SMF 99 WLM Decision-Making Traces

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

Creators of Pivotor®

Peter Enrico

Email: Peter.Enrico@EPStrategies.com

Enterprise Performance Strategies, Inc.
3457-53rd Avenue North, #145
Bradenton, FL 34210
http://www.epstrategies.com
http://www.pivotor.com

Voice: 813-435-2297
Mobile: 941-685-6789

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

SMF 99 WLM Decision-Making Traces

For the last 30 years, WLM has made trace data available to help us understand the decisions it makes when managing WLM service class periods when meeting their goals.

During this webinar, **Peter Enrico** will review the available WLM trace measurements, discuss common terms used by WLM when managing workloads, and walk us through some WLM traces to show how WLM decisions are made. This webinar will be interesting to those who want to have a better understanding of how WLM works.

Note from Peter: In past webinars and conference presentations I have presented how WLM works. While this presentation draws from that 'how WLM works' presentation, you can think of that presentation as a primer to understanding this Decision-Making Traces presentation.

EPS: We do z/OS performance…

- **Pivotor** – z/OS performance reporting and analysis software and services
  - Not just SMF reporting, but analysis-based reporting based on expertise
  - www.pivotor.com
- **Education and instruction**
  - We teach our z/OS performance workshops all over the world
  - Want a workshop in your area? Just contact me.
- **z/OS Performance War Rooms**
  - Intense, concentrated, and highly productive on-site performance group discussions, analysis and education
  - Amazing feedback from dozens of past clients
- **MSU Reduction Exercises**
  - The goal is to reduce the MSU consumption of your applications and environment
- **Information**
  - We present around the world and participate in online forums
  - https://www.pivotor.com/content.html
    https://www.pivotor.com/webinar.html

z/OS Performance workshops available

During these workshops you will be analyzing your own data!

- WLM Performance and Re-evaluating Goals
  - May 12 - 16, 2025 (4 days)
- Parallel Sysplex and z/OS Performance Tuning
  - October 21 - 22, 2025 (2 days)
- Essential z/OS Performance Tuning
  - September 22 - 26, 2025 (4 days)
- Also… please make sure you are signed up for our free monthly z/OS educational webinars! (email contact@epstrategies.com)

Pivotor's Comprehensive Report Sets for Immediate Performance Analysis

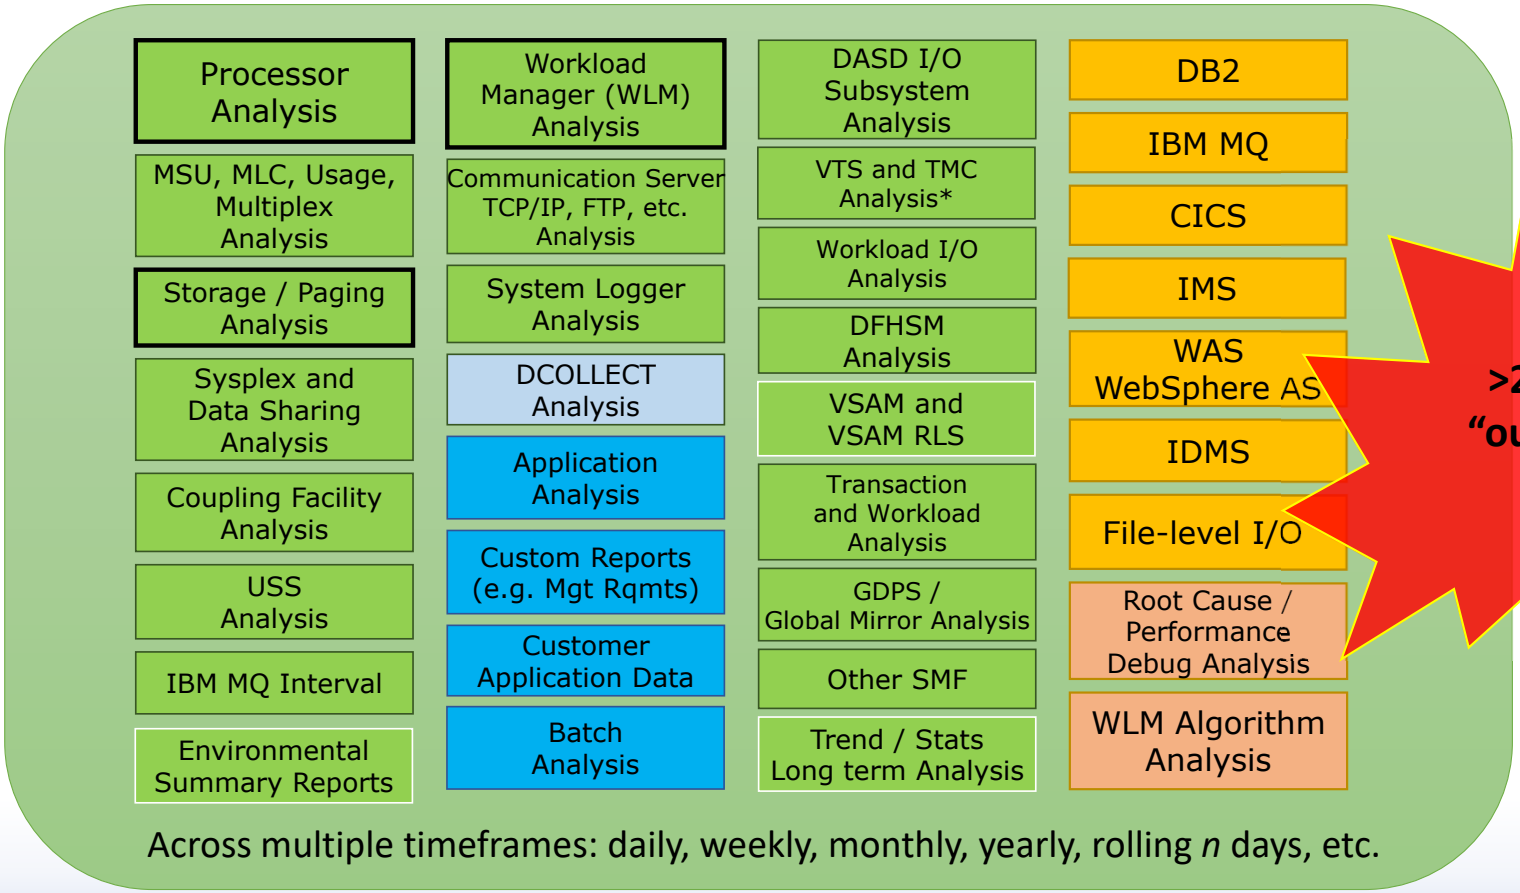

Like what you see?
Spring / Summer 2025 Webinars

- Free z/OS Performance Educational webinars!
  - The titles for our Spring / Summer 2025 webinars are as follows:
    - ✓ *Overseeing z/OS Performance Management With Your Outsourcer*
    - ✓ *Back to basics - Processor Consumption Analysis*
    - ✓ *Pivotor Pointers*
    - ✓ *WLM and CPU Critical Control*
    - ✓ *Back to Basics - Evaluating Latent Demand*
    - ✓ *Understanding SMF 98 Locking Measurements*
    - ✓ *Standard Measurements when Monitoring Transactions*
    - ✓ *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*
    - *Optimization in the Application Trenches*
    - *Understanding SMF 98 Address Space Consumption Measurements*
    - *z/OS Performance Management in an AI World*
    - *Understanding z/Architecture Processor Topologies*

- 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

Summer / Fall 2024 Webinars

- Free z/OS Performance Educational webinars!
  - The titles for our Summer / Fall 2024 webinars are as follows:
    - ✓ *What a z/OS Guy Learned About AWS in 10 Years*
    - ✓ *Advantages of Multiple Period Service Classes*
    - ✓ *Understanding z/OS Connect Measurements*
    - ✓ *WLM and SMF 99.1 – System Measurements Deeper Dive*
    - ✓ *WLM and SMF 99.2 – Service Class Period Measurements Deeper Dive*
    - ✓ *Optimizing Performance at the Speed of Light: Why I/O Avoidance is Even More Important Today*
    - ✓ *Understanding MVS Busy % versus LPAR Busy % versus Physical Busy %*
    - ✓ *Rethinking IBM Software Cost Management Under Tailored Fit Pricing*
    - ✓ *Understanding Page Faults and Their Influence on Uncaptured Time*
    - ✓ *Response Time Goals: Average or Percentiles?*
    - ✓ *Understanding and Using Enclave*
  - We would love to hear your suggestions

Like what you see?

- 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
- We also have a ***free*** Pivotor offering available as well
  - 1 System, SMF 70-72 only, 7 Day rolling retention
  - That still encompasses over 100 reports!
- We offer free bi-monthly z/OS performance educational Webinars
  - Sign our guestbook to get invited
  - Or send an email to
    - Peter.Enrico@EPStrategies.com

Pivotor Software as a Solution (SaaS)

We do everything for you so, you can focus on Analysis!

Because we also have access to your data, 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, outsourcers, or other vendors

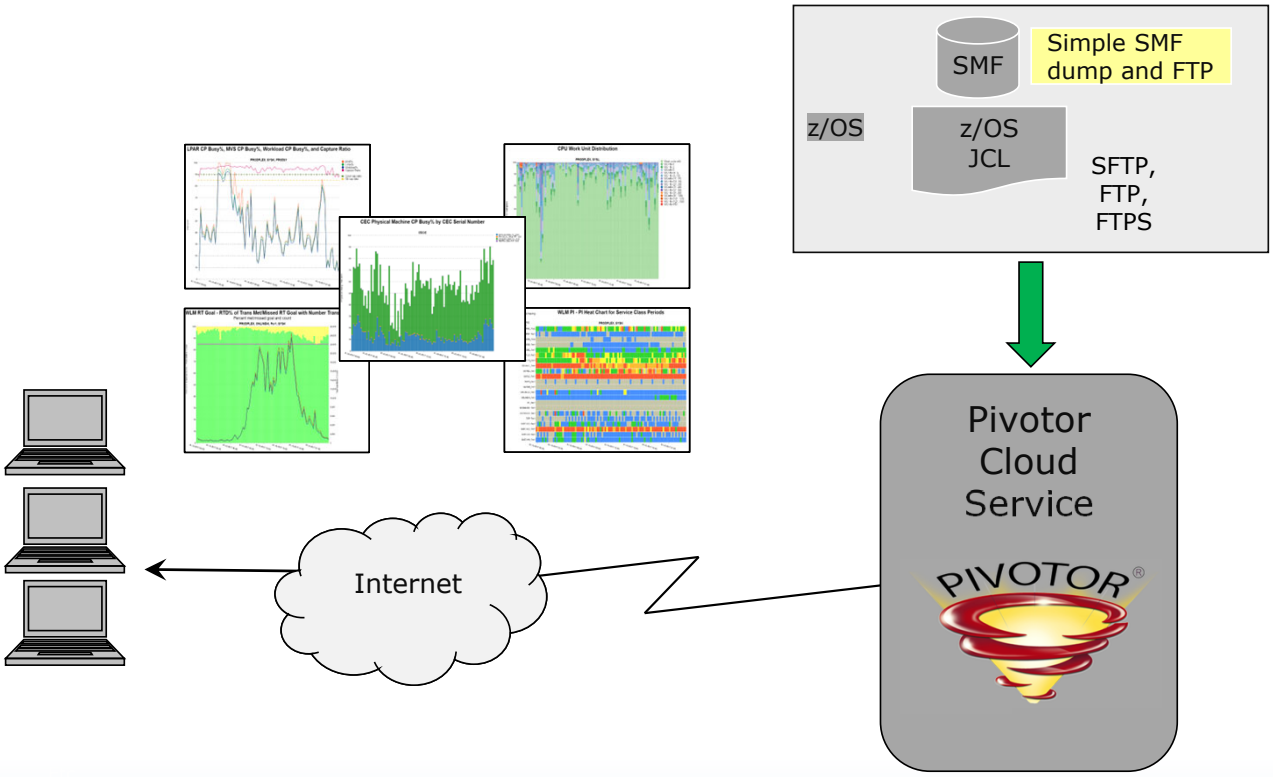

Cookbook Approach to Looking at WLM

- High level steps for revisiting your WLM setup and service definition
  - Step 0: Learn the basic concepts of WLM (and the oddities)
  - Step 1: Inventory Your Managed Resources
  - Step 2: Inventory System Workloads
  - Step 3: Understand Current WLM Definition and controls (like IEAOPTxx, etc.)
  - Step 4: Clean Up Your Service Definition
  - Step 5: Learn How to Interpret WLM Measurements
  - Step 6: Verify Properness of WLM Controls
  - Step 7: Determine Effectiveness of Controls
  - Step 8: Examine Workload Mixtures
  - Step 9: Re-evaluate Assigned Goals and Importance
  - Step 10: Explore Exploiting New and optional Functions
  - Step 11: Start to Tackle Those Difficult Issues

Subject Matter Note

- This presentation discusses the WLM traces that are meant to help us understand WLM decision making
  - This is a portion of WLM that is rarely looked at or investigated
  - However, it is also a topic more to show you what information and data is out there if you want to take a deeper look into WLM and how it works
  - It is not something I would suggest anyone spend a lot of time
  - But it can be, and at times ***is*** rather interesting

Recap of SMF 99 Subtypes

- Subtype 1
  - System level measurement data used for decision input
  - Trace of WLM actions
  - Written every 10 seconds (i.e. policy adjustment interval)
- Subtype 2
  - Service class period measurement data used for decision input
  - Written every 10 seconds (i.e. policy adjustment interval)
- Subtype 3
  - Service class period plot data
  - Written every 10 seconds (i.e. policy adjustment interval)
- Subtype 4
  - Service class device cluster information
  - Written every 10 seconds (i.e. policy adjustment interval)
- Subtype 5
  - Data about monitored address spaces
  - Written every 10 seconds (i.e. policy adjustment interval)

Recap of SMF 99 Subtypes

- Subtype 6
  - Service class period settings and measurements
  - Written every 10 seconds (i.e. policy adjustment interval)
- Subtype 7
  - Enterprise Storage Server ® (ESS) with Parallel Access Volumes (PAVs)
  - Written every 30 seconds (i.e. 3 policy adjustment intervals)
- Subtype 8
  - Information about LPAR CPU management
  - Written every 10 seconds (i.e. policy adjustment interval)
- Subtype 9
  - Information about dynamic channel path management
  - Written every 10 seconds (i.e. policy adjustment interval)
- Subtype 10
  - Information about dynamic processor speed changes
  - Written when speed changes

Recap of SMF 99 Subtypes

- Subtype 11
  - Information about Group Capacity Limits
  - Written every 5 minutes
- Subtype 12
  - HiperDispatch interval data
  - Written every 2 seconds (i.e. policy adjustment interval)
- Subtype 13
  - HiperDispatch IBM internal use only (so undocumented)
- Subtype 14
  - HiperDispatch topology data
  - Written every 5 minutes

SMF 99 Recommendations

- Pivotor Customers
  - Send us all your SMF 99 data everyday for heath check analysis and, if needed, performance debug

- Non-Pivotor Customers
  - Consider regularly collecting the following SMF 99 subtypes
    - Subtype 6 - Service class period settings and measurements
    - Subtype 11 - Information about Group Capacity Limits
    - Subtype 12 - HiperDispatch interval data
    - Subtype 14 - HiperDispatch topology data

    - Collectively these records typically produce about 40MiB/system/day
    - They contain the most interesting and useful data of the 99s

  - Records to collect for problem periods of time, or when doing a study to better understand WLM decision making
    - Subtype 1 - System level measurement and trace data used for decisions
    - Subtype 2 - Service class period measurement data used for decision input
    - Subtype 3 - Service class period plot data
    - Subtype 5 - Data about monitored address spaces

- If in need, contact Peter Enrico, Scott Chapman, or support@epstrategies.com to process with Pivotor and discuss

Ways of Approaching a WLM analysis

- There are multiple ways to approach a WLM analysis
  - Real Time monitors
    - OK, but usually for 'bubble gum' for the eyes and to see what is going on 'now'
  - SMF 72, subtype 3 records
    - When you need to wade into the waters of WLM
    - Great for daily monitoring and for WLM health checks
  - SMF 99, subtypes 6, 12, 14
    - When you need gain insights into transient performance issues and periods of time
    - Examples:
      - How the WLM controls may have changed on a 10 second basis
      - Processor busy and HiperDispatch on a 2 second basis, etc.
  - SMF 99, all subtypes (including subtype 1 and 2)
    - When you need to take a super deep dive into WLM decisions when you need to gain an understanding of 'what the heck really happened'
    - To look for patterns to help gaining deep insights into system performance

When we look at SMF 99 records

- Walking straight into an performance problem or concern
  - We dive into the data to understand what happed
  - This is ‘walking straight into’ a WLM analysis with the purpose of tuning or debug

- Backing into a performance analysis
  - There may not be a performance problem or concern, but to better understand WLM or enable better WLM and system tuning recommendation
  - To help squeeze out even better performance or CPU savings
  - When we want to understand how new recommendation or a change to WLM code has affected the algorithms

SMF 99.1 Trace Records

Workload Manager Algorithms

- Internally to the z/OS workload manager are a series of algorithms that are used to ensure the installations performance objectives are met

- The SMF 99.1 trace records provide insights into the execution of these algorithm
  - Although note: IBM WLM development has gotten 'forgetful' over time to provide important trace information with new releases, and they do not always update the documentation as needed

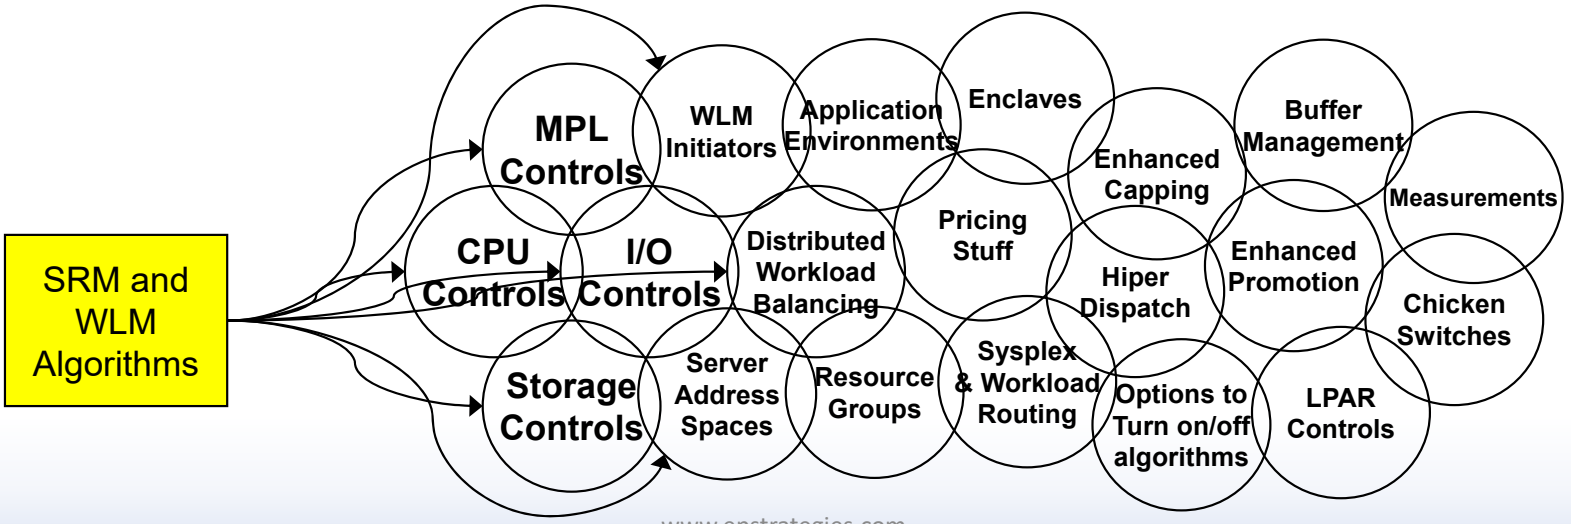

SMF 99 Traces:
WLM Policy Adjustment – 'The Loop'

- Summarize data for state of the system and workloads
- Select a receiver period (highest importance missing goal the most)
- Find the receiver's largest bottleneck
  - Determine fix for receiver's bottleneck
    - Determine if needed resources can be gotten from unused resources
    - Find donor(s) of resource that receiver needs
    - Assess effect of reallocating resources from donor(s) to receivers
    - If allocation has both net and receiver value
      Then commit change
      Else don't make change
  - If reallocation was done
    then jump to Exit and allow change to be absorbed
  - If reallocation was not done
    then try to fix receiver's next largest bottleneck
- If cannot help receiver
  then look for next receiver (highest importance missing goal the most)
- Exit
  - Housekeep current set of controls

Useful Terms:
WLM algorithm phases of processing

- To understand the WLM traces in the SMF 99 records, one must remember that there are two primary phases of WLM algorithms

- Policy Adjustment (PA)
  - Done approximately every 10 seconds (AKA 'PA interval')
  - Objectives include:
    - Summarize state of system and resources
    - Help work meet goals by setting resource controls
    - Housekeep resource controls that may be out of date

- Resource Adjustment (RA)
  - Done approximately every 2 seconds (AKA 'RA interval')
  - Objectives include:
    - improve efficiency of system resources
    - avoided if at the expense of goals

Useful Terms:
When WLM chooses Receivers and Donors

- Receiver – a service class period to help
  - WLM will help only one receiver during each policy adjustment interval

Goal Receiver	-	Period with goal that needs help
Resource Receiver	-	Period to whom WLM will give the resources in order to help the 'goal receiver'
Secondary Receiver	-	Period that is helped indirectly due to an action to help the goal receiver

- Donor – a service class period to potentially 'donate' resources to help receiver
  - WLM may take from multiple donors during each policy adjustment interval

Goal Donor	-	Period whose goals may be impacted by resource donation
Resource Donor	-	Period to donate resources
Secondary Donor	-	Period that donates indirectly when receiver is helped

Understanding Who is Selected as Receiver

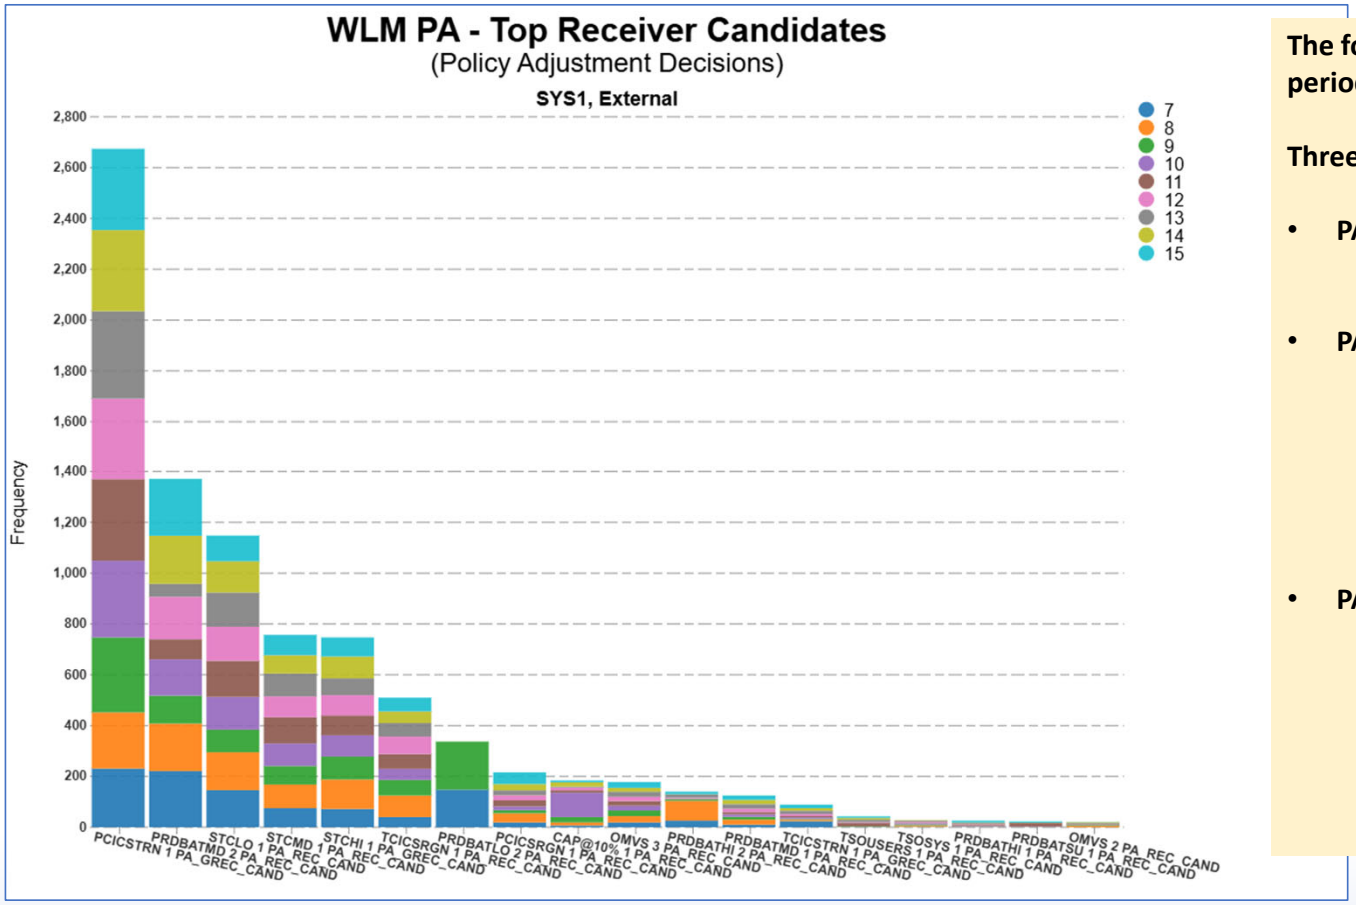

The following chart shows which service class periods were selected the most often as a receiver

Three types of 'primary' receivers:

- **PA_REC_CAND**
  - **When receiver is both a goal and resource receiver**
- **PA_GREC_CAND**
  - **When receiver is just a goal receiver**
  - **Typically, an CICS, IMS, or WAS transaction service class**
  - **Also, could be period executing a stored procedure**
- **PA_RREC_CAND**
  - **When receiver is just a resource receiver**
  - **Typically, a $SRMSxxx period**

Not on this chart, but there are also secondary receivers of each type

Maps Server Topology For Enclave Servers

- WLM can determine a server topology for enclave servers

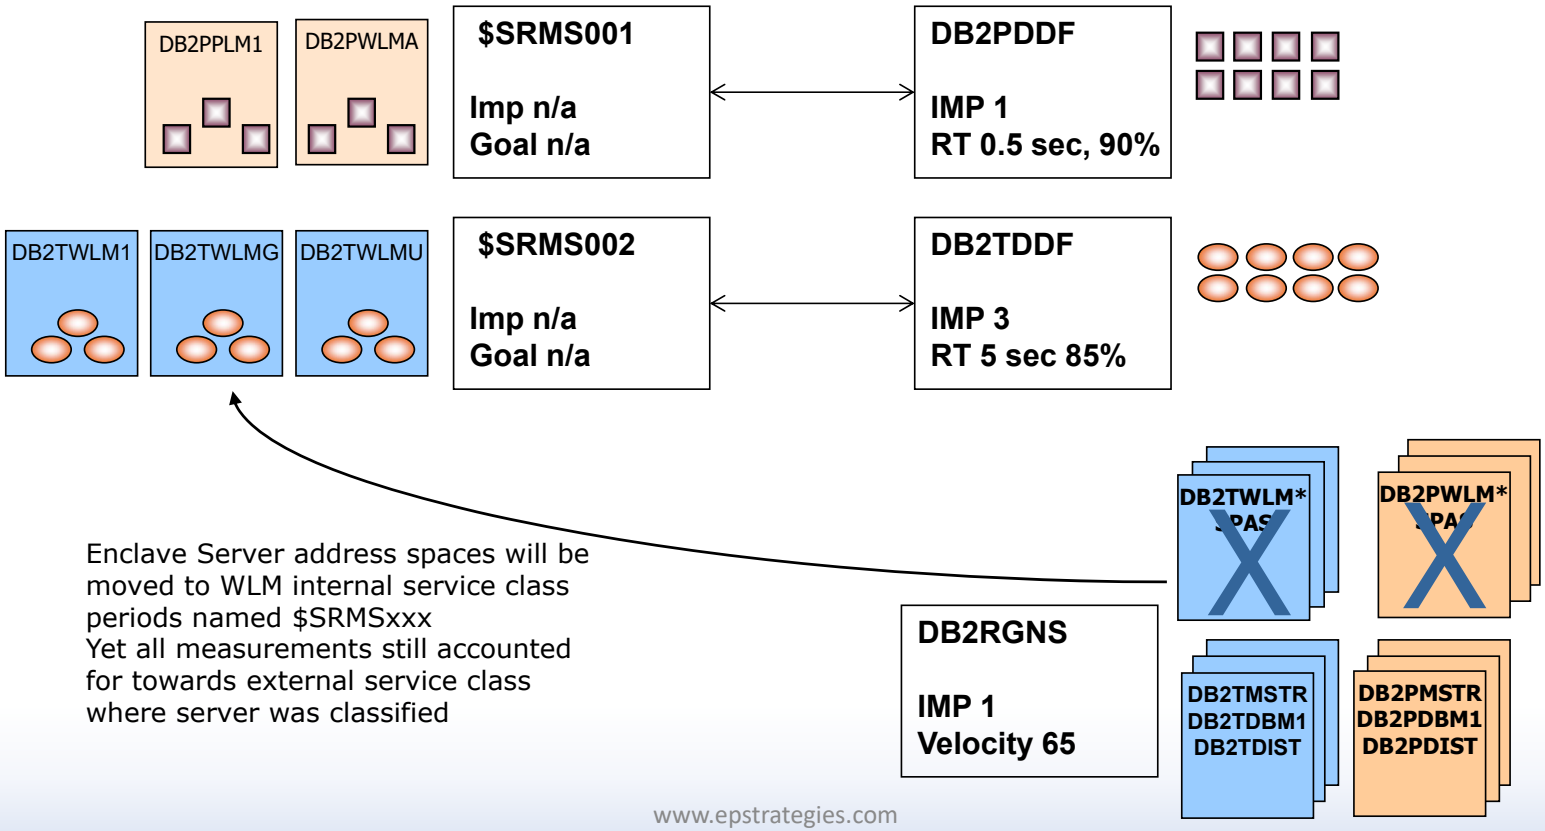

Server Topology for CICS– with BOTH option

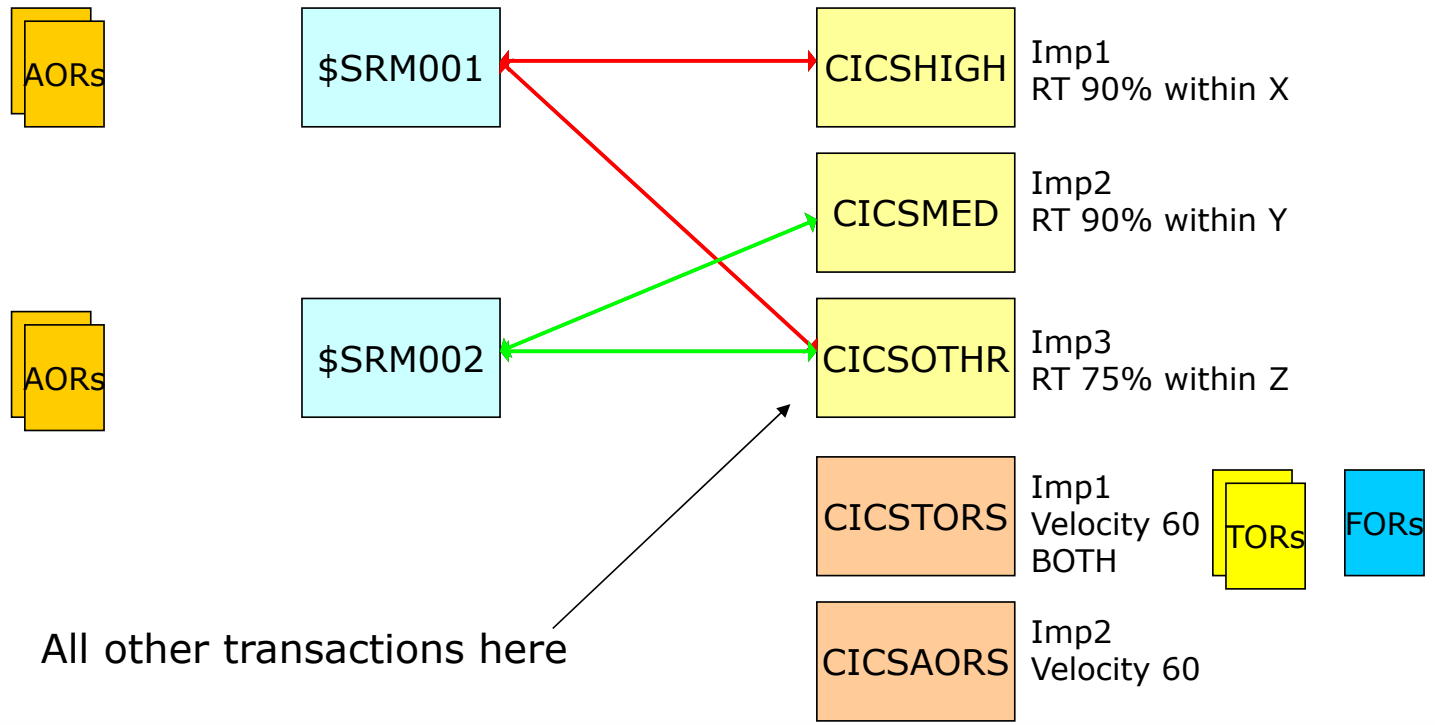

Understanding Who is Selected as Donor

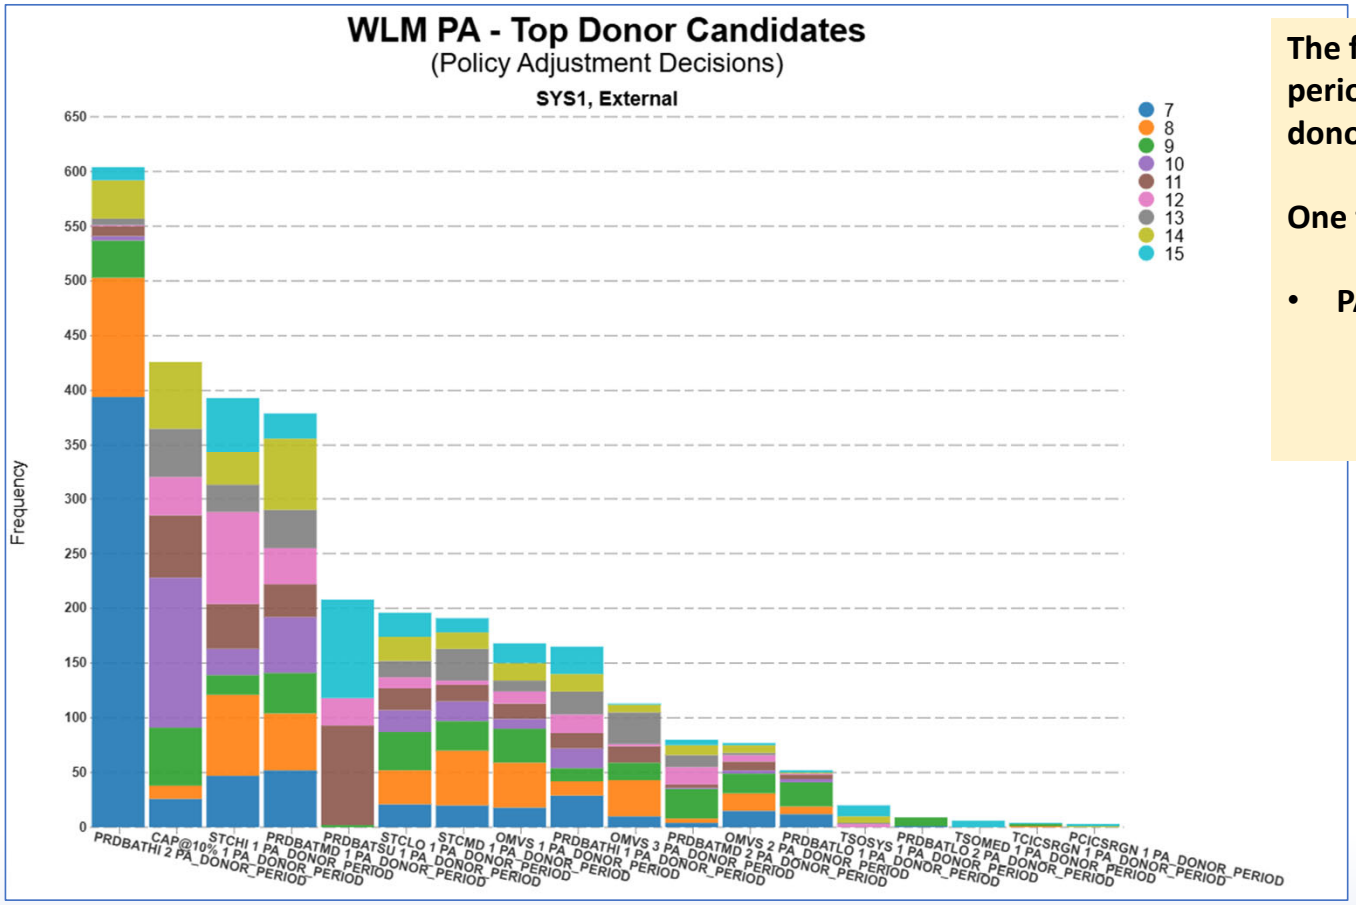

The following chart shows which service class periods were selected the most often as a donors

One type of 'primary' donor:

- **PA_DONOR_PERIOD**
  - **When receiver is both a goal and resource receiver**

Understanding Who is Selected as Receiver

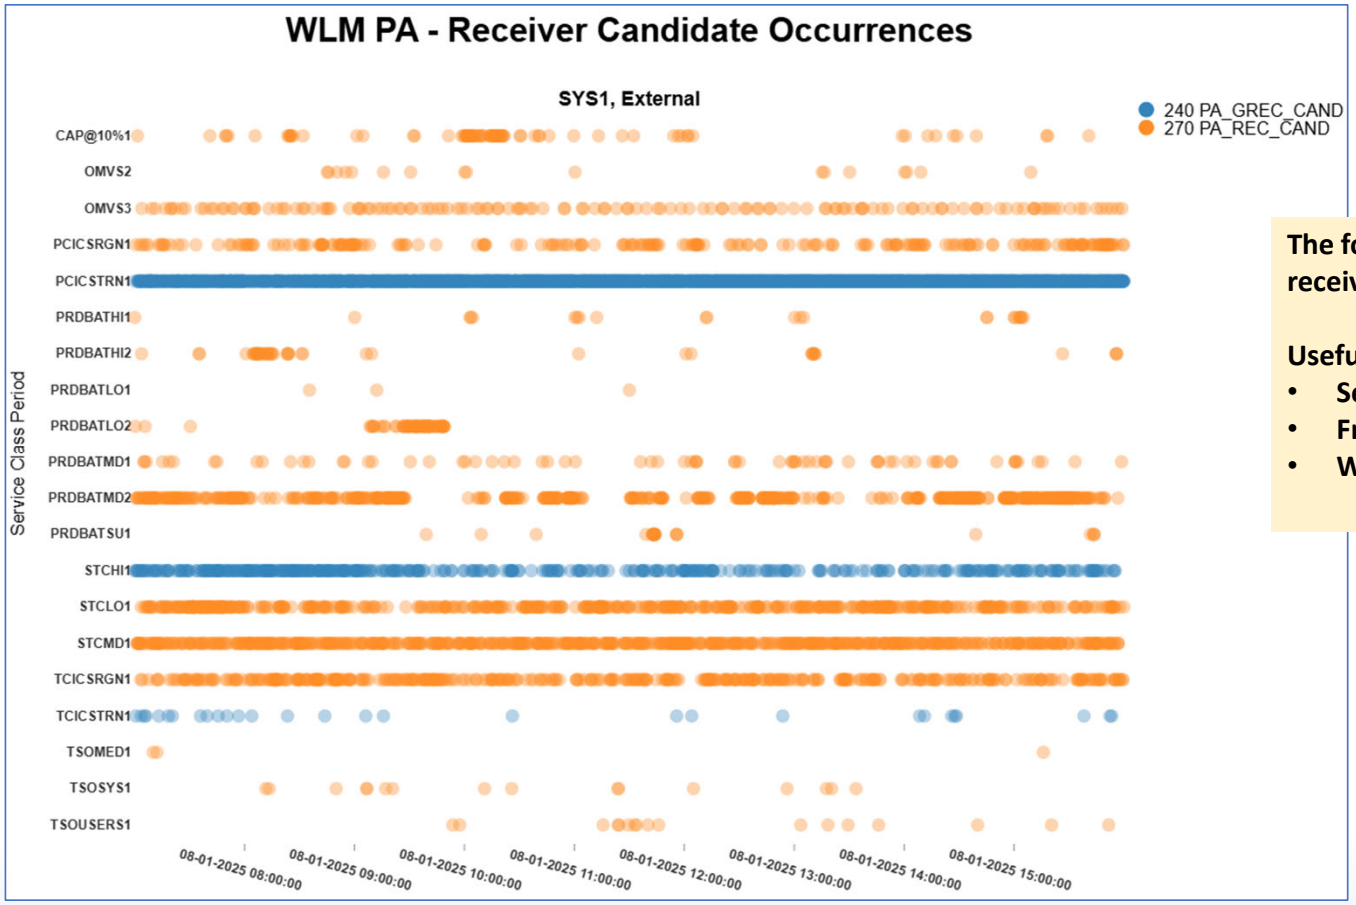

The following chart shows is a visualization of when receivers are selected every 10 seconds.

Useful for

- **Seeing selection patterns**
- **Frequency**
- **Who is being selected when**

Periods whose goals may be too difficult

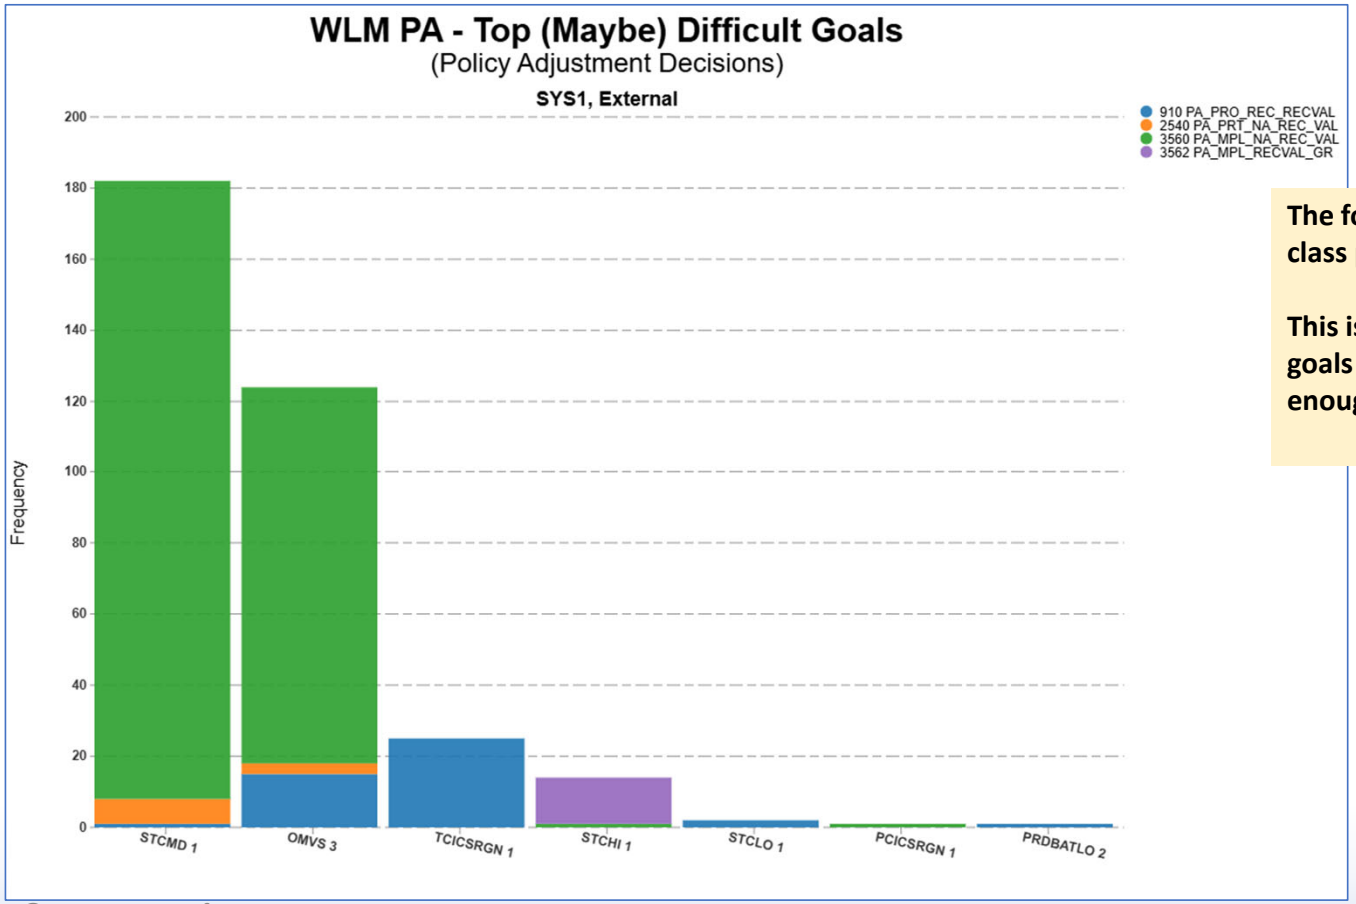

The following chart shows is a visualization of service class periods whose goals may be too difficult.

This is gotten from those periods that were missing goals but were never changed because there was not enough 'receiver value'.

Velocity Achieved for STCMD

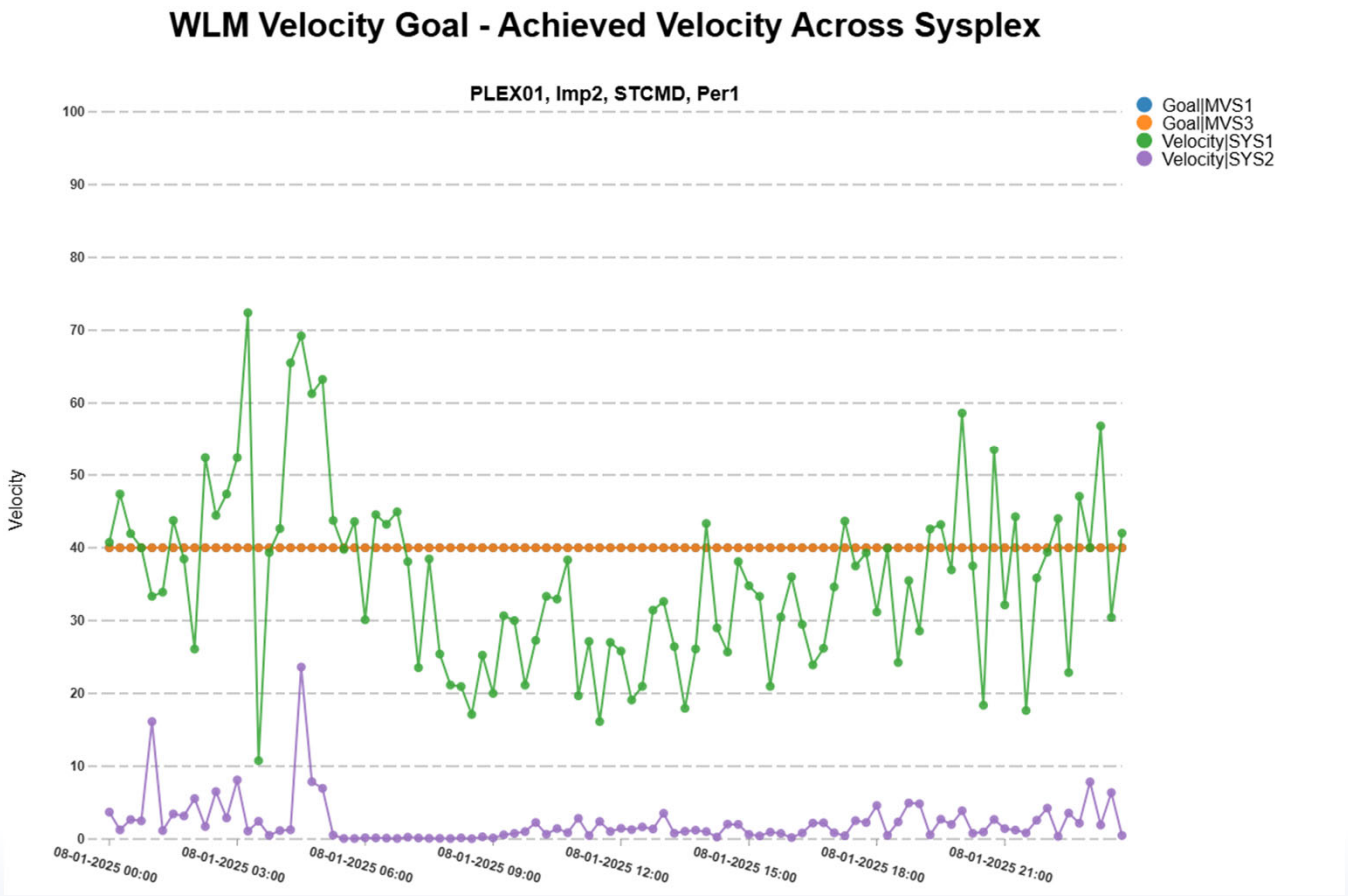

WLM PA Loop: Receiver Value Check

- Receiver Value
  - Receiver helped only if there is projected to be sufficient *receiver value*
    - Designed to reject 'small or marginal improvements'
    - Allows WLM to get on to addressing larger problems for other periods
  - Minimum projected improvement to make change worth the effort
    - Projected PI improvement
    - or projected minimum group service increase
    - or some other projected minimum criteria
- Guideline:
  - Projected PI improvement is the larger of (10% of the PI change to meet goal) or (0.05)
  - Or Reduction in delay samples is at least half of the largest delay
- Example:
  - PRODTSO period 1 PI = 3.5
  - WLM algorithms suggest improvements can bring PI to 3.46
  - Don't take action

WLM PA Loop: Net Value Check

- Receiver is only helped by a specific donor if there is projected to be sufficient *net value*
  - Designed to reject changes that will harm the donor more than the projected improvement to the receiver
  - Allows WLM to assess taking from other donors
- All external service policy specifications are considered for both primary and secondary donors
  - goals
  - importance
  - resource group minimums and maximums
- Example
  - PRODBAT PI = 4.0
  - WLM algorithms suggest improvements can bring PI to 3.0
  - Change hurts donor more then helps receiver

WLM PA Loop: Net Value Check

- Net Value Check - Very high level logic
  - If receiver is more important than donor
    - Make the move if receiver is missing goal
  - If receiver is less important than donor
    - Never make the move if donor is missing goal
    - or is projected to miss goal
  - If receiver and donor are equal importance
    - Receiver's PI benefit is more than donor's loss, and
    - less disparity in projected PIs
- Resource group minimums and maximums make this even more complicated

CPU Adjustment Algorithms

Example: WLM Possible WLM Actions - CPU

- Dispatching Priority
  - Priority adjustment for
    - Periods with goals or server period
    - Discretionary periods in a resource group
  - Small consumer
    - For periods that use very little CPU
    - Gets these periods 'out of the way' of critical adjustments
  - Actions include:
    - Increase Receiver's priority
    - Decrease Donor's priority
      - Decreased service consumption and/or increased wait-to-using ratio
    - Both

255	SYSTEM
254	SYSSTC
253 249	'Unused' (SYSSTC1-5)
248	Small Consumer
247 203	Priorities Used for RT or Velocity Periods (i.e. Imp 1 – 5)
202	Unused
201 192	Discretionary (MTTW)
191	Quiesce

Example: WLM Possible WLM Actions - CPU

- WLM will model and project effects of dispatching priority adjustments
  - Objective: Increase Receiver's CPU using, or decrease Receiver's CPU delay
  - Interesting concepts:
    - Wait-to-Using ratio - ratio of CPU delay samples to CPU using samples (change in ratio used to determine change in CPU delay)
    - Maximum demand
      - Theoretical maximum percentage of total processor time a period can consume if it had no CPU delay
    - Achievable maximum demand
      - Percentage of total processor time a service period is projected to consume, taking into account demand of all higher work
  - Some possible actions

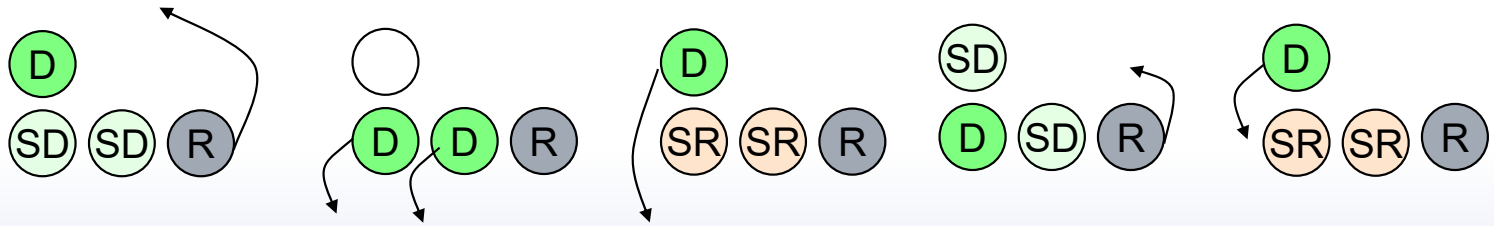

Example of CPU DP Changes

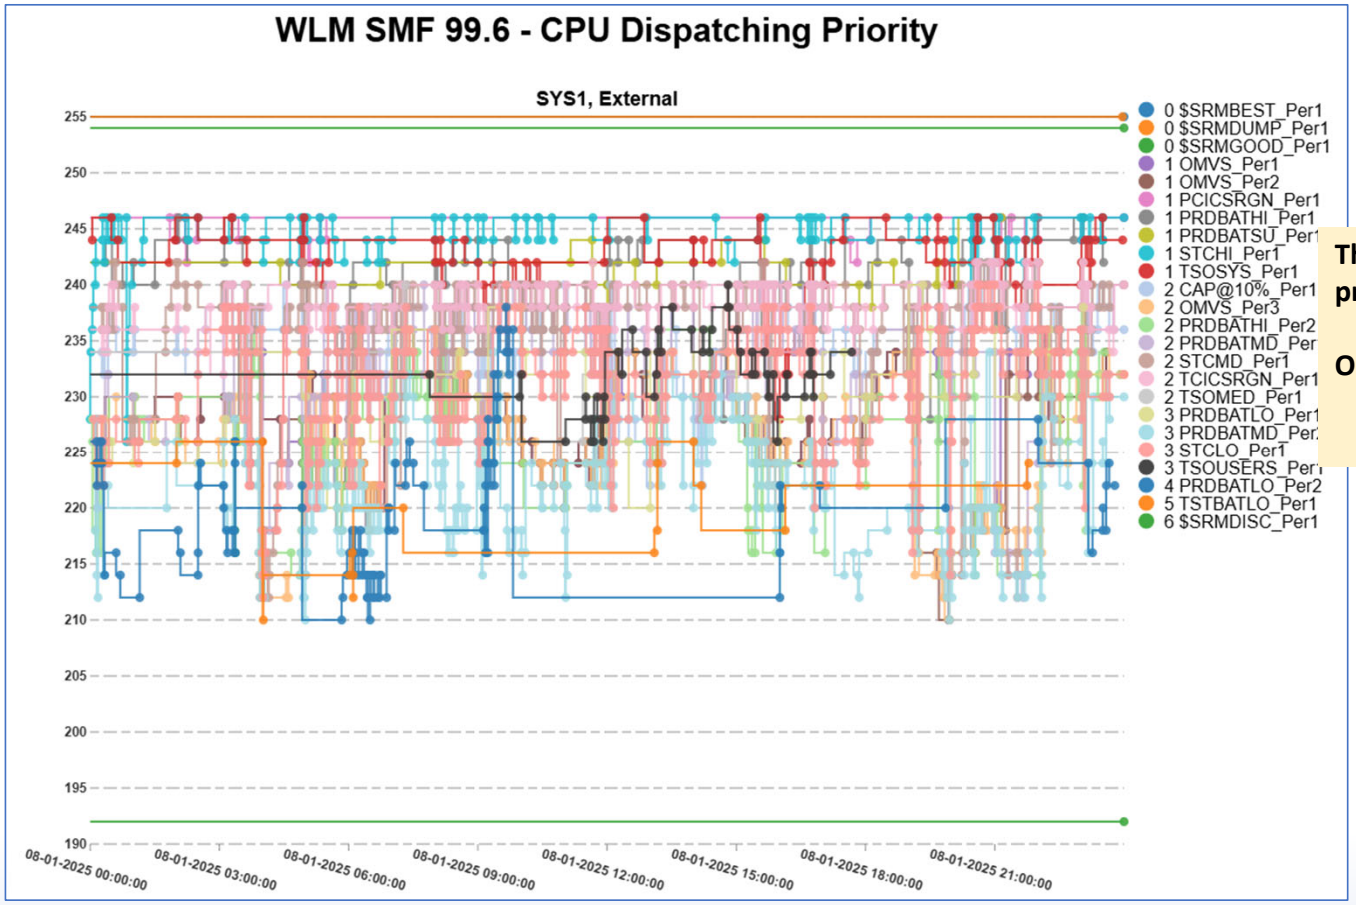

The following chart shows the CPU dispatching priority changes for 24 hours

Only shown are external service class periods

Example of CPU DP Changes for Hour 8

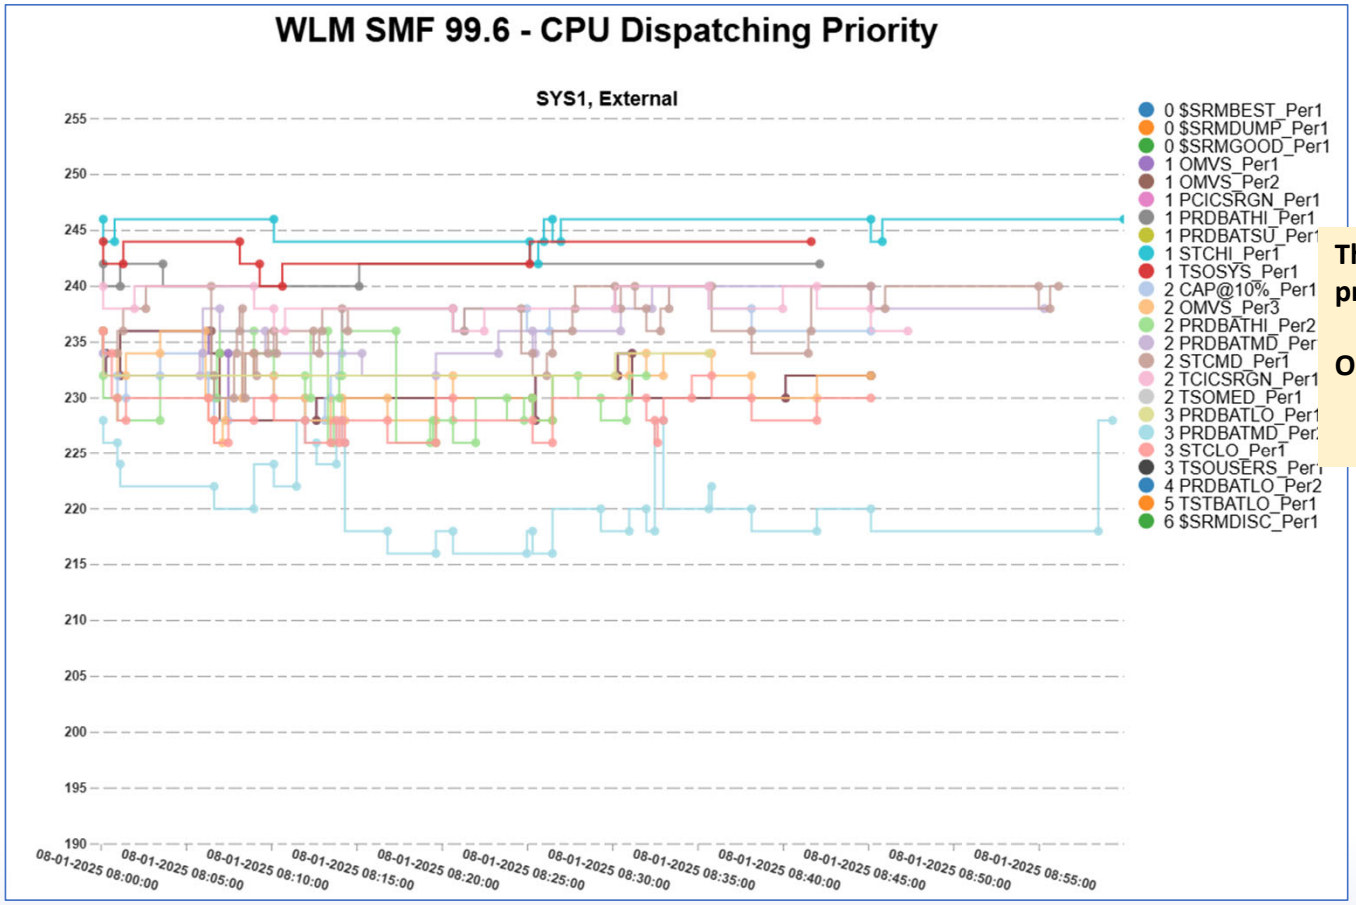

The following chart shows the CPU dispatching priority changes for hour 8.

Only shown are external service class periods

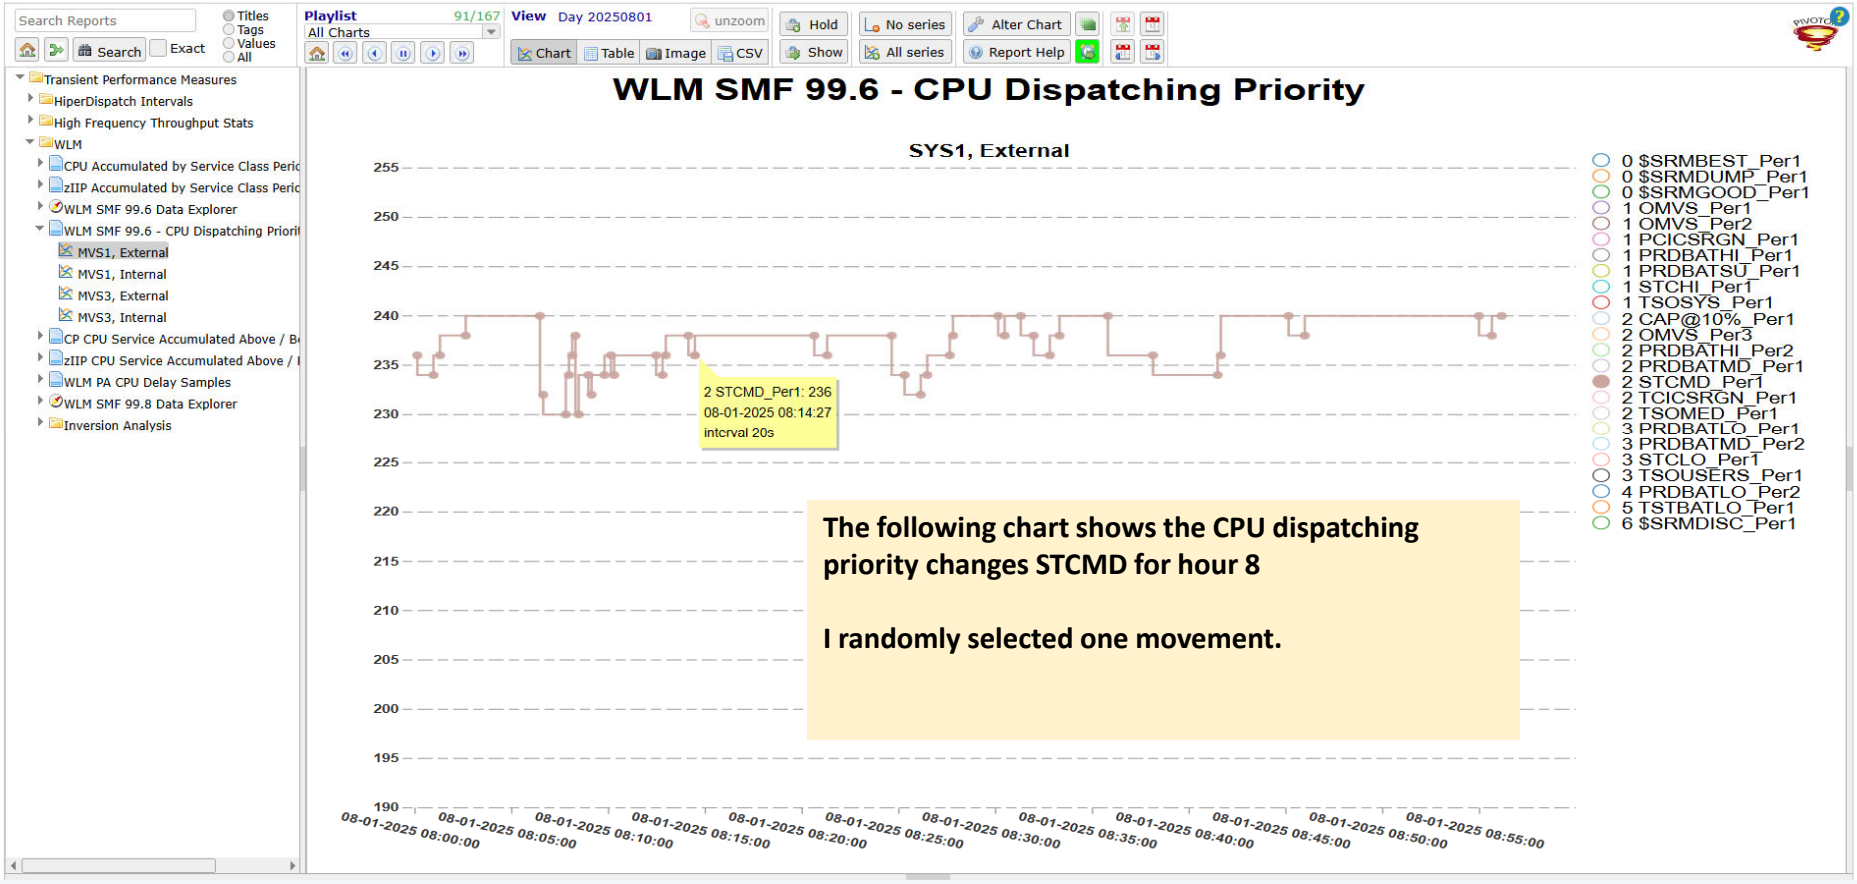
Search Reports
Titles
Tags
Values
All
Search
Exact
Playlist 91/167
All Charts
View Day 20250801
unzoom
Hold
No series
Alter Chart
Chart
Table
Image
CSV
Show
All series
Report Help
Transient Performance Measures
HiperDispatch Intervals
High Frequency Throughput Stats
WLM
CPU Accumulated by Service Class Perio
zIIP Accumulated by Service Class Perio
WLM SMF 99.6 Data Explorer
WLM SMF 99.6 - CPU Dispatching Priorit
MVS1, External
MVS1, Internal
MVS3, External
MVS3, Internal
CP CPU Service Accumulated Above / B
zIIP CPU Service Accumulated Above / I
WLM PA CPU Delay Samples
WLM SMF 99.8 Data Explorer
Inversion Analysis
WLM SMF 99.6 - CPU Dispatching Priority
SYS1, External
255
250
245
240
235
230
225
220
215
210
205
200
195
190
2 STCMD_Per1: 236
08-01-2025 08:14:27
interval 20s
08-01-2025 08:00:00
08-01-2025 08:05:00
08-01-2025 08:10:00
08-01-2025 08:15:00
08-01-2025 08:20:00
08-01-2025 08:25:00
08-01-2025 08:30:00
08-01-2025 08:35:00
08-01-2025 08:40:00
08-01-2025 08:45:00
08-01-2025 08:50:00
08-01-2025 08:55:00
0 $SRMBEST_Per1
0 $SRMDUMP_Per1
0 $SRMGOOD_Per1
1 OMVS_Per1
1 OMVS_Per2
1 PCICSRGN_Per1
1 PRDBATHI_Per1
1 PRDBATSU_Per1
1 STCHI_Per1
1 TSOSYS_Per1
2 CAP@10%_Per1
2 OMVS_Per3
2 PRDBATHI_Per2
2 PRDBATMD_Per1
2 STCMD_Per1
2 TCICSRGN_Per1
2 TSOMED_Per1
3 PRDBATLO_Per1
3 PRDBATMD_Per2
3 STCLO_Per1
3 TSOUSERS_Per1
4 PRDBATLO_Per2
5 TSTBATLO_Per1
6 $SRMDISC_Per1
The following chart shows the CPU dispatching priority changes STCMD for hour 8
I randomly selected one movement.

Periods whose CPU DPs have changed

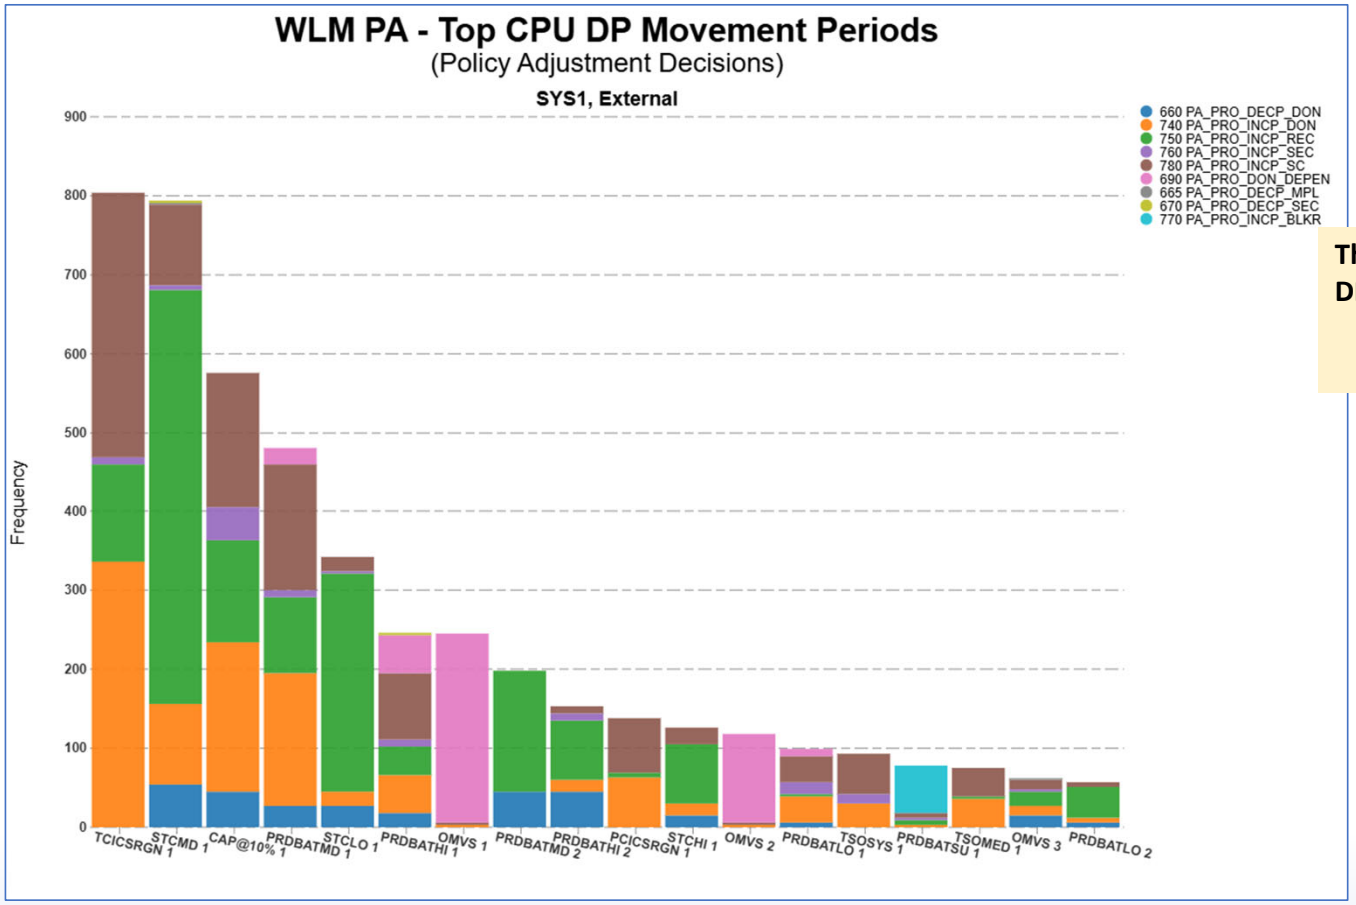

The following chart shows the periods whose CPU DPs changed the most often, and for which reason

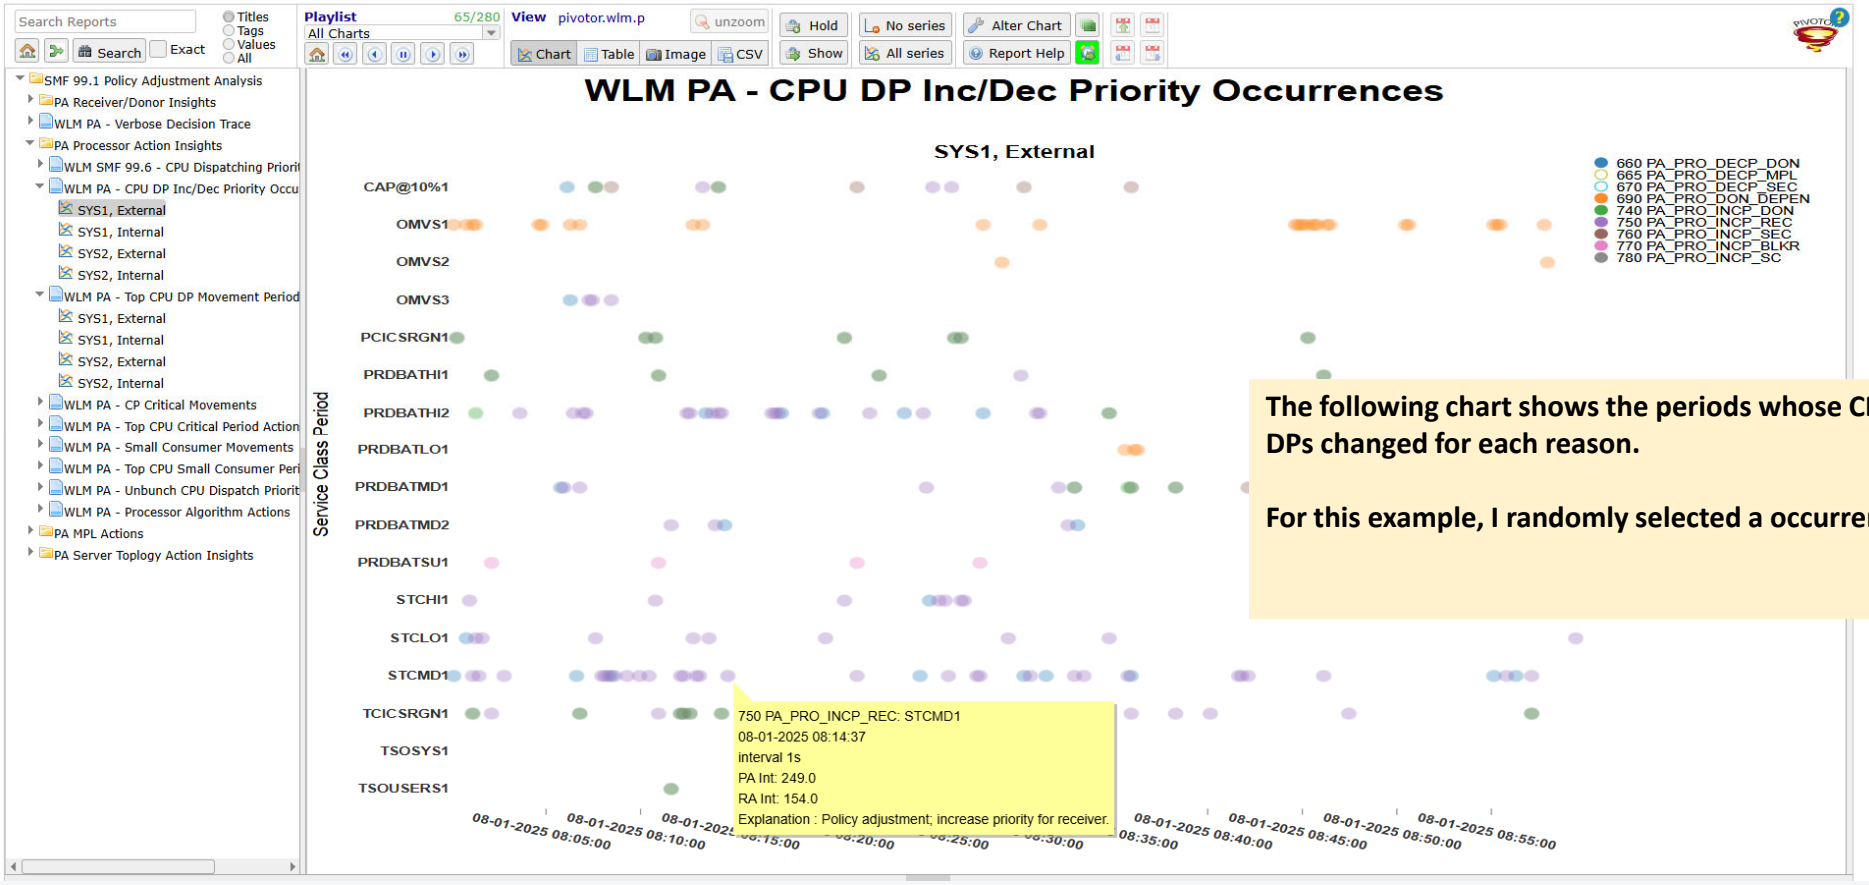

The following chart shows the periods whose CPU DPs changed for each reason.

For this example, I randomly selected a occurrence.

Example of Trace

smftime	SMF99_Trace_Event_Number	SMF99_TPID	SMF99_TRID	SMF99_TCOD	SMF99_TCNM	SMF99_TPER	SMF99_TLPI	SMF99_TSPI	SMF99_TJOB	SMF99_equate_symbol	SMF99_description
8:14:37 AM	23	249	154	9331		0	0	0		PA_LPCAP_PATTERN	Logical partition is to have capping turned on and off to enforce a soft cap.
8:14:37 AM	24	249	154	4747	$SRMS027	1	3	3	B0HOGPF	inc_ipi_tar_blw_bw	The protective processor target cannot be increamented the new target would be below threshol
8:14:37 AM	25	249	154	4747	$SRMS027	1	3	3	CI11PA10	inc_ipi_tar_blw_bw	The protective processor target cannot be increamented the new target would be below threshol
8:14:37 AM	26	249	154	4747	$SRMS027	1	3	3	CI11PA11	inc_ipi_tar_blw_bw	The protective processor target cannot be increamented the new target would be below threshol
8:14:37 AM	27	249	154	4747	$SRMS027	1	3	3	CI11PT10	inc_ipi_tar_blw_bw	The protective processor target cannot be increamented the new target would be below threshol
8:14:37 AM	28	249	154	4747	$SRMS027	1	3	3	CICSCM12	inc_ipi_tar_blw_bw	The protective processor target cannot be increamented the new target would be below threshol
8:14:37 AM	29	249	154	4747	$SRMS027	1	3	3	CICSPA18	inc_ipi_tar_blw_bw	The protective processor target cannot be increamented the new target would be below threshol
8:14:37 AM	30	249	154	4747	$SRMS027	1	3	3	CICSWU12	inc_ipi_tar_blw_bw	The protective processor target cannot be increamented the new target would be below threshol
8:14:37 AM	31	249	154	4747	$SRMS028	1	0.01	0.01	D0POGT2	inc_ipi_tar_blw_bw	The protective processor target cannot be increamented the new target would be below threshol
8:14:37 AM	32	249	154	3620		0	0	0		TDH_NA_INI_BAL	Balancing of initiators; no action.
8:14:37 AM	33	249	154	270	STCMD	1	40	40		PA_REC_CAND	Policy adjustment; receiver candidate selected.
8:14:37 AM	34	249	154	975	TCICSRGN	1	4.44	4.44		PA_SDO_DONFAIL_SPC	Policy adjustment; select donor failed selecting period as the donor because period is small consu
8:14:37 AM	35	249	154	975	TSOSYS	1	1	1		PA_SDO_DONFAIL_SPC	Policy adjustment; select donor failed selecting period as the donor because period is small consu
8:14:37 AM	36	249	154	975	CAP@10%	1	1.21	1.21		PA_SDO_DONFAIL_SPC	Policy adjustment; select donor failed selecting period as the donor because period is small consu
8:14:37 AM	37	249	154	308	STCHI	1	0.6	0.6		PA_DONOR_PERIOD	Policy adjustment; donor period.
8:14:37 AM	38	249	154	880	STCHI	1	0.6	0.6		PA_PRO_RDON_CAND	Policy adjustment; processor resource donor candidate selected.
8:14:37 AM	39	249	154	620	STCMD	1	40	40		PA_PMUO_REC	Policy adjustment; assess moving primary processor receiver up to occupied priority.
8:14:37 AM	40	249	154	620	STCMD	1	40	40		PA_PMUO_REC	Policy adjustment; assess moving primary processor receiver up to occupied priority.
8:14:37 AM	41	249	154	620	STCMD	1	40	40		PA_PMUO_REC	Policy adjustment; assess moving primary processor receiver up to occupied priority.
8:14:37 AM	42	249	154	960	PRDBATHI	2	1.08	1.08		PA_PRO_UNC_SEC_DON	Policy adjustment; unchanged secondary donor.
8:14:37 AM	43	249	154	960	PRDBATHI	2	1.08	1.08		PA_PRO_UNC_SEC_DON	Policy adjustment; unchanged secondary donor.
8:14:37 AM	44	249	154	960	PRDBATHI	2	1.08	1.08		PA_PRO_UNC_SEC_DON	Policy adjustment; unchanged secondary donor.
8:14:37 AM	45	249	154	960	TCICSRGN	1	4.44	4.44		PA_PRO_UNC_SEC_DON	Policy adjustment; unchanged secondary donor.
8:14:37 AM	46	249	154	960	TCICSRGN	1	4.44	4.44		PA_PRO_UNC_SEC_DON	Policy adjustment; unchanged secondary donor.
8:14:37 AM	47	249	154	960	TCICSRGN	1	4.44	4.44		PA_PRO_UNC_SEC_DON	Policy adjustment; unchanged secondary donor.
8:14:37 AM	48	249	154	940	STCHI	1	0.6	0.6		PA_PRO_UNC_DON	Policy adjustment; unchanged donor.
8:14:37 AM	49	249	154	940	STCHI	1	0.6	0.6		PA_PRO_UNC_DON	Policy adjustment; unchanged donor.
8:14:37 AM	50	249	154	940	STCHI	1	0.6	0.6		PA_PRO_UNC_DON	Policy adjustment; unchanged donor.
8:14:37 AM	51	249	154	960	CAP@10%	1	2	2		PA_PRO_UNC_SEC_DON	Policy adjustment; unchanged secondary donor.
8:14:37 AM	52	249	154	960	CAP@10%	1	2	2		PA_PRO_UNC_SEC_DON	Policy adjustment; unchanged secondary donor.
8:14:37 AM	53	249	154	960	CAP@10%	1	2	2		PA_PRO_UNC_SEC_DON	Policy adjustment; unchanged secondary donor.
8:14:37 AM	54	249	154	750	STCMD	1	0.4	0.8		PA_PRO_INCP_REC	Policy adjustment; increase priority for receiver.
8:14:37 AM	55	249	154	750	STCMD	1	0.4	0.8		PA_PRO_INCP_REC	Policy adjustment; increase priority for receiver.
8:14:37 AM	56	249	154	750	STCMD	1	0.4	0.8		PA_PRO_INCP_REC	Policy adjustment; increase priority for receiver.

CPU Critical Control- Background

- Some installations are concerned that WLM will not react fast enough for high priority work

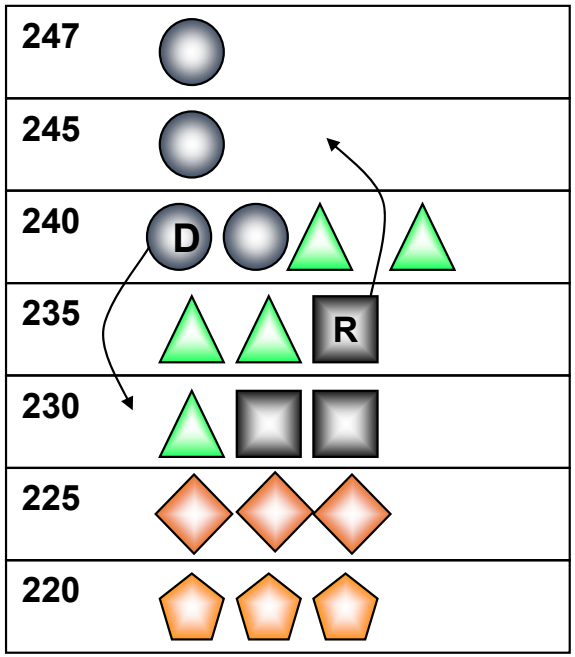

Note: To make the point, just a few priorities between DP 203 and DP247 are shown.

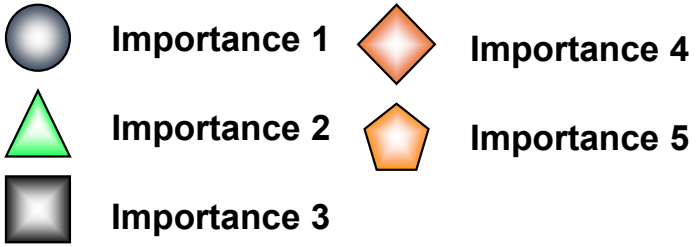

- With well set predictable goals, DPs tend to be ordered by importance
- If work is missing its goal WLM may decide to adjust its DP equal or above a higher importance period
- The problem occurs when this lower importance period starts to consume more CPU and causes the higher importance period to miss its goal
- WLM will recognize this condition and fix it … but it can be slow to react

CPU Critical Actions

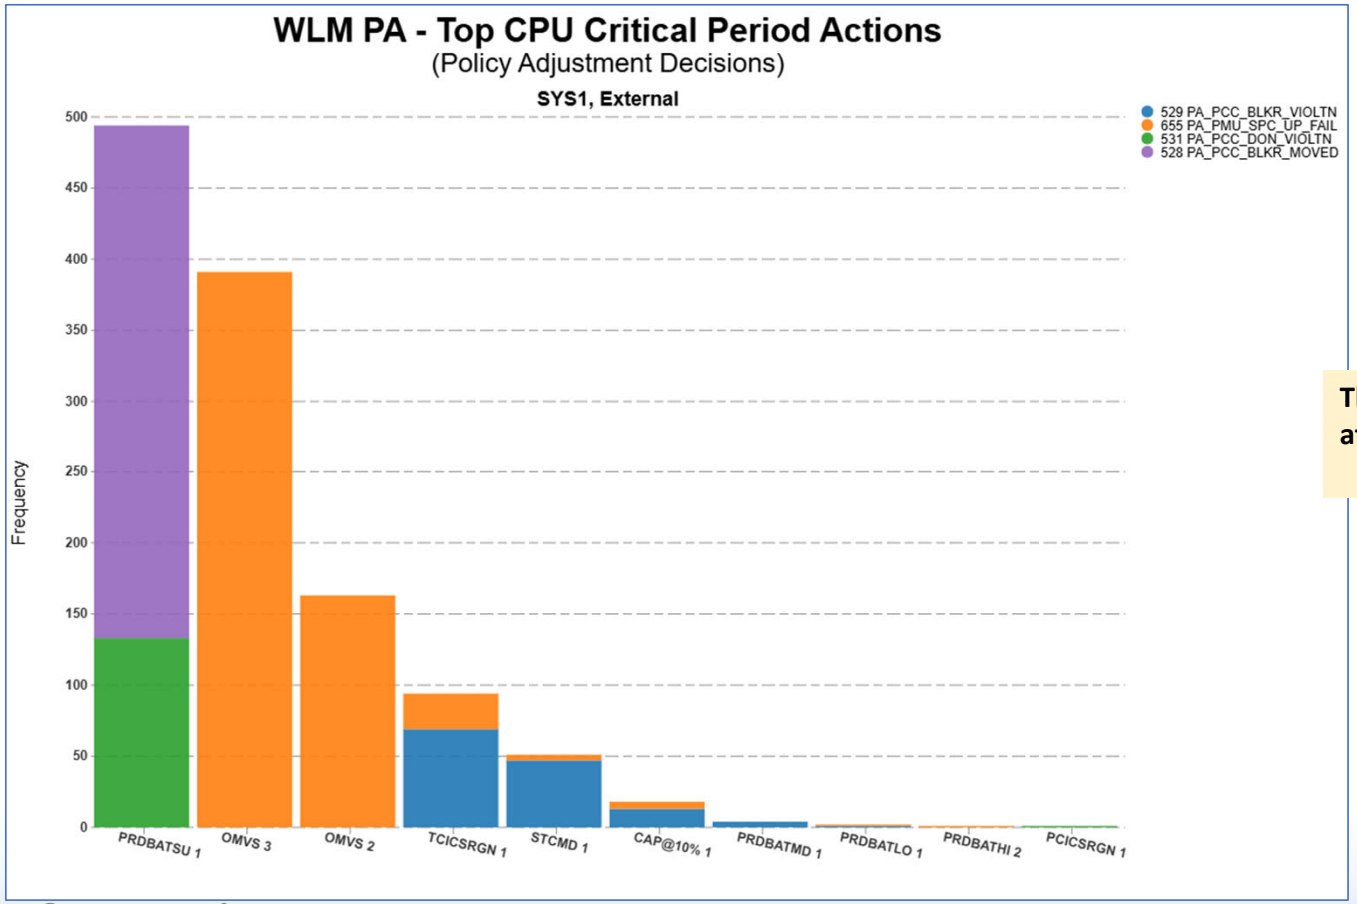

The following chart shows the periods whose were affected by the CPU critical setting

CPU Critical Actions

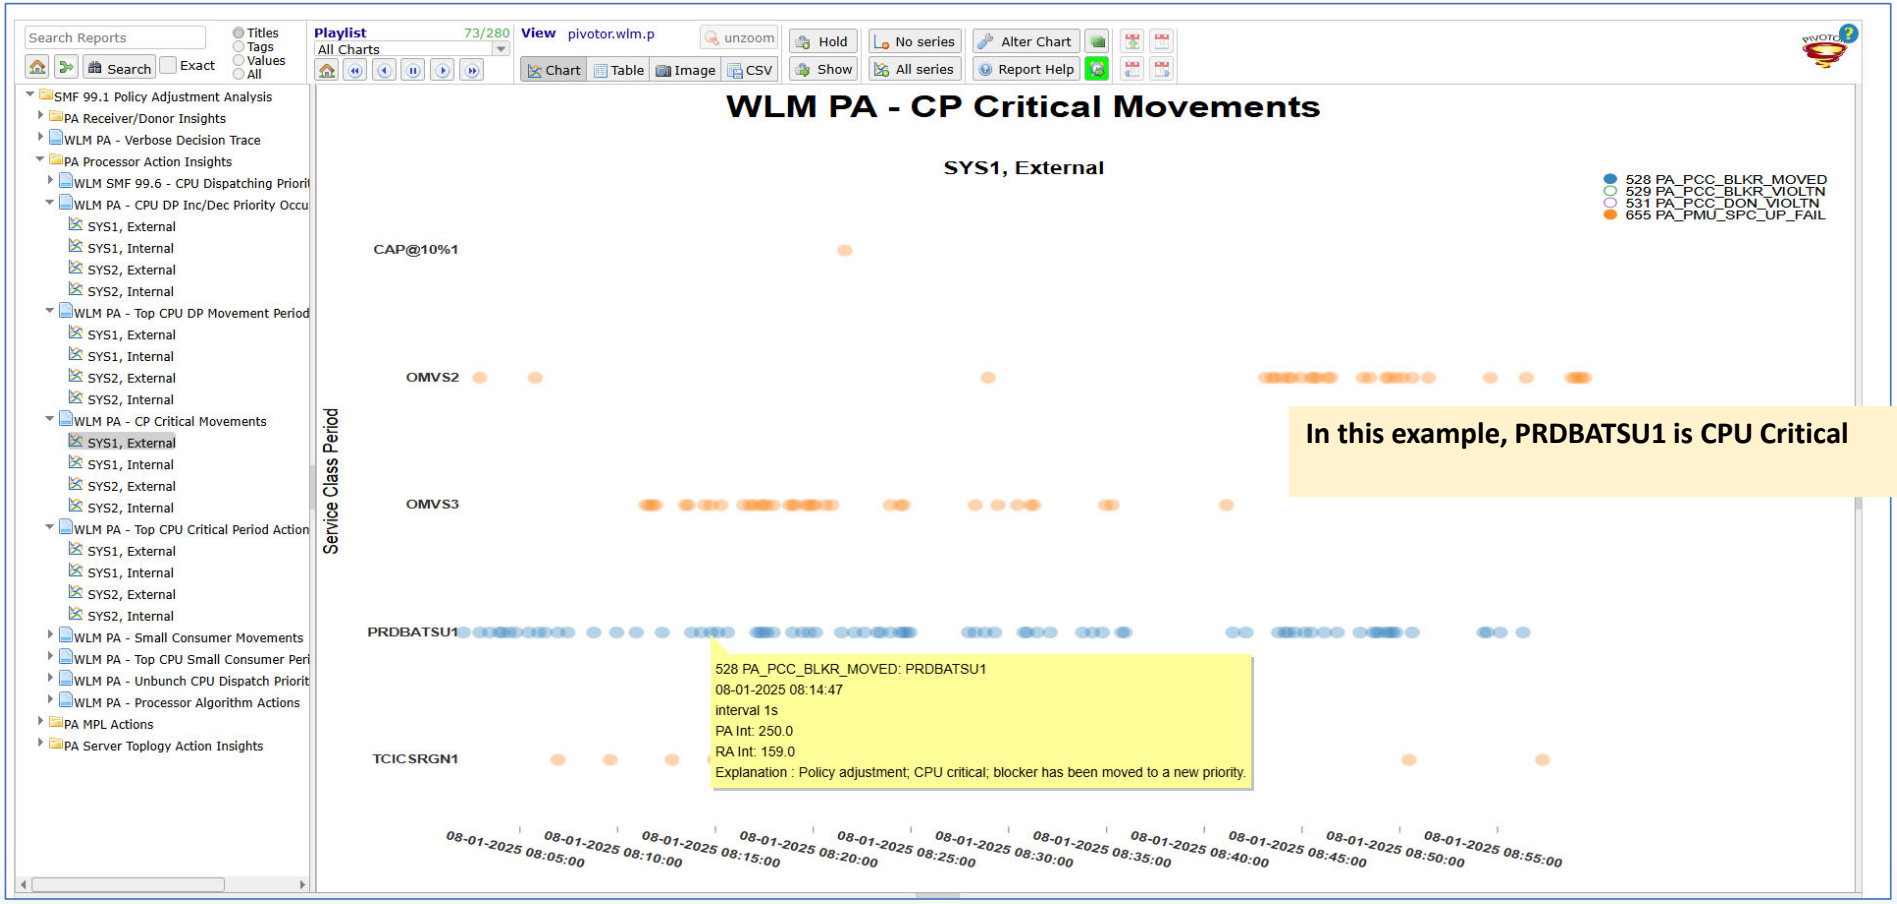

Small Consumer Periods

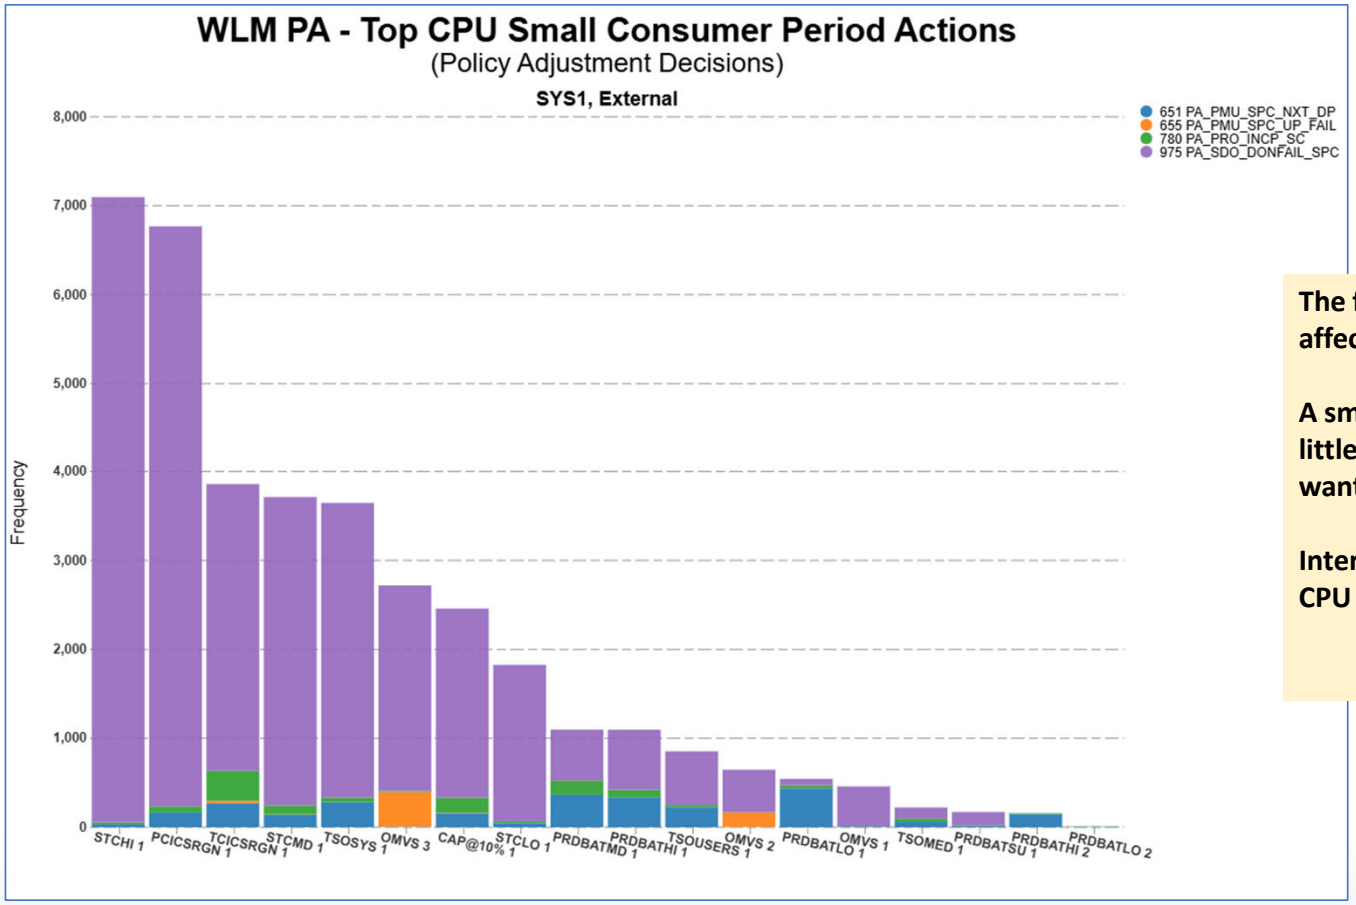

The following chart shows the periods that were affected due to small CPU consuming periods.

A small CPU consuming period is one that uses so little CPU, so WLM tries to just give it the CPU it wants to 'get it out of the way'

Interesting to note which periods do not have much CPU activity.

Server Topology for CICS– with BOTH option

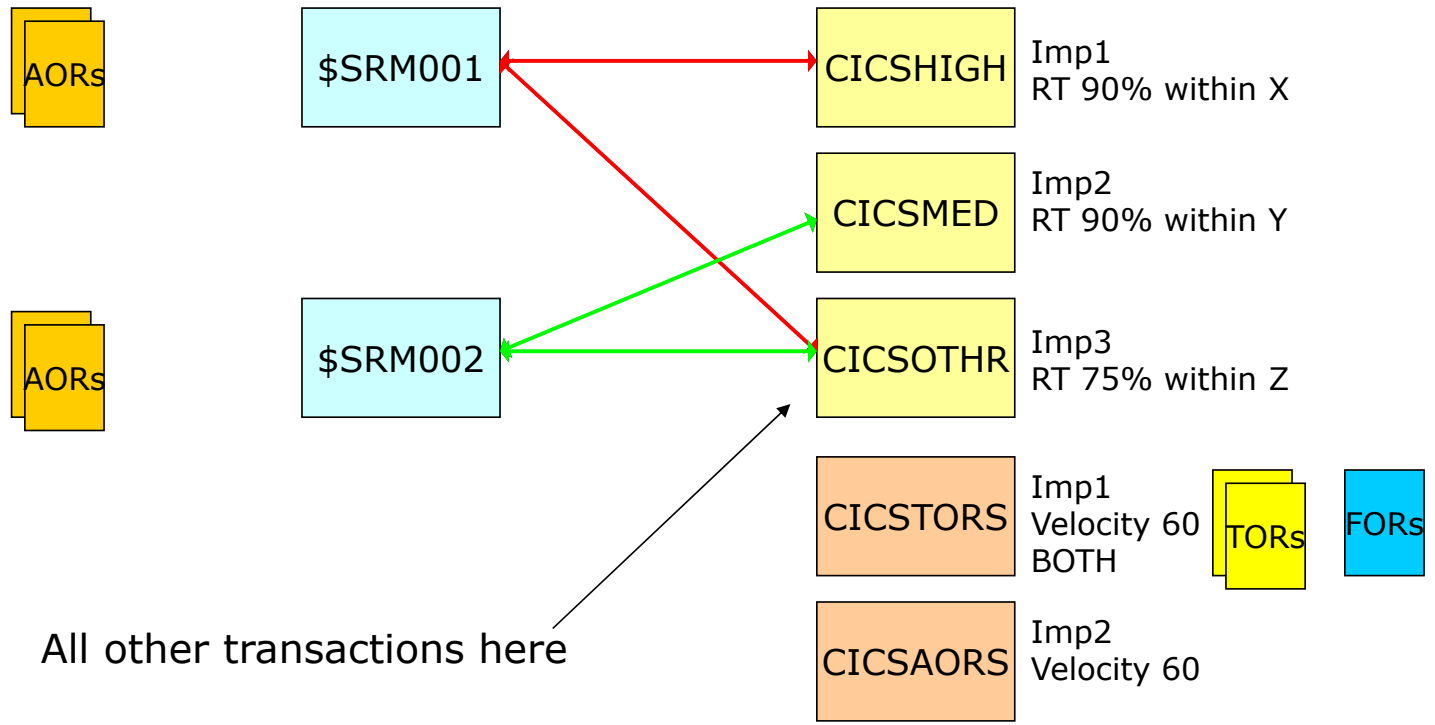

Server Topology – External Period View

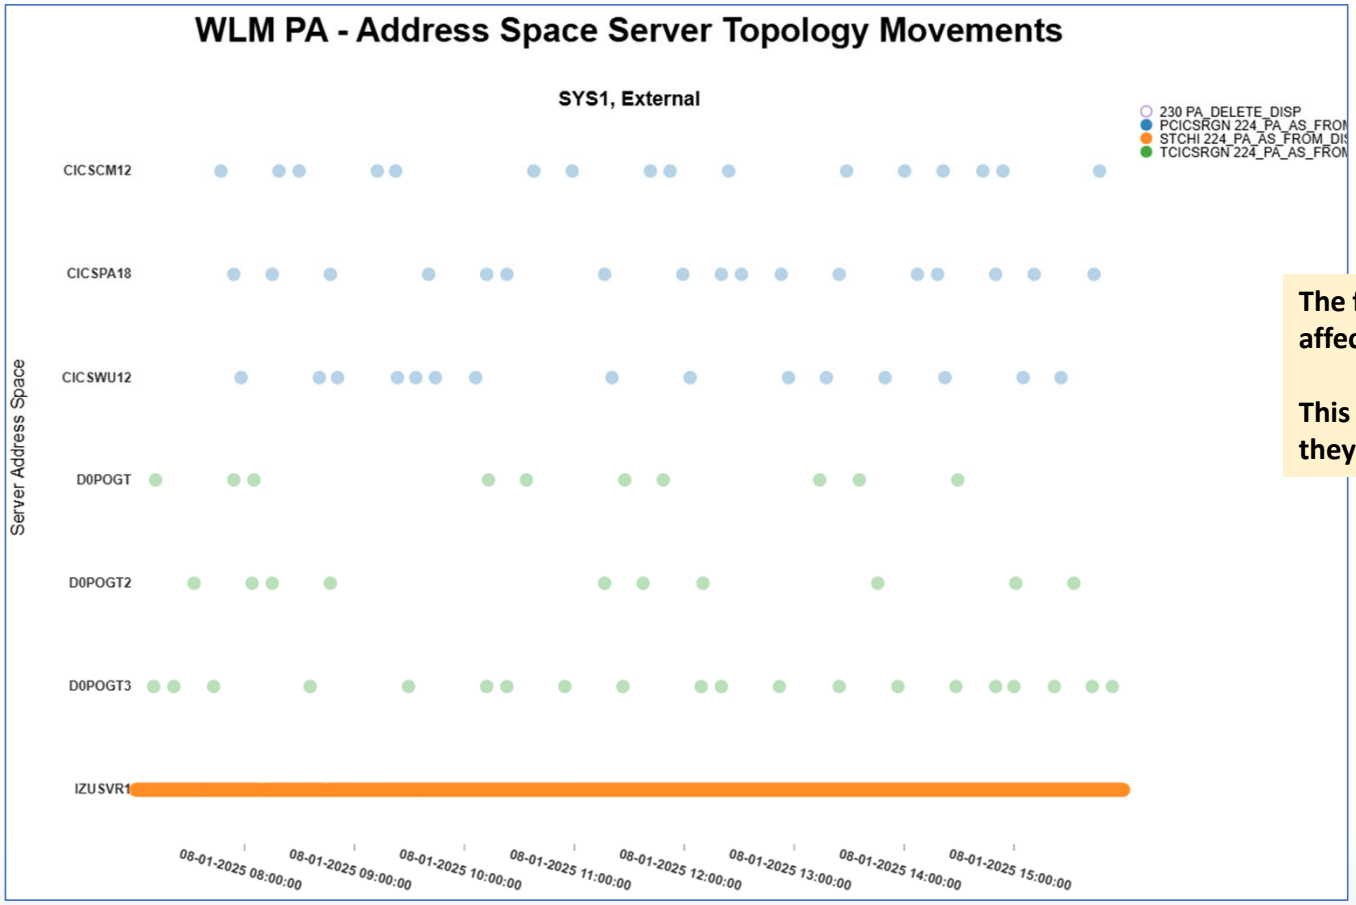

The following chart shows address spaces that were affected by the WLM server topology.

This show address spaces and which external periods they were related to.

Server Topology – Internal Period View

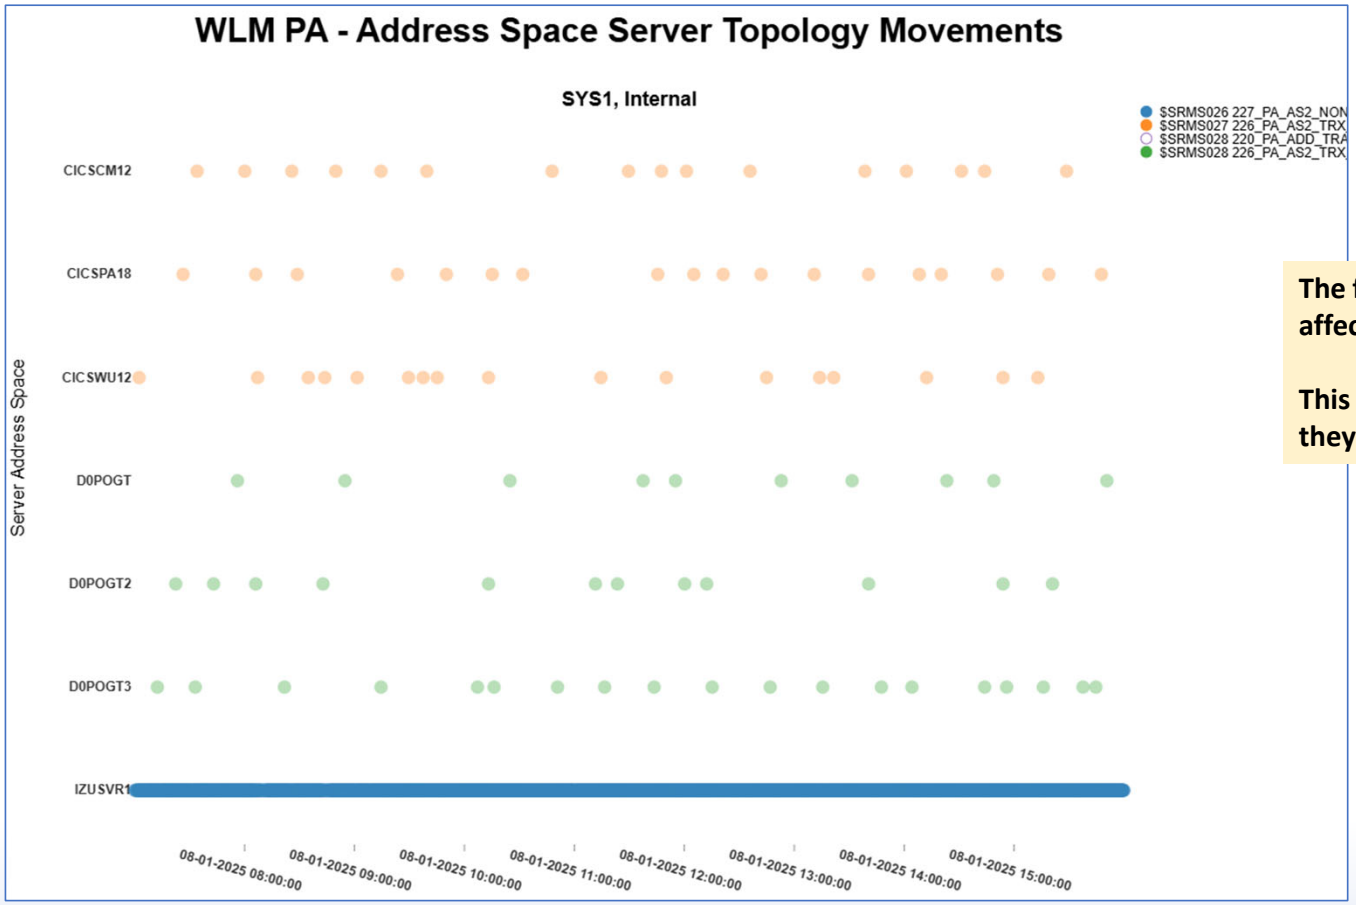

The following chart shows address spaces that were affected by the WLM server topology.

This show address spaces and which internal periods they were related to.

Workload Manager Algorithms

- Internally to the z/OS workload manager are a series of algorithms that are used to ensure the installations performance objectives are met

- The SMF 99.1 trace records provide insights into the execution of these algorithm
  - Although note: IBM WLM development has gotten 'forgetful' over time to provide important trace information with new releases, and they do not always update the documentation as needed

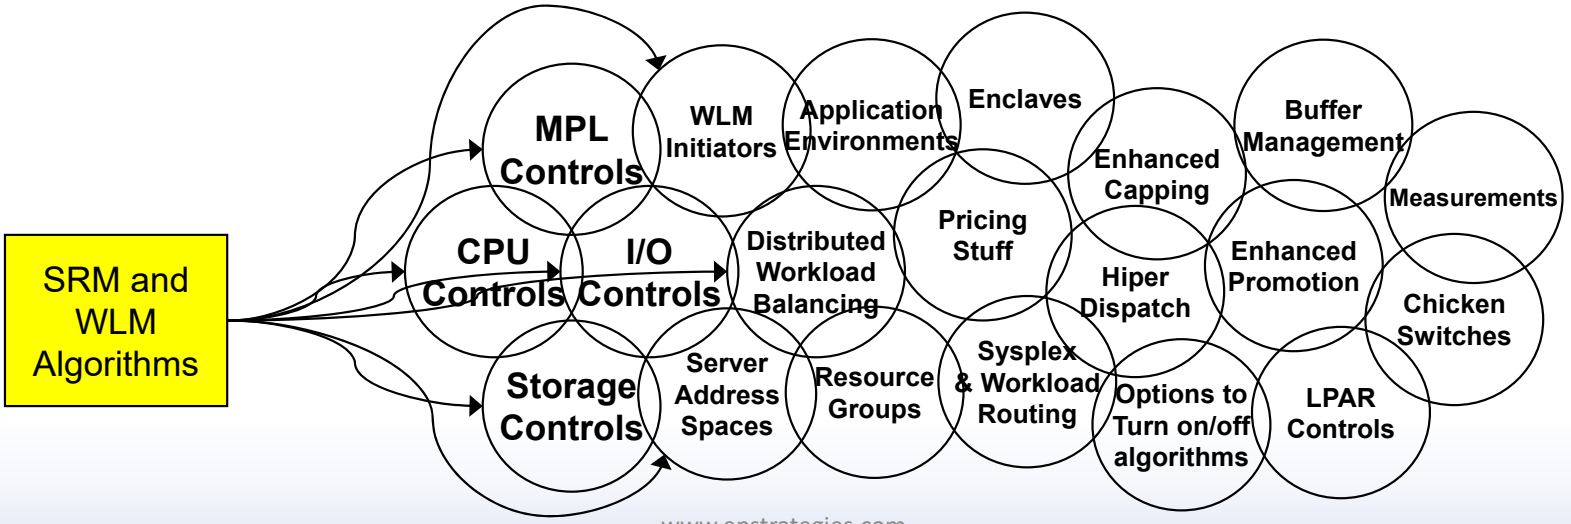

Questions?