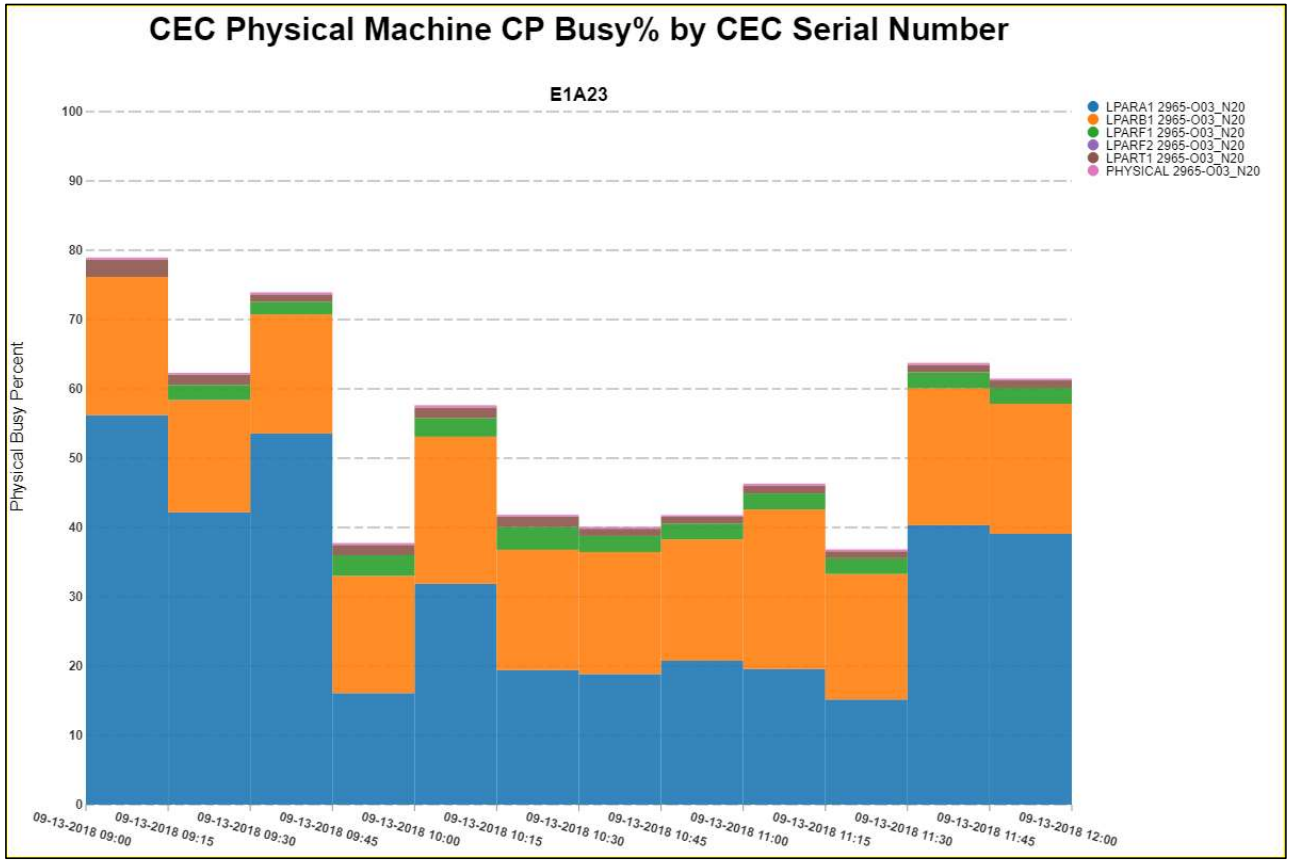

Problem Statement:

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

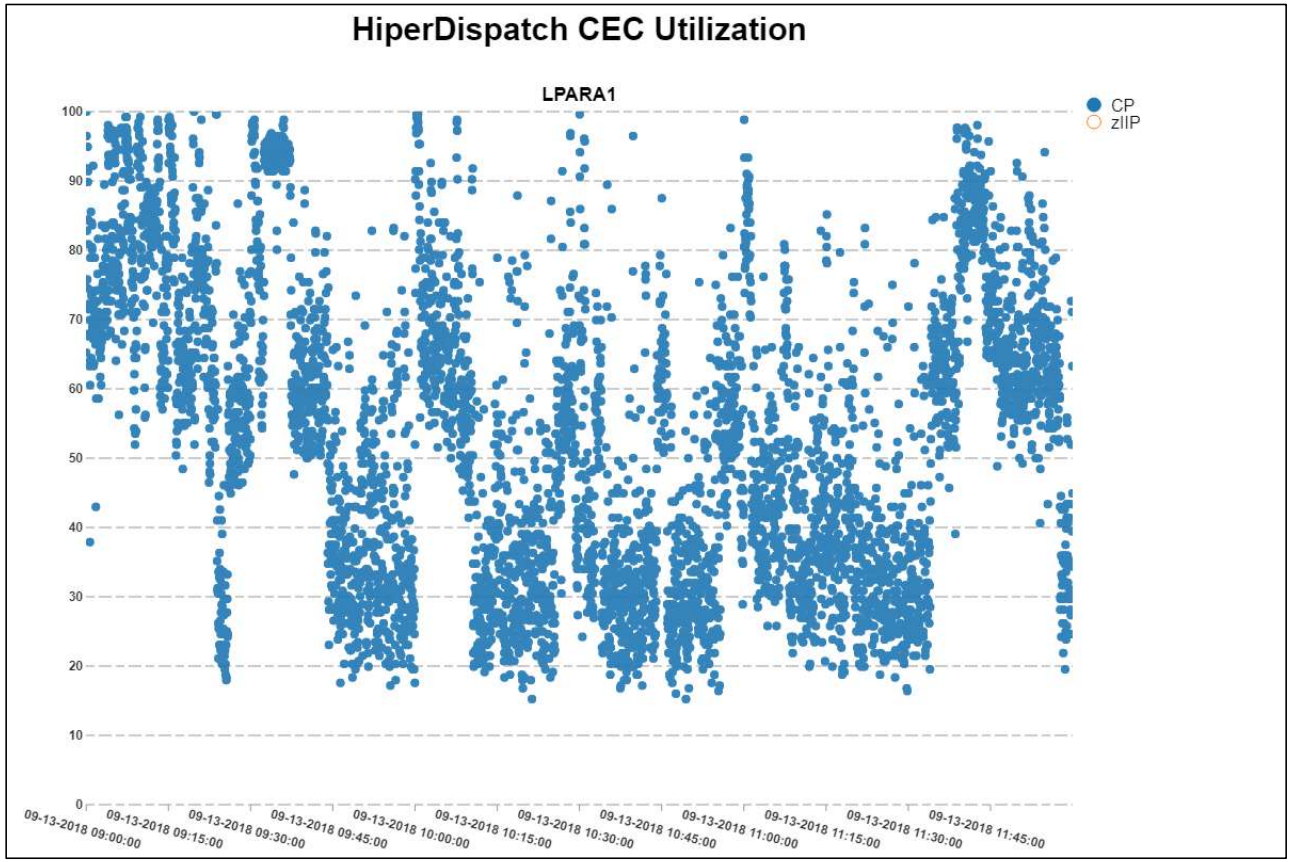

High Frequency CEC Utilization:

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.

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