Introduction to z Processor Measurements

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

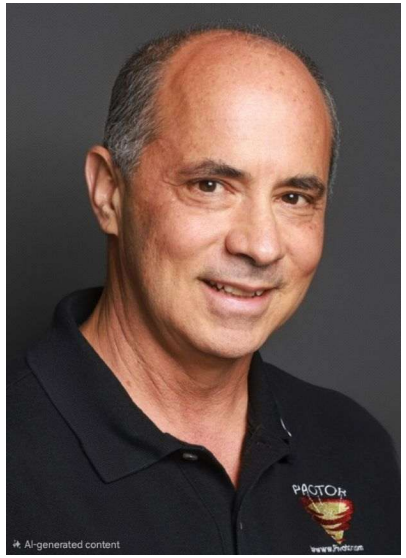

z/OS Performance Education, Software, and Managed Service Providers

Creators of Pivotor®

Contact, Copyright, and Trademark Notices

Questions?

Send email to Peter at Peter.Enrico@EPStrategies.com, or visit our website at http://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

Introduction to z Processor Measurements

On the z platform, there are more processor measurements than any other computing platform. What are all these processor measurements measuring, where do they come from, and how can they be used?

During this presentation, **Peter Enrico** will introduce to you the many processor measurements available and show you how many of them can be used during any processor performance analysis or capacity planning exercise.

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!

- Essential z/OS Performance Tuning
  - March 30 – April 3, 2026 (4 days, excl Wednesday the 1st)
- WLM Performance and Re-evaluating Goals
  - July 27 - 31, 2026 (4 days, excl Wednesday the 29th)
- 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 Latent Demand 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 to show you the results
  - See also: http://pivotor.com/cursoryReview.html

Overview of Processor Measurements on z

- z/OS Processor Measurements
  - Compared to any other computing platform, zArchitecture and z/OS have vastly more processor related measurements than any other platform
    (So many, in fact, that there is not even a worthwhile comparison)

- We can think of the data sources of measurements as follows:
  - **Physical/hardware level** – measuring hardware / machine / CEC view of processor
  - **Operating system level** – measuring the z/OS systems view of processor
  - **Workload/policy level** – measuring workload groupings
  - **Address-space/job level** – measuring interactive, batch, and long-term address spaces
  - **Transaction / Program level** – measuring online trans programs, database transactions
  - **Sampling-frequency/diagnostic level** – fine granularity of any of the above

Pivotor Processor Analysis Reportsets

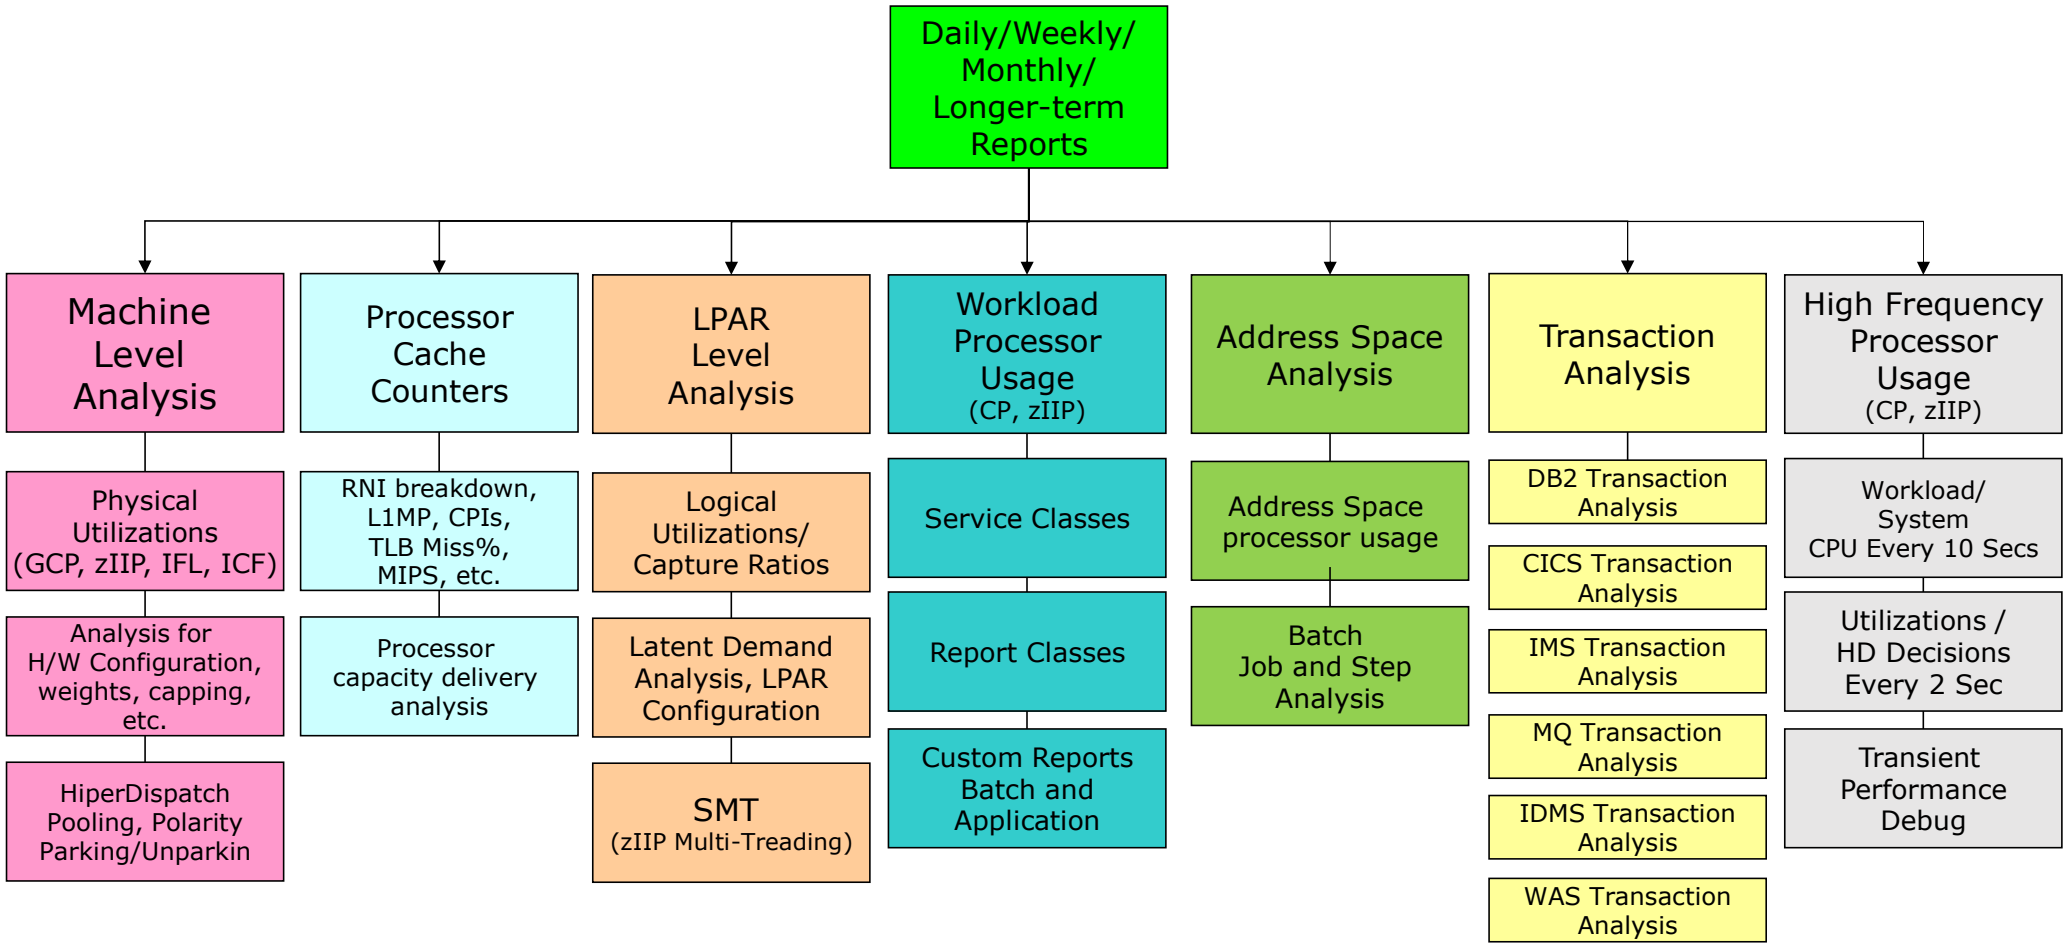

Physical/hardware level

Measuring hardware / machine / CEC view of processor

PR/SM Machine Level
- Measuring processor at the CEC Level

- Many usages of machine / hardware level processor measurements
  - Includes…
    - Delivered capacity at the CEC level
    - Hardware capacity planning to help right-size while planning a processor upgrade
    - Chargeback / MSU usage across LPARs
    - Evaluating LPAR processor usage at the physical machine level
    - Capping
    - Evaluating LPAR configuration settings such as weights and capping controls
    - Etc..

- SMF 70 – Processor Activity record
  - Most common SMF records for measuring the physical processor usage
    - Specifically, PR/SM (hypervisor) view of the entire CEC/machine
    - Source is the Diagnose x’204’ instruction

Breakdown of General Purpose Processor

- We always needed to understand the break down of CP CPU consumption

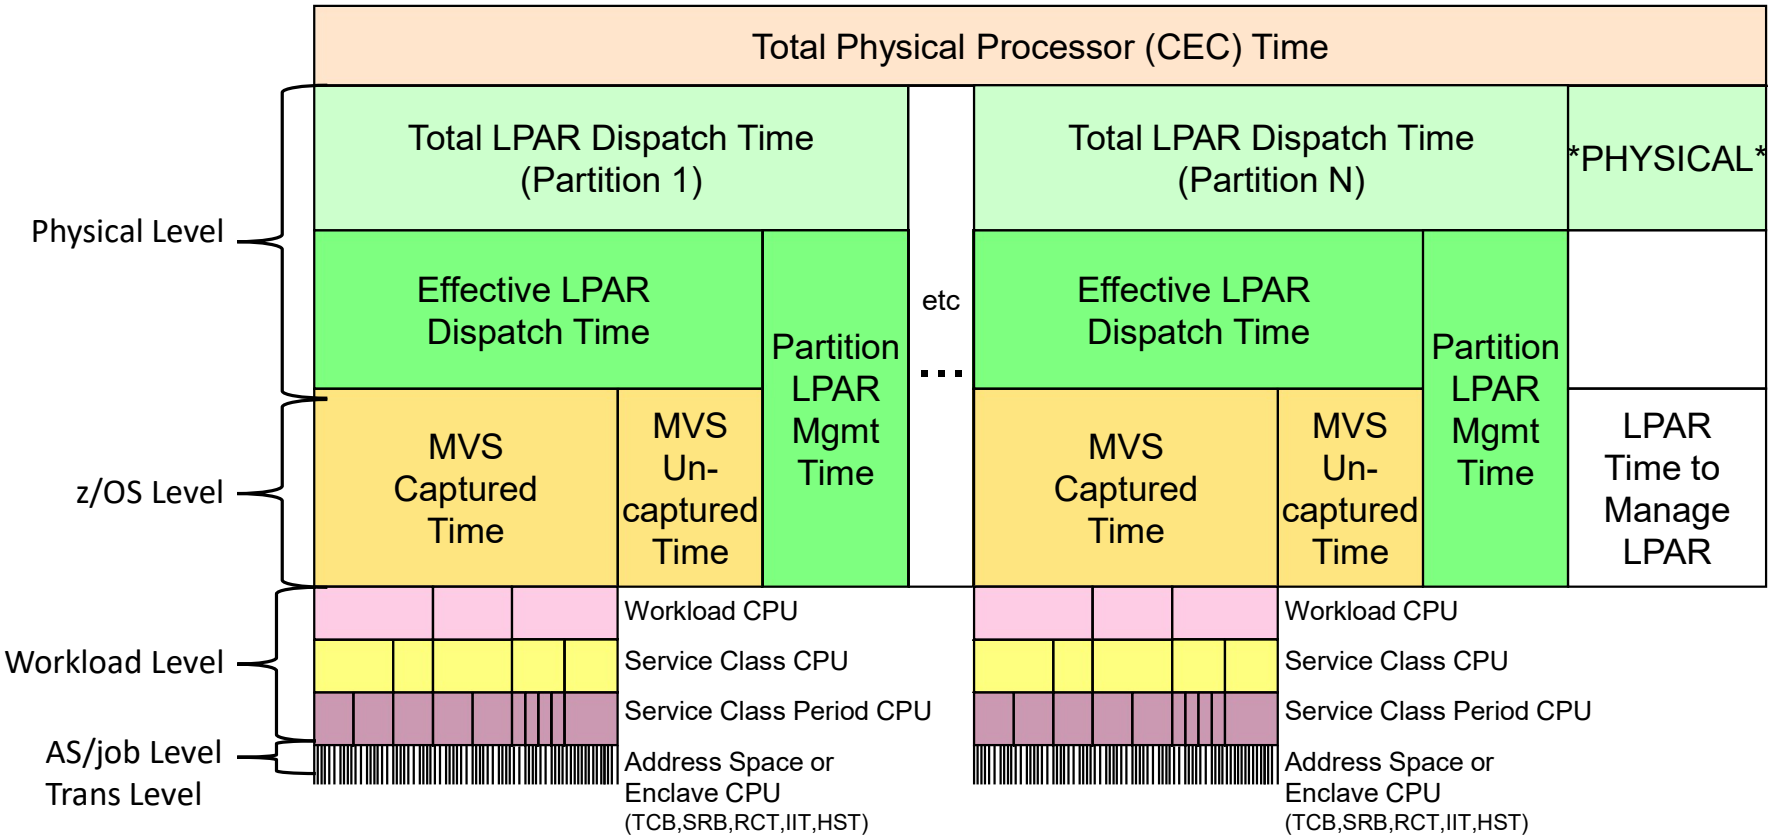

Breakdown of zIIP Engine Time

- We need to understand how PR/SM allocates the zIIP processor resource
  - In all measurements zIIPs

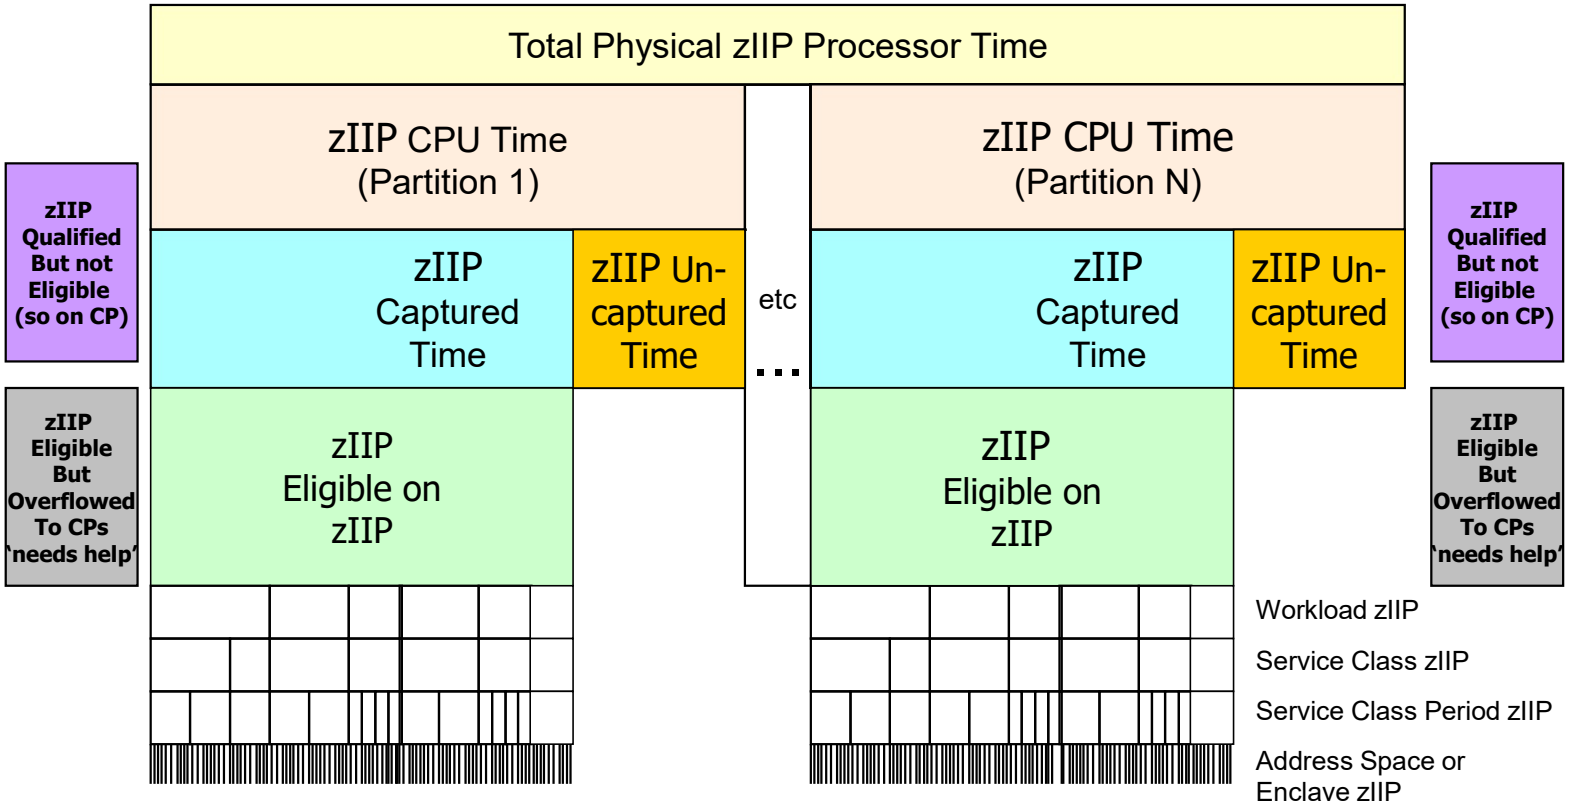

Lots of Possible reports

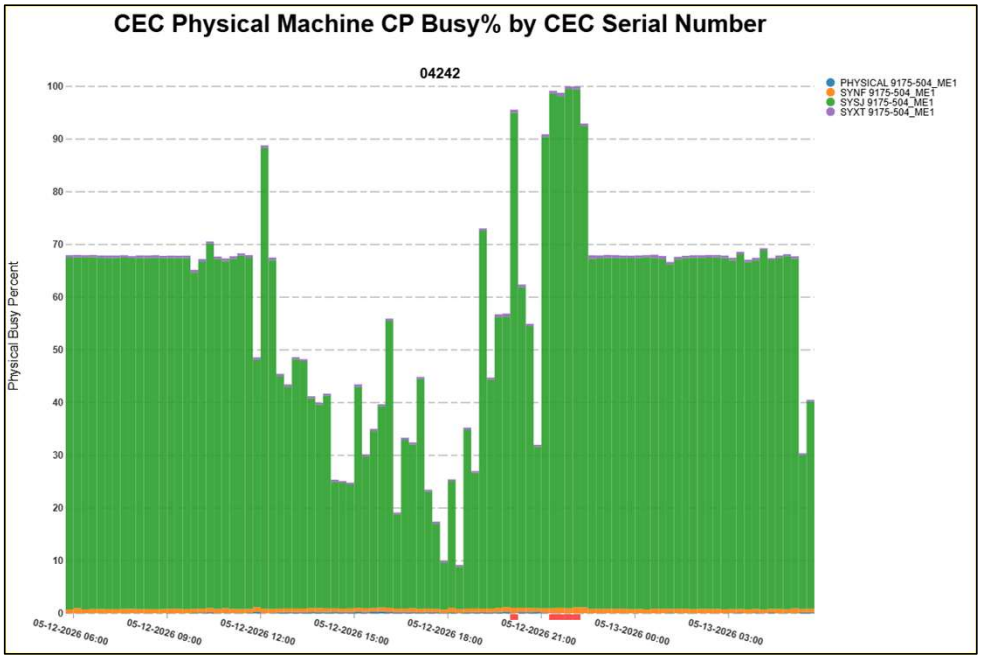

- Note the CPU utilization patterns on this chart. The center of chart is a typical pattern of CPU utilization The left and right-hand side of the chart shows plateaus.

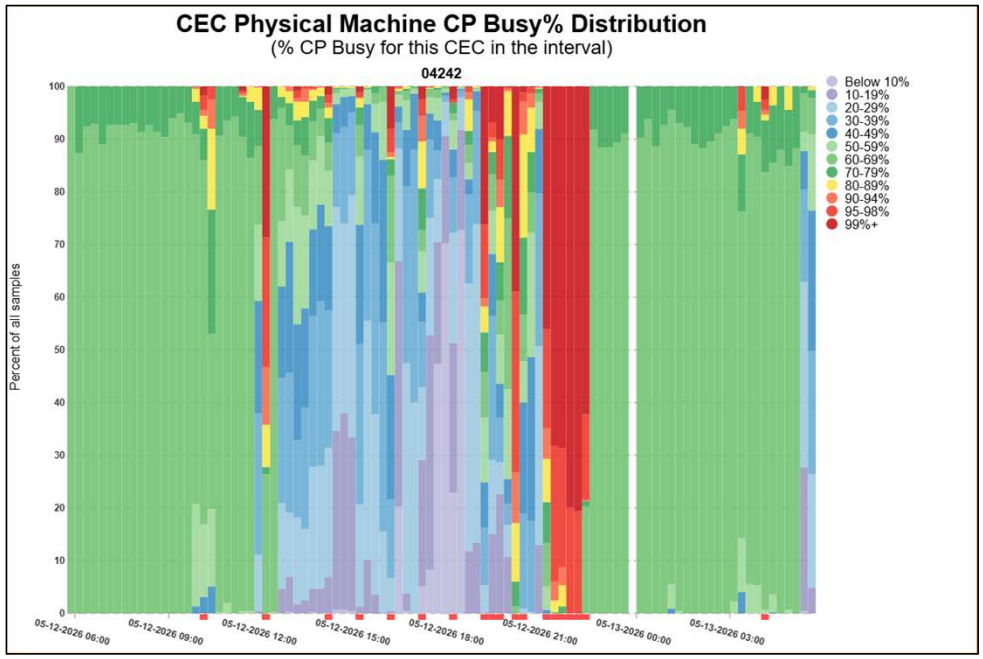

- This chart show what percentage of the measurement interval the utilization of the CEC was a certain percentage. In this case, note only the period-of-time the CEC was near 100% busy.

Physical Processor Utilizations

- Physical utilizations
  - Helps to understand the utilization of the constraint due to the number of physical processors active on the machine

- Physical Processor Utilization Effective
  - Percentage of the measurement interval that the partition was utilizing a physical processor on behalf of itself
  - Online time is related to the interval time. A single CPU cannot be online longer than measurement interval

$$\frac{\sum \text{Partition Effective Dispatch Times}}{\text{No of Physcial Processors} * \text{Online Time}} * 100$$

- Physical Processor Utilization Total
  - Percentage of the measurement interval that the partition was utilizing a physical processor on behalf of itself and for LPAR management time attributed to the partition
  - Online time is related to the interval time. A single CPU cannot be online longer than measurement interval

$$\frac{\sum \text{Partition Total Dispatch Times}}{\text{No of Physical Processors} * \text{Online Time}} * 100$$

Lots of Possible reports

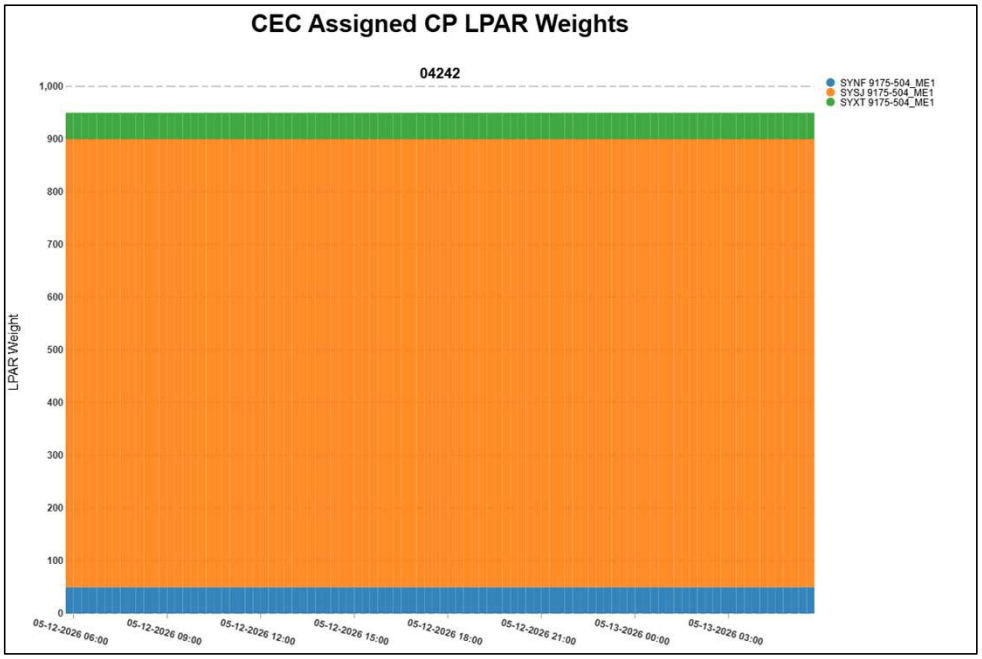

- This chart shows the PR/SM weights assigned to each LPAR to allow PR/SM to assign a guaranteed share of the processor resource.

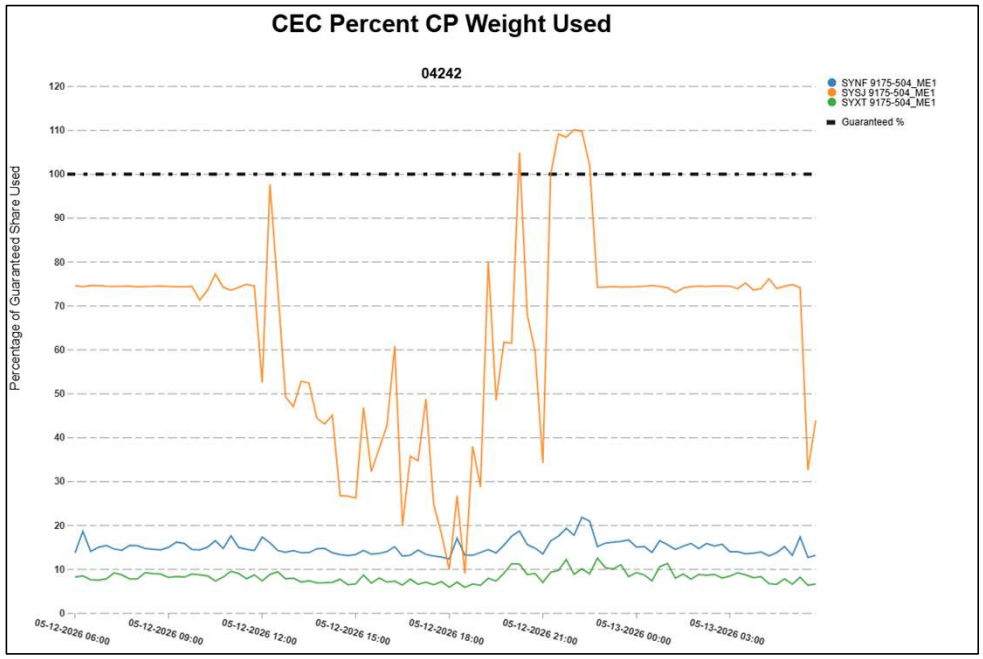

- This chart show what percentage of each LPAR's guaranteed share that was consumed by each LPAR.

Capping and Weight Enforcement affect HiperDispatch Pooling

- When capping is enforced, it is very possible that HiperDispatch changes the pooling of the CPUs
  - In this example, 2 high and 1 mediums turn into 1 highs, 2 mediums
  - There are no low pool processors, so all are un-parked.
  - This matters since high pool CPUs tend to have longer dispatch intervals than medium pool CPUs

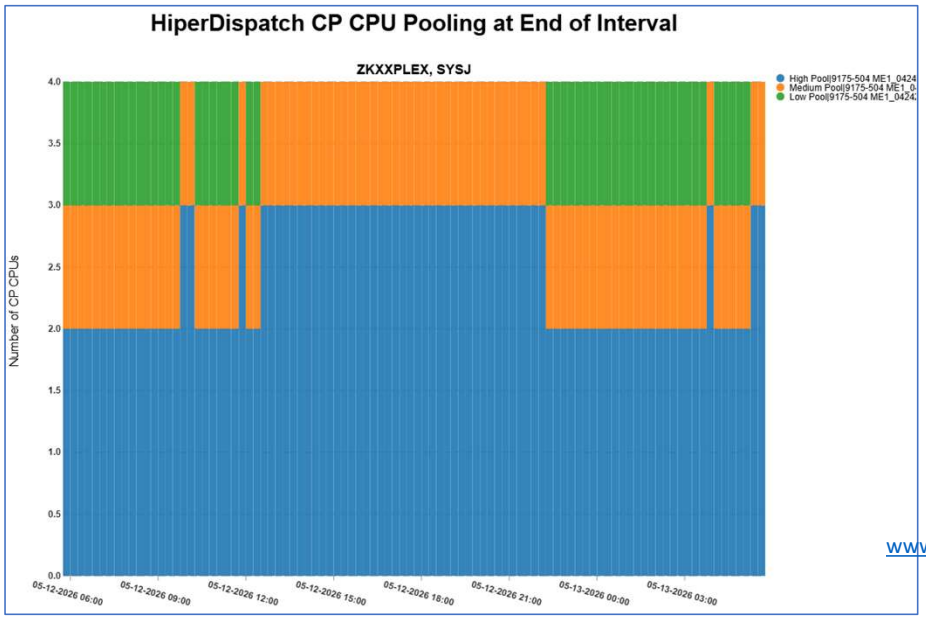

www.pivotor.com
charts

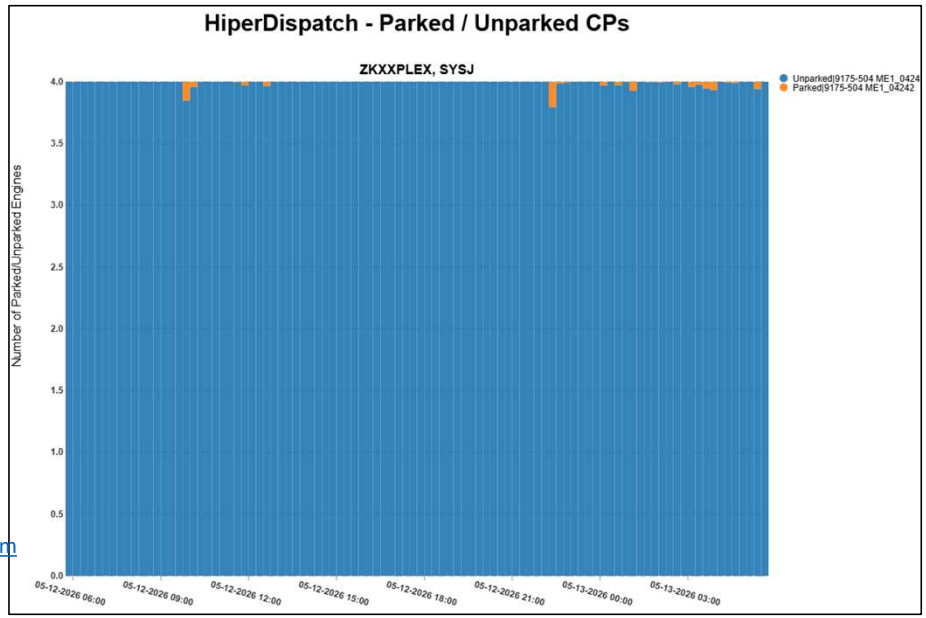

Processor Caches
– Measuring processor at the cache level

- Many usages of processor cache measurement
  - Includes
    - Evaluation of processor memory hierarchy
    - Software efficiency analysis — is this workload cache-friendly?
    - LSPR-style capacity-sizing inputs to help right-size while planning a processor upgrade
    - Evaluating the cost of Dynamic Address Space Translations by examining TLB Miss CPU%
    - Evaluating positive / negative changes to hardware changes by examining a variety of counts
    - Etc..

- SMF 113 – Hardware Instrumentation
  - The SMF record for measuring the physical processor cache usage
  - Collected by the Hardware Instrumentation Services (HIS) started task
  - Source are cache counts directly from the hardware (Basic, Problem, Crypto Extended)

COUNTER SET= BASIC / PROBLEM-STATE

- Basic and Extended counters are the most interesting
- Most important derived values from the measurements include
  - L1MP – Level 1 cache misses per 100 instructions
  - CPI – Cycles per Instruction
  - Relative Nest Intensity
  - TLB CPU Miss Percent of CPU
  - CPU speeds and normalization factors

- Lesser used value (from this record)include:
  - Actual MIPS
  - CPU utilization
  - Cache sourcing at the instruction and data levels

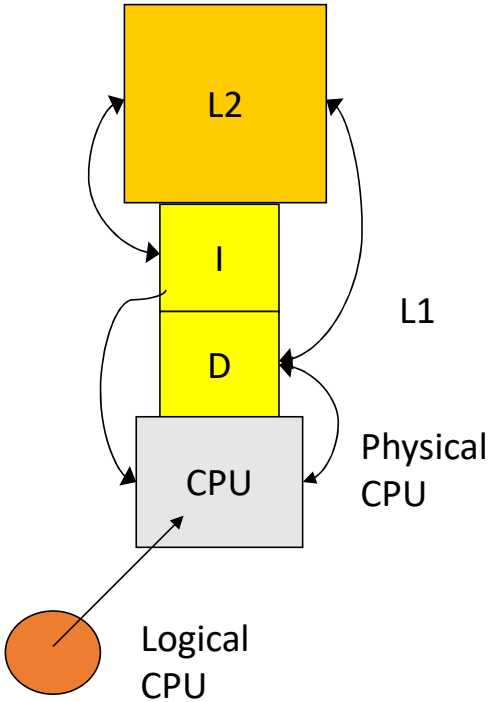

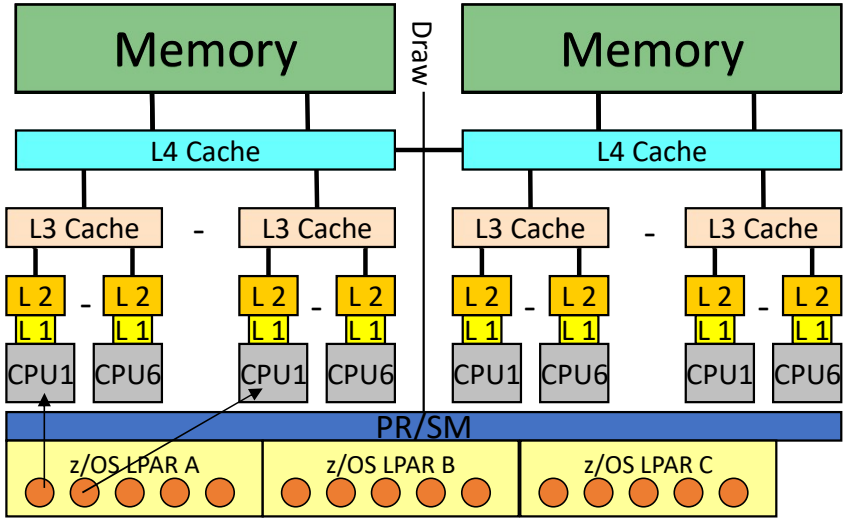

There are 4 key cache values (Pivotor Example)

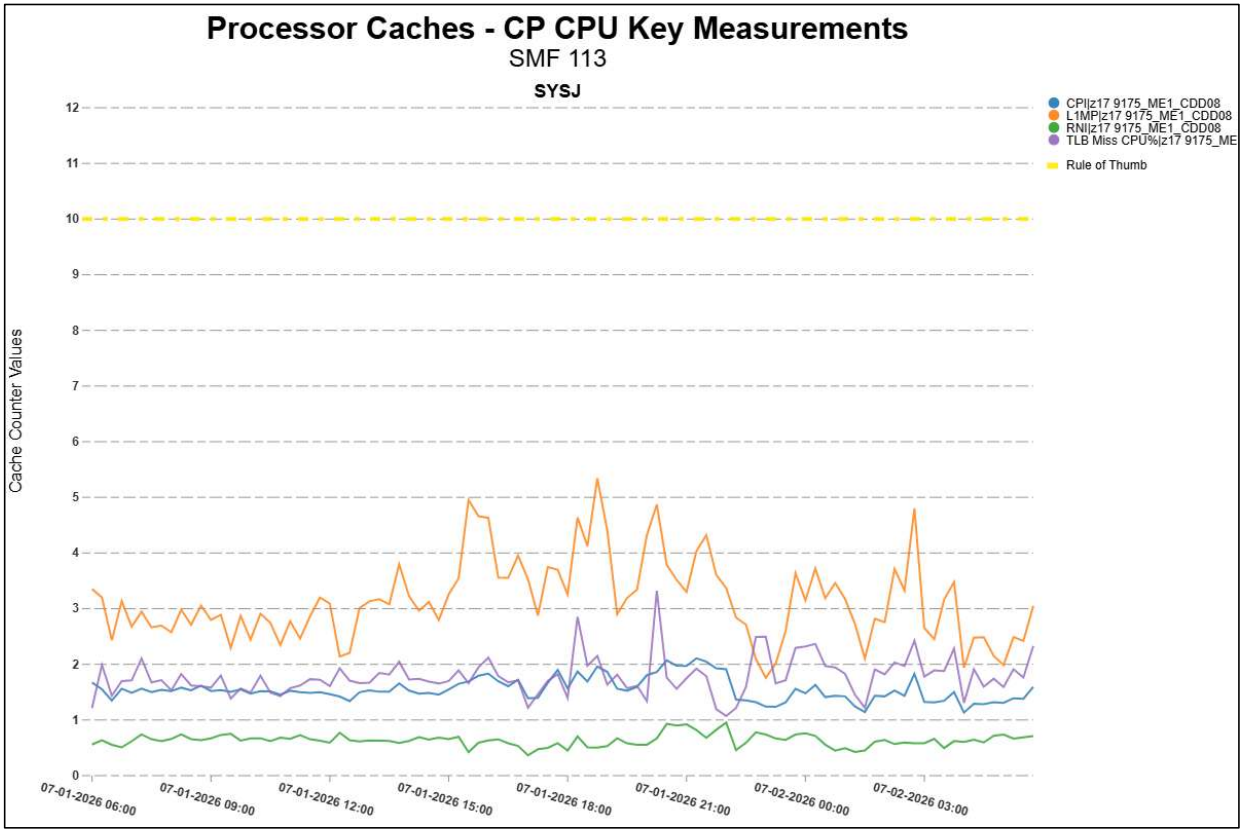

4 key calculated (CP CPU example):

- CPI: Cycles per Instruction
  - Used to gauge processor contention, as well as instruction mix consistency
- L1MP: L1 misses per 100 instructions
  - Think of this as a 'miss percentage'
- RNI: Relative Next Intensity
  - Workload 'signature' that gauges the pressures being put on the upper caches and memory
- TLB Miss CPU Percentage
  - Total percent of the CPU consumed by the LPAR that goes to dynamic address translation (DAT) due to a translation look-aside buffer miss

Lots of Possible Reports

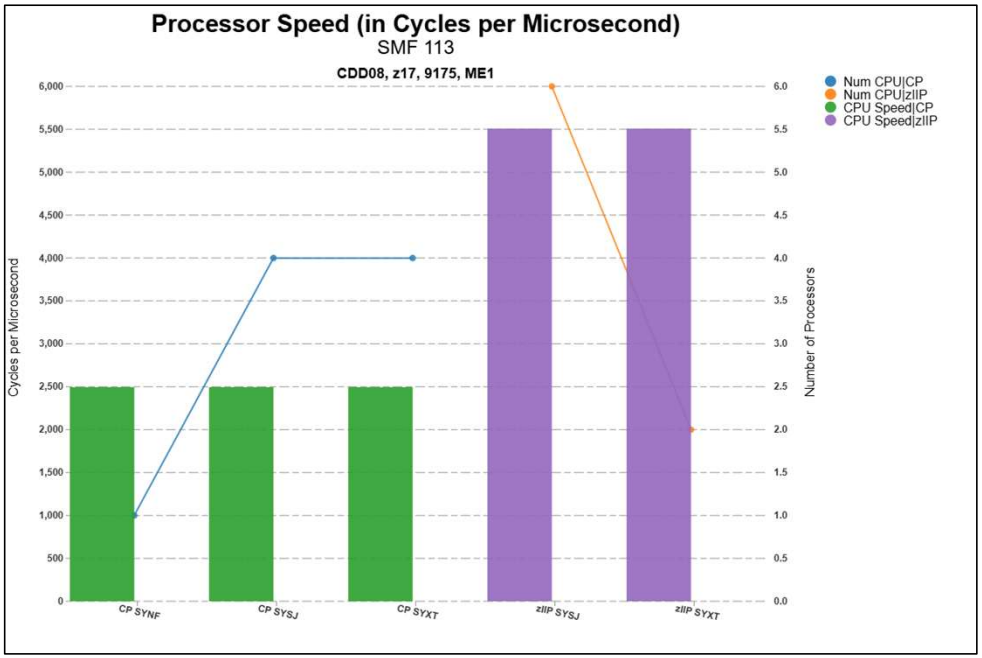

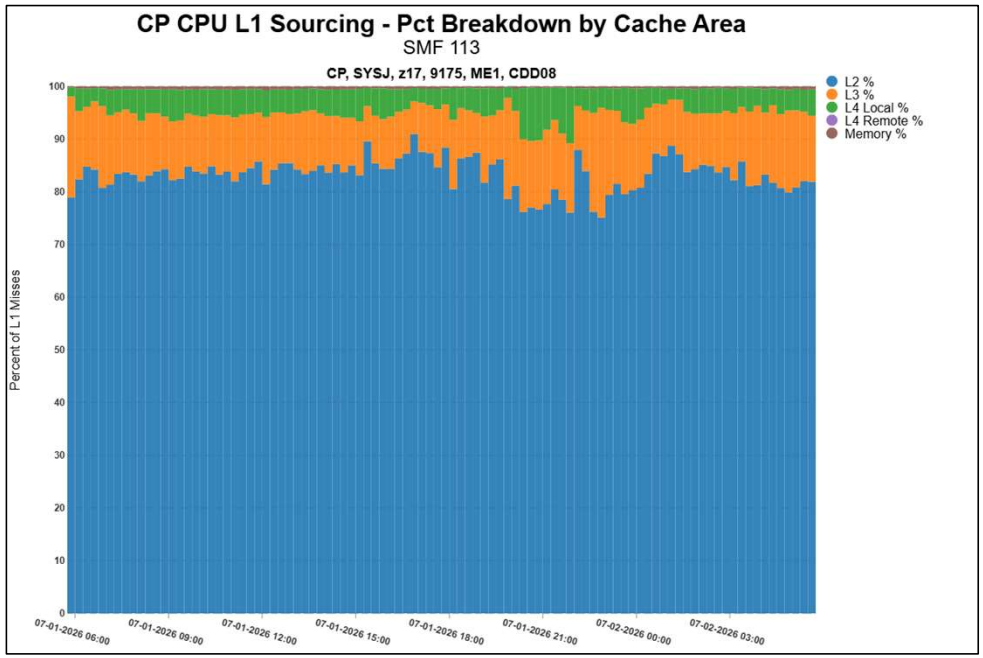

Operating system level

Measuring the z/OS systems view of processor

z/OS System – Measuring processor at z/OS System Level

- Many usages of z/OS operating system level processor measurements
  - z/OS image's own accounting of its dispatched CPU time independent of what the hypervisor promised it.
  - Includes
    - Delivered capacity at the LPAR level – LPAR Busy %, MVS Busy %, Capture Ratios
      - Verifying the z/OS processor PR/SM is providing the system
    - Understanding effects of the logical constraints placed on the LPAR
    - Chargeback / MSU usage for an LPAR
    - Evaluating latent demand and work unit queues
    - Certain processor performance parameters and algorithms (such as CPENABLE)
    - Boosting
    - Etc…

- SMF 70 – Processor Activity record
  - Most common SMF records for measuring the physical processor usage
    - Specifically, some PR/SM (hypervisor) view and z/OS operating system measurement
    - Both Diagnose x'204' instruction for PR/SM inputs and z/OS operating system for software inputs

LPAR Busy % vs MVS Busy %

- LPAR Busy % measures how busy the LPAR kept its logical processors
- MVS Busy % measures how much of the logical resource the LPAR wanted
- Differences indicate latent demand
  - Flat lines usually indicate capping or weight enforcement

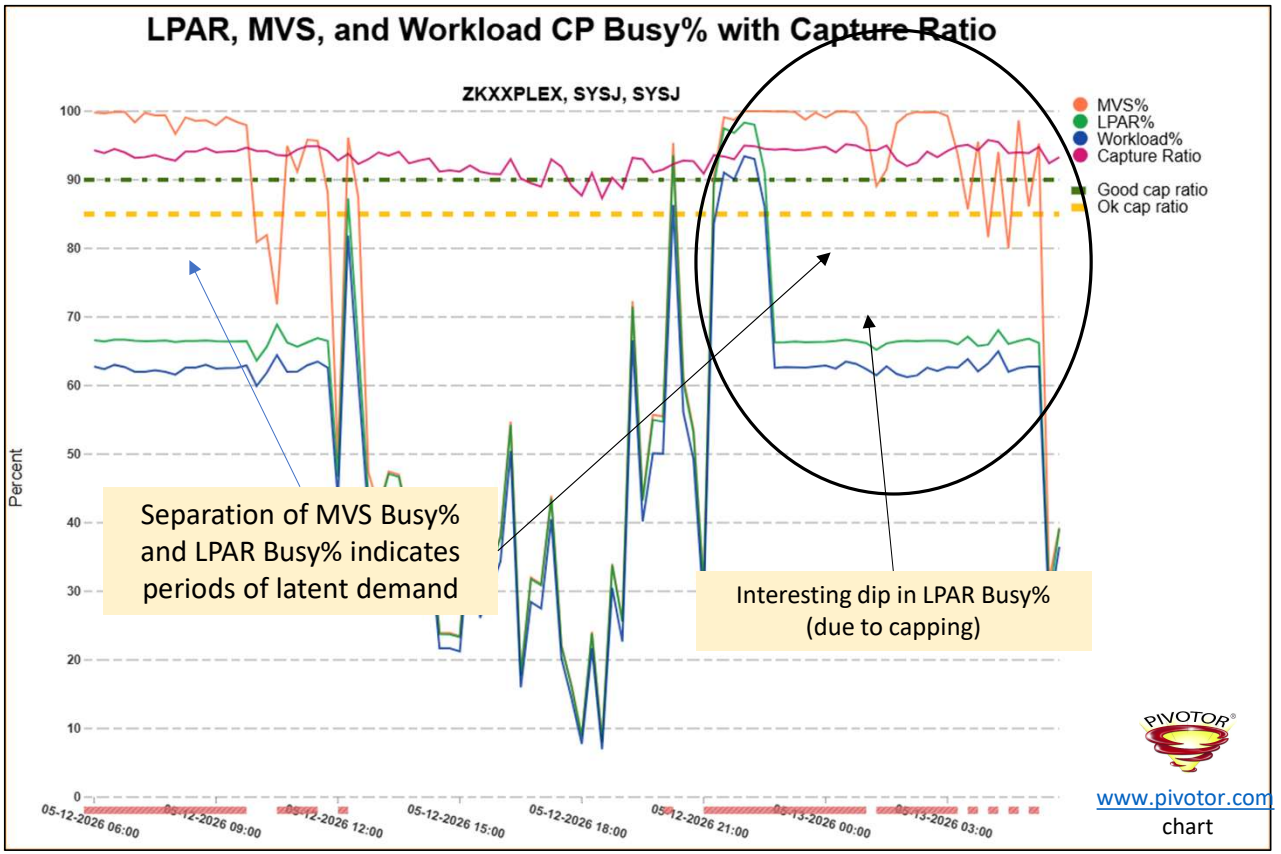

Logical Processor Utilizations

- Logical utilizations
  - Helps to understand the utilization of the constraint due to the number of logical processors assigned to the partition

- Logical Processor Utilization Effective
  - Percentage of the measurement interval that the partition was utilizing a logical processor on behalf of itself

$$\frac{\sum \text{Partition Effective Dispatch Times}}{\text{No of Logical Processors} * \text{Online Time}} * 100$$

- Logical Processor Utilization Total
  - Percentage of the measurement interval that the partition was utilizing a logical processor on behalf of itself and for LPAR management time attributed to the partition

$$\frac{\sum \text{Partition Total Dispatch Times}}{\text{No of Logical Processors} * \text{Online Time}} * 100$$

Understanding Dispatching to Gain Insights to MVS Busy %

- Dispatch Time
  - Time logical processor is associated with a physical processor

- MVS Time
  - Time z/OS was busy before voluntarily giving up a processor

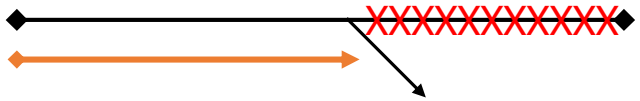

Voluntary Wait
- z/OS voluntarily gives up the processor
- MVS time equals dispatch time

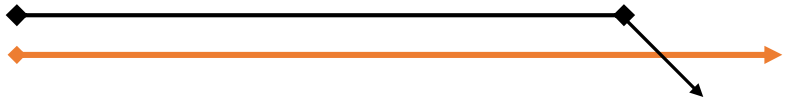

Involuntary Wait (mostly on vertical mediums)
- z/OS does not give up the processor voluntarily
- Instead PR/SM un-dispatches the partition
- MVS time will be greater than dispatch time

Lots of Possible reports

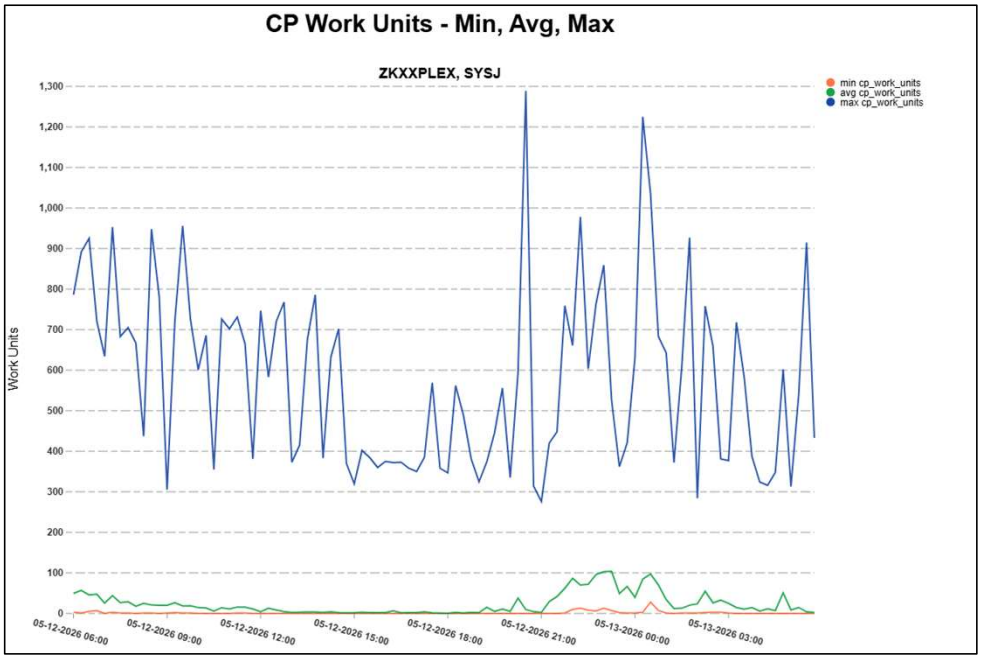

- Min, Max, Average of work unit distribution of work unit queue lengths

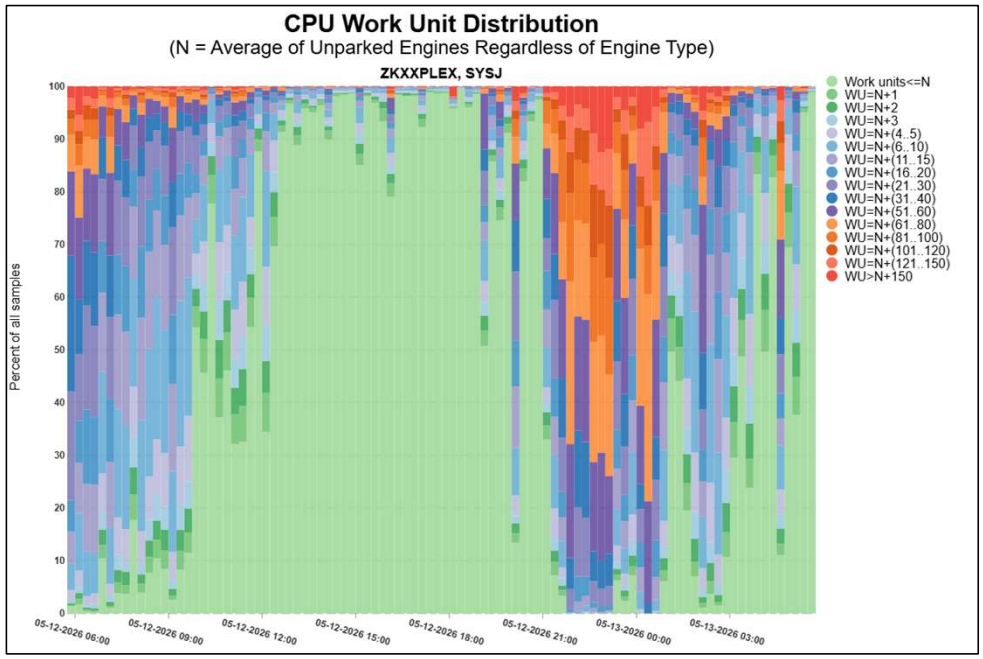

- Each bucket of the distribution represents the percentage of the measurement interval the queue of work waiting to use the CPUs is a certain length

Workload/policy level

Measuring workload groupings

z/OS Workload / WLM – Measuring processor at workload level

- Many usages of z/OS workload processor measurements
  - Includes
    - CPU consumed at the WLM Service Class Period level
    - CPU consumed at the WLM Report Class Period level for more granule reporting
    - Summarization of all workload measurements for a system level point-of-view
    - Both GCP and zIIP processor usage at the TCB, SRB, IIT, HSP, and RCT task levels
    - Processor using and delays samples for input to WLM management
    - Etc…

- SMF 72 – Workload Activity Level
  - Most common SMF records for measuring processor usage at the workload level in the form of WLM Service Class period, Report Class periods, and overall system workload CPU usage

Lots of Possible reports

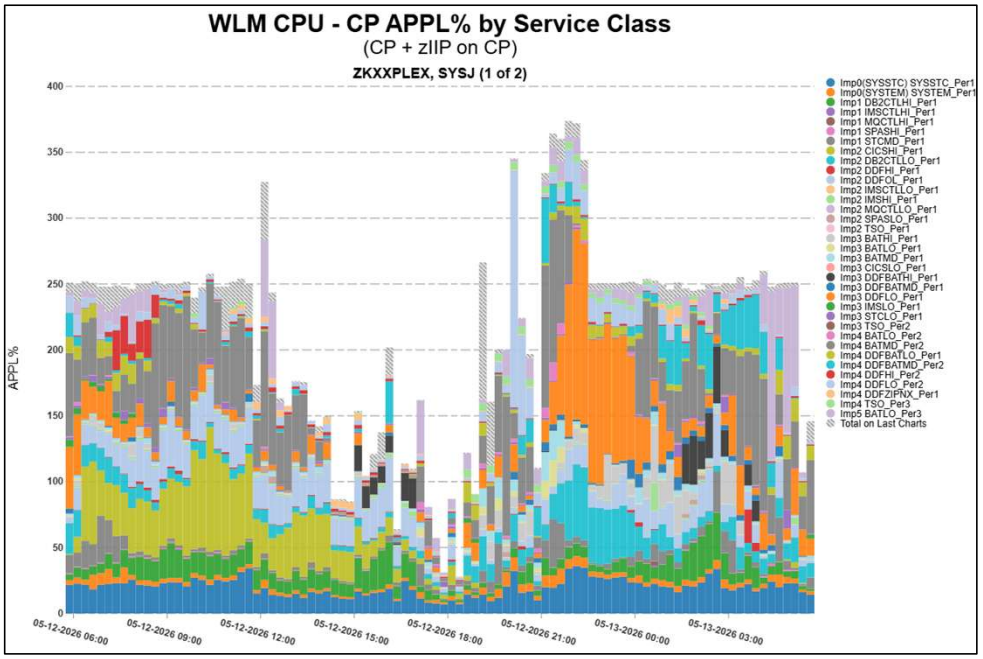

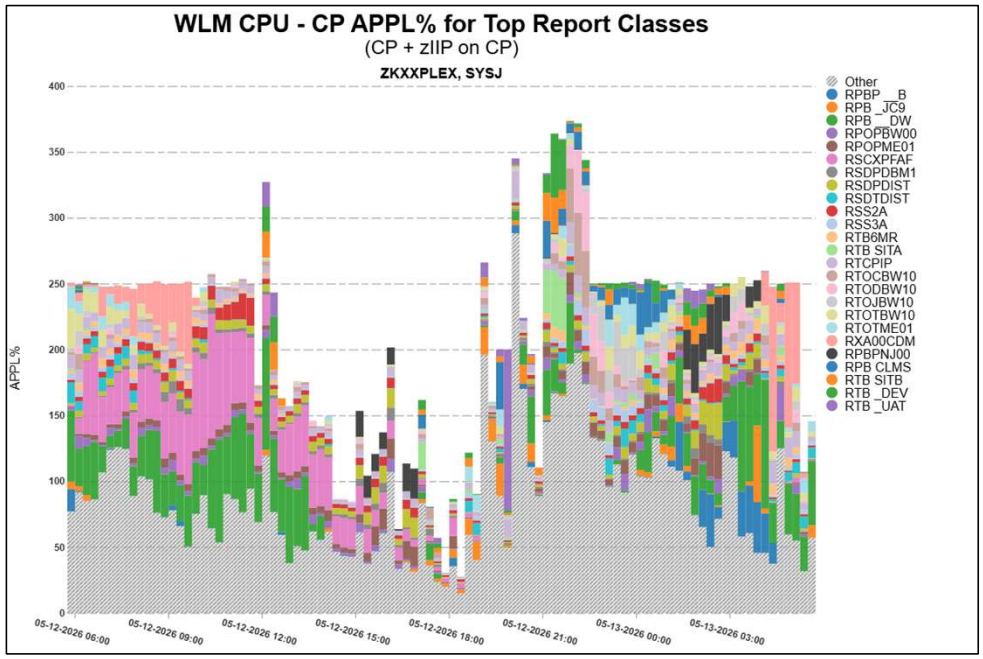

- Reported as a percentage of 1 GCP, this report shows the CPU consumption at the WLM service class period level.
- Reported as a percentage of 1 GCP, this report shows the CPU consumption at the WLM report class level.

Lots of Possible reports

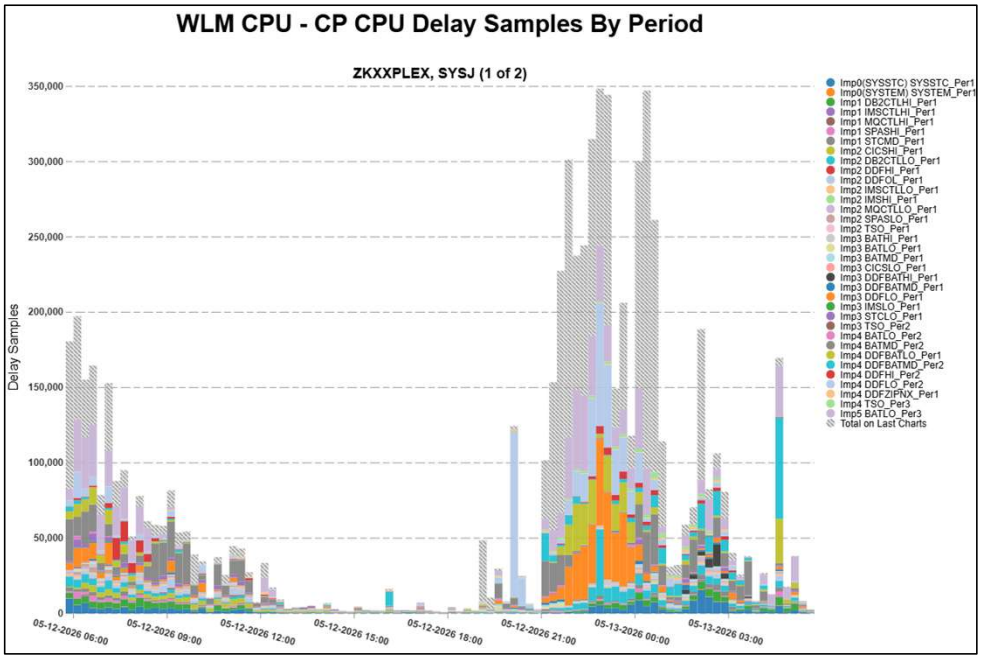

- This report shows the CPU delay samples at the WLM WLM service class period level.

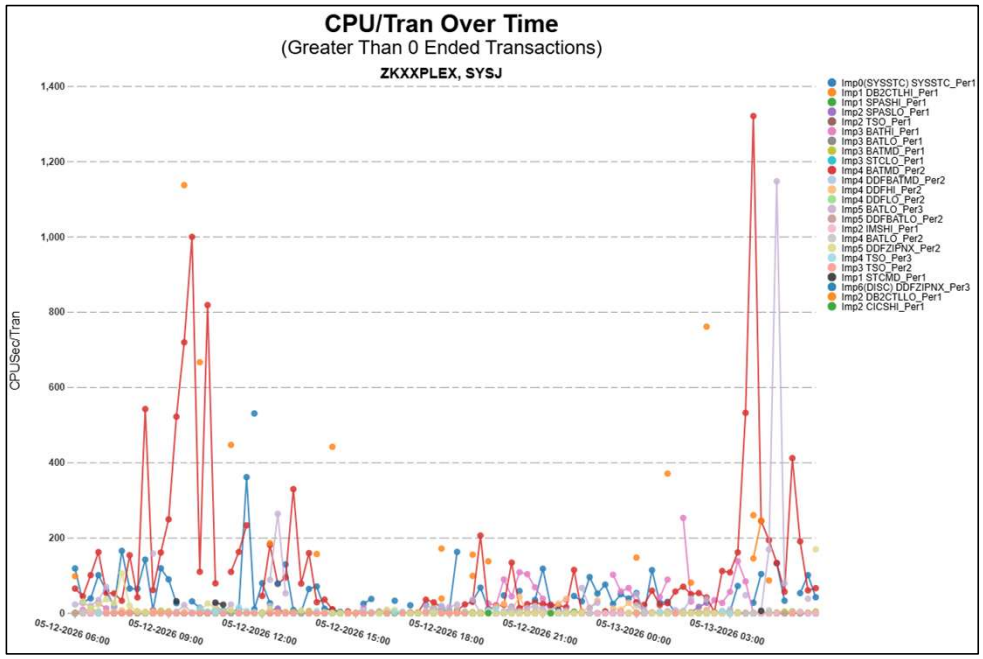

- In the SMF 72 record we know the number of ended transactions for each WLM service class. With this information we can determine the average Internal Throughput Rate of CPU/Tran and Tran/CPU.

Address-space/job level

Measuring interactive, batch, and long-term address spaces

z/OS Address Space and Job – Measuring processor at Address Space Level

- Many usages of z/OS address space processor measurements
  - Includes
    - Evaluating the CPU consumption of each address space, batch job, and batch job step/program
    - Understanding how much CPU is consumed in address space, in enclaves, in other address spaces but charged to this address space, and more
    - Both GCP and zIIP processor usage at the TCB, SRB, IIT, HSP, and RCT task levels
    - CPU intensity values, zIIP offload, etc.
    - Chargeback at the address space and job level
    - Etc…

- SMF 30 – Address Space Level
  - Most common SMF records for measuring processor usage at the address space level
  - Based on options, turned on for each address space, and cut on the interval basis, job end basis, step end basis, and when address space is terminated.

Summary of Dispatchable Unit Types

- Dispatchable units represent the programs that run on the CPUs

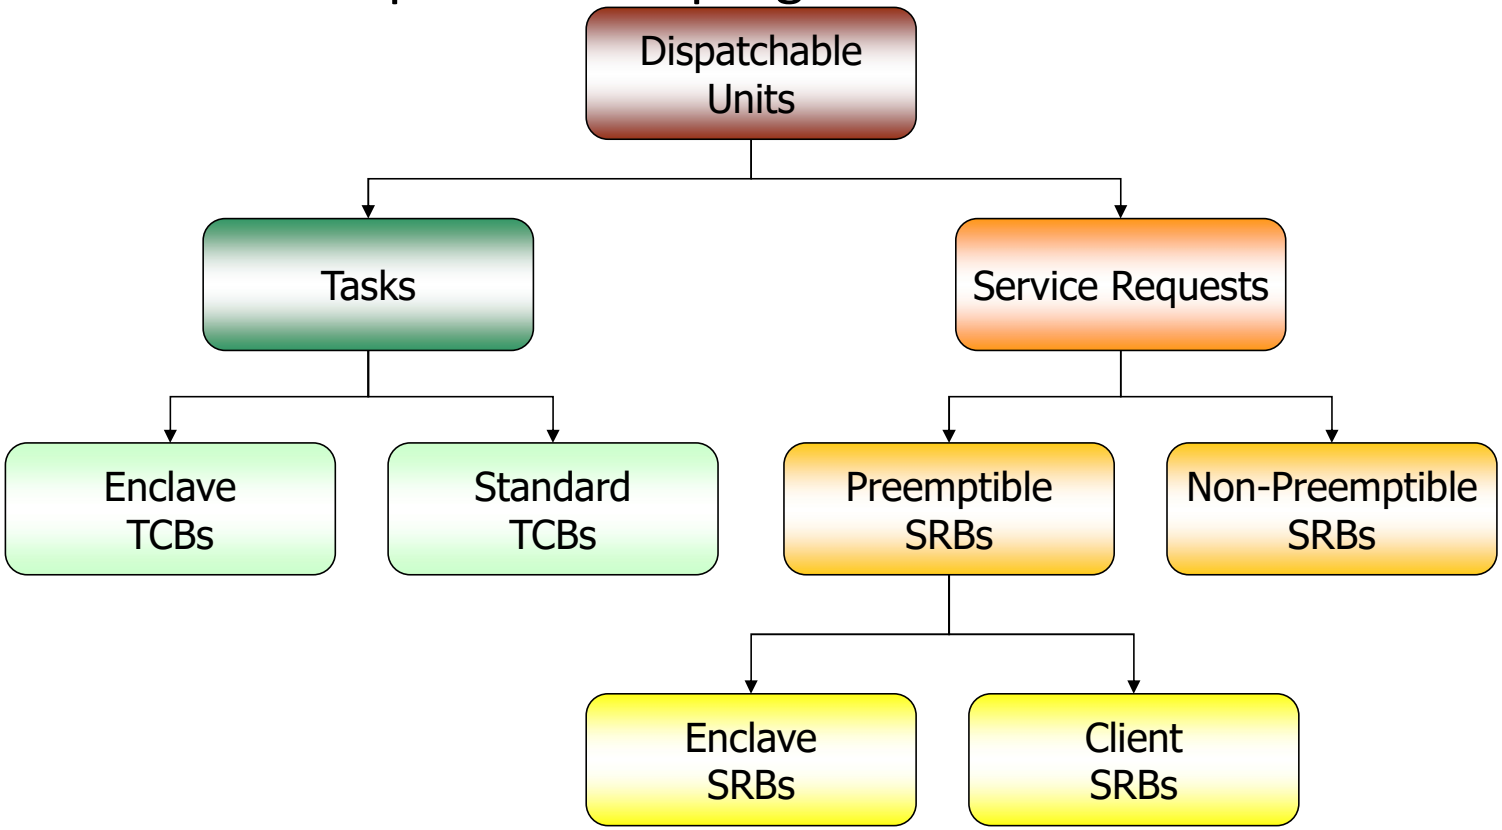

Lots of Possible reports

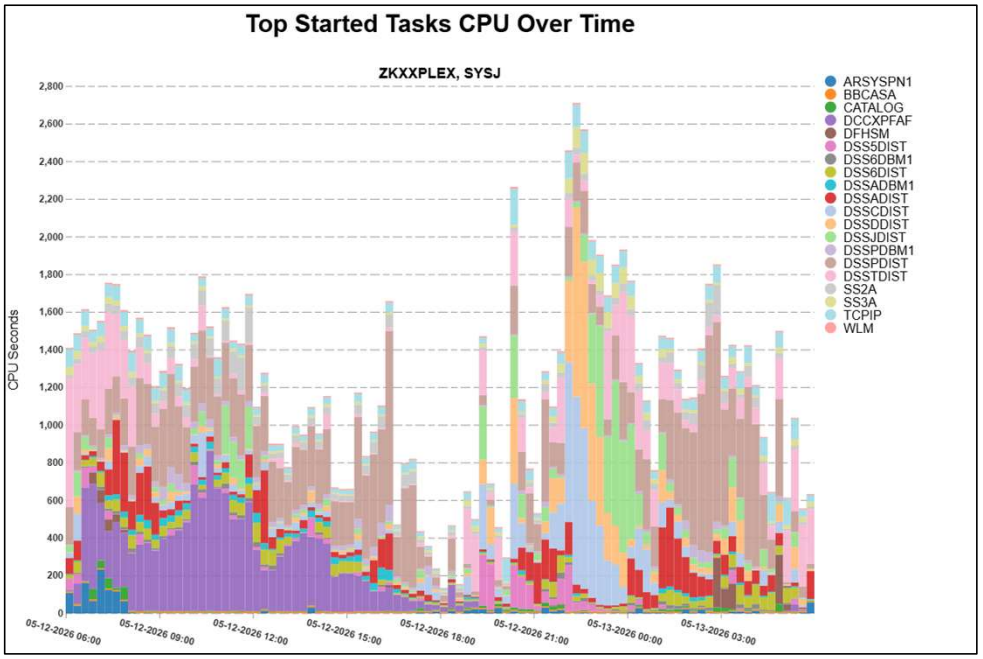

- This report shows the top started task CPU times

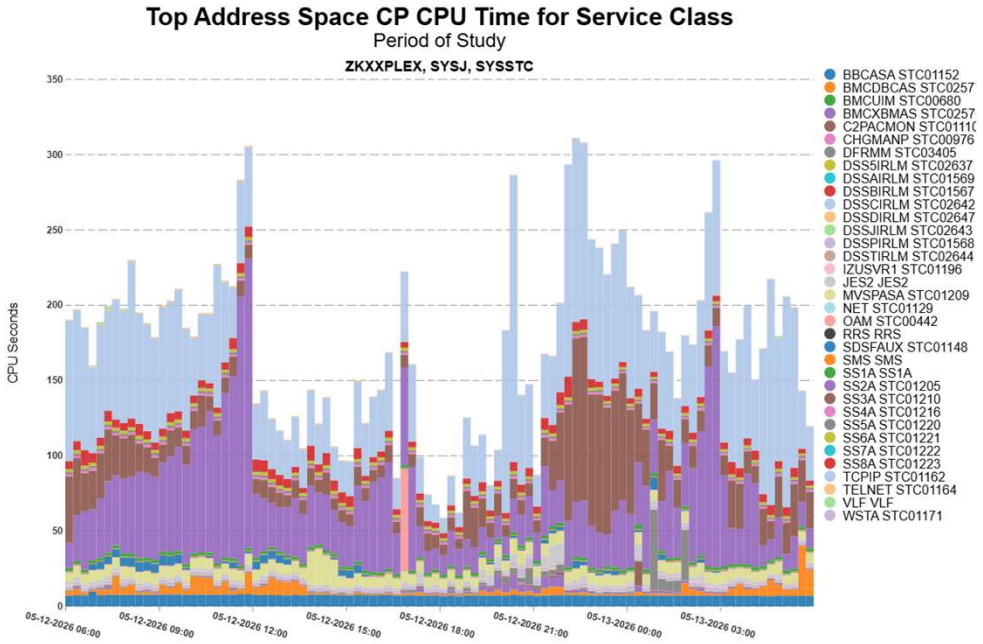

- This report shows the top started task CPU times for just the SYSSTC service class

Lots of Possible reports

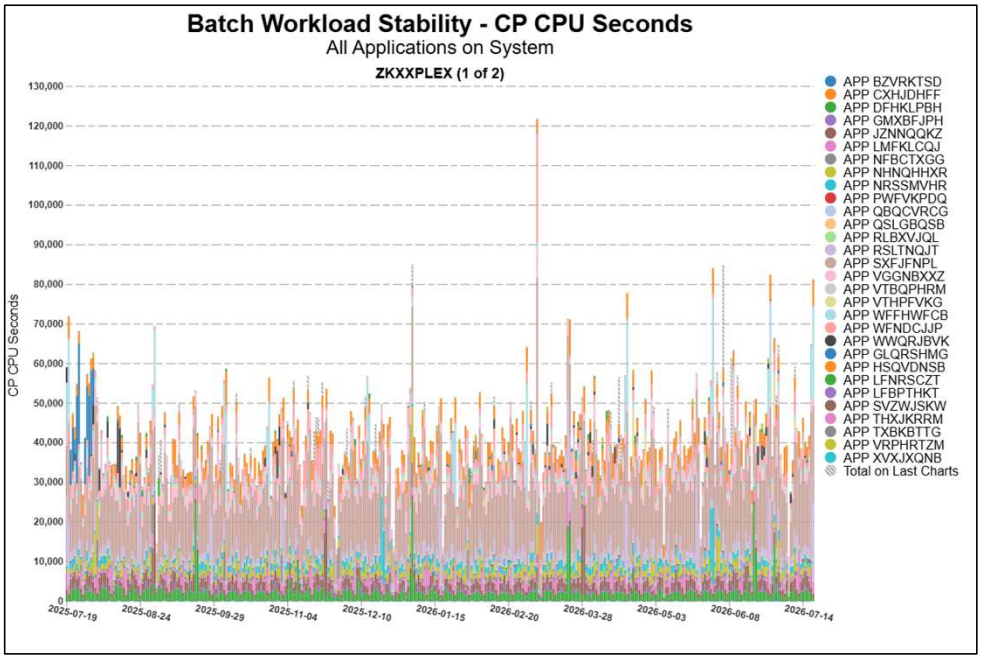

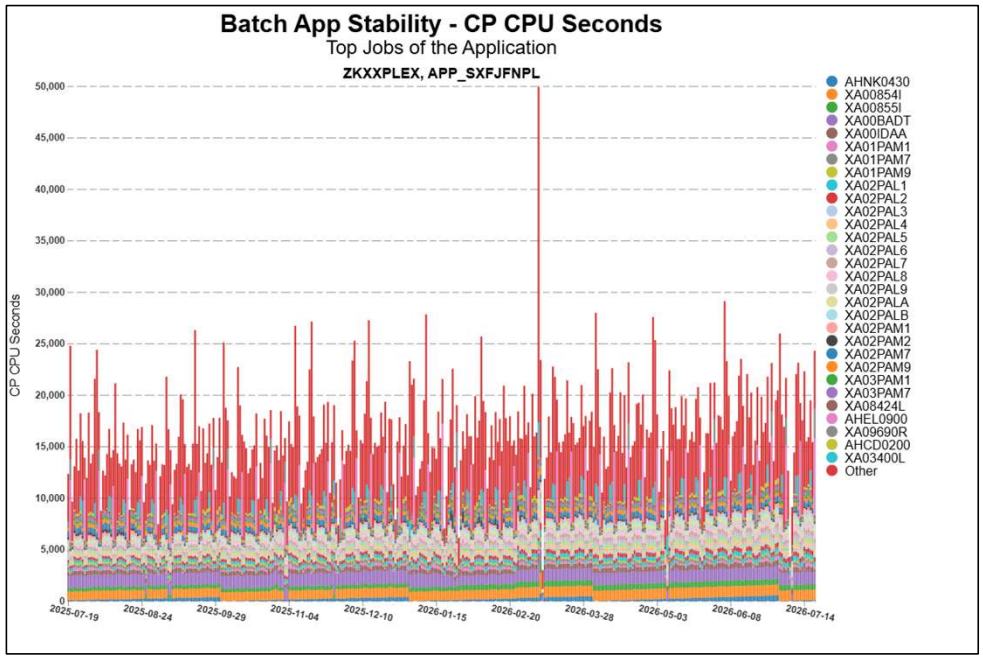

- This report show the CPU time consumed by the top applications over a period of a year

- This report show the CPU time consumed by a specific application and broken down by the top CPU consuming jobs

Lots of Possible reports

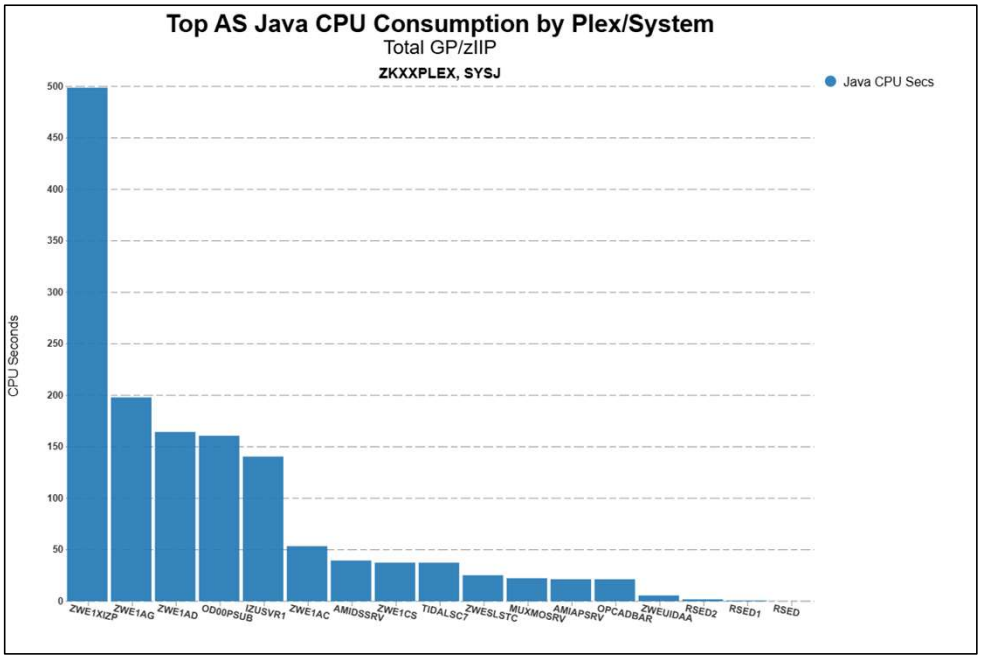

- This report shows the top Java CPU consuming address spaces

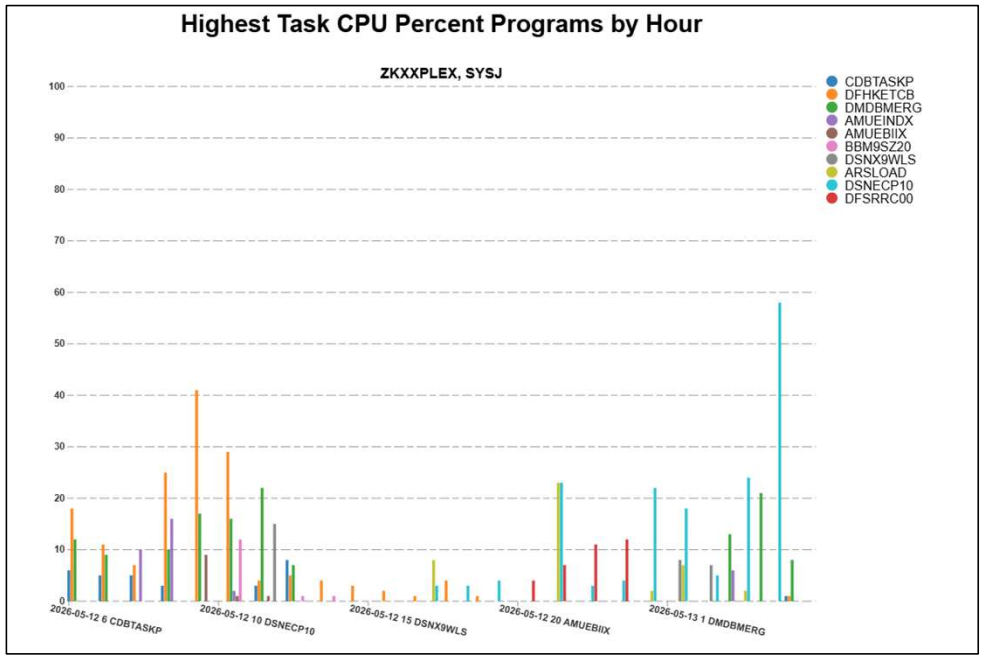

- This report shows the highest CPU task CPU intensities. This chart would be used to help to determine if you should migrate to a processor with more slower CPUs or fewer faster CPUs.

Sampling-frequency/diagnostic level

Measuring on a finer granularity level of other area discussed

Sampling-frequency/diagnostic level – Granule processor measurements

- Many usages of z/OS high frequency measurements
  - SMF 98 Includes metrics such as Utilization, Concurrency, Efficiency, Contention, Queuing
  - Processor utilization at the 2 second level
  - HiperDispatch pooling and decision points at the 2 second level
  - Processor speed changes at the 2 second level
  - HiperDispatch Topologies at the 2 second level
  - CPU consumption and CPU dispatch priorities of the WLM service class every 10 seconds
- SMF 98 – High-frequency Throughput Statistics (HFTS)
- SMF 99 – SRM/WLM details
  - Most common SMF records for measuring processor usage at the address space level
  - Based on options, turned on for each address space, and cut on the interval basis, job end basis, step end basis, and when address space is terminated.

Lots of Possible reports

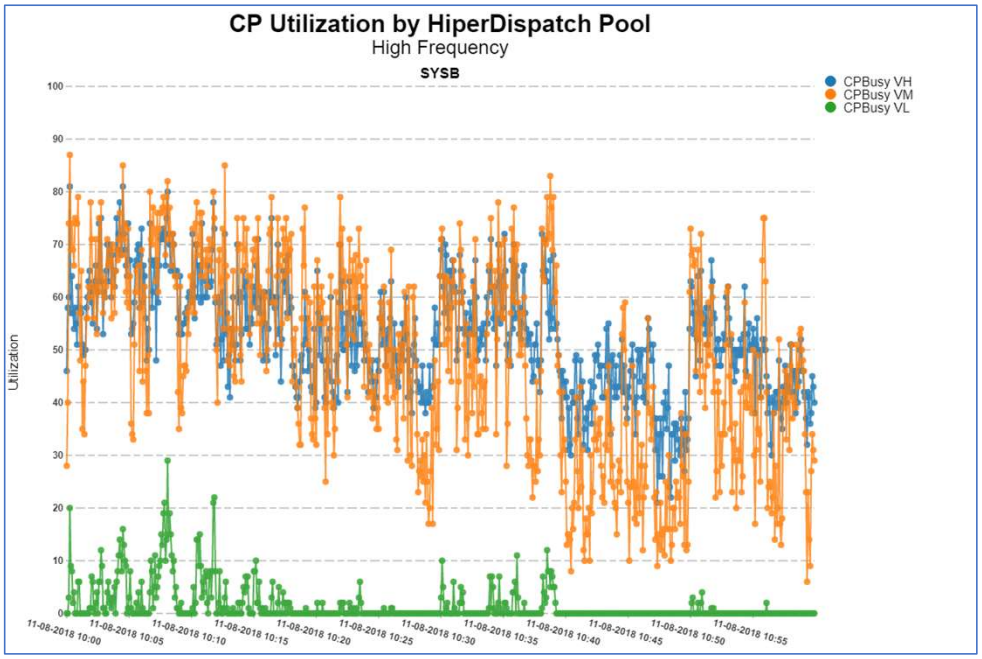

Here's an example of one hour of data showing the CP CPU busy by higher dispatch pool

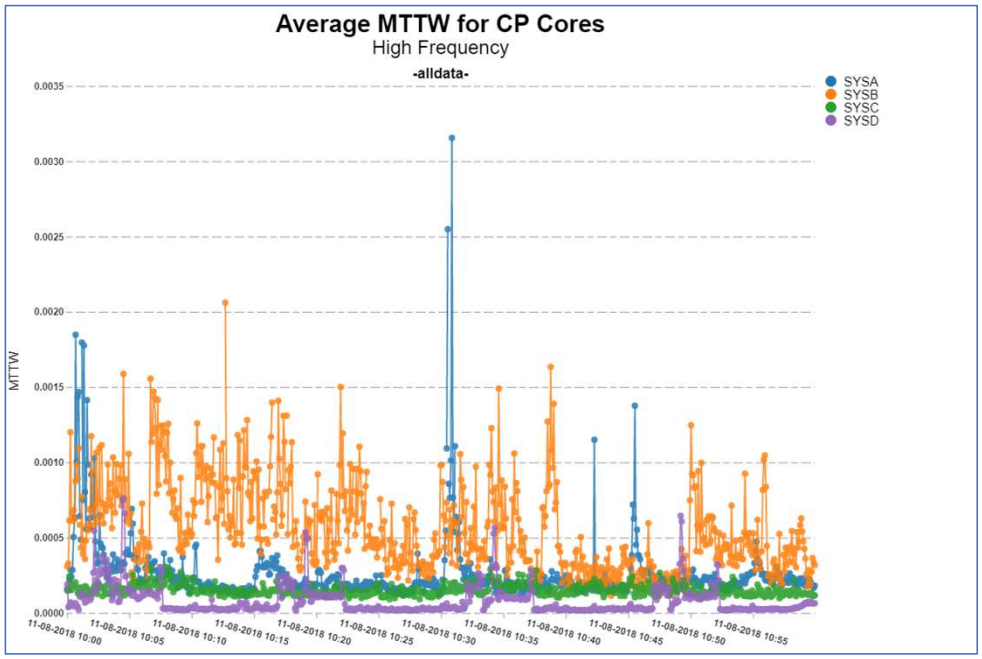

- Average Mean Time to Wait for CP Cores for each system. Kind of interesting to see how long it take for a CORE to go into a wait state

Lots of Possible reports

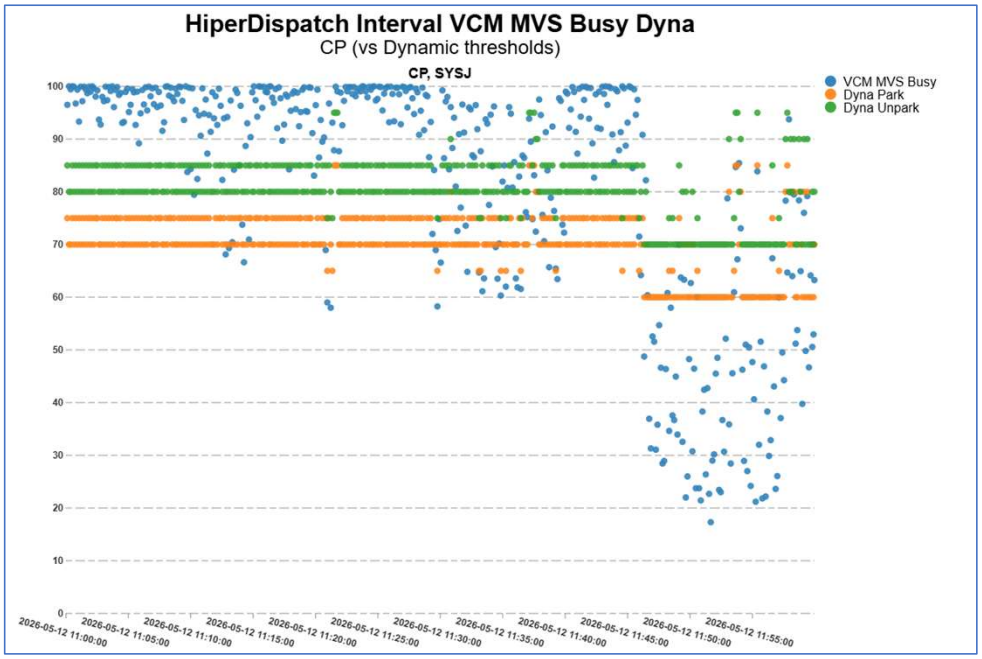

- SMF 99 HiperDispatch decision points with z/OS utilization on the 2 second level

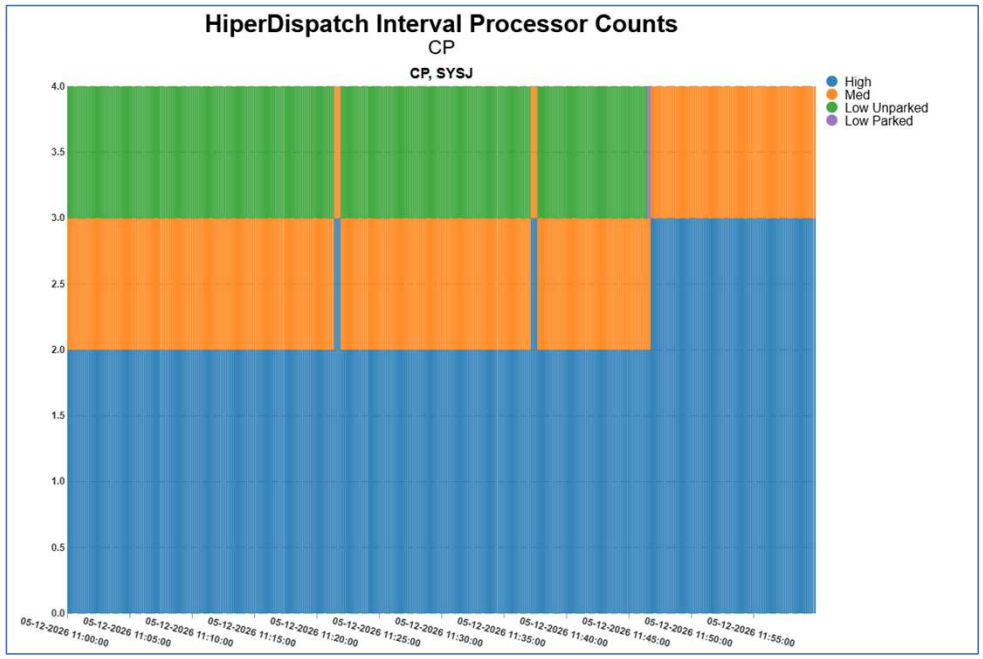

- SMF 99 HiperDispatch Parking and Unparking on a 2 second level

SMF 99
Looking at CPU Dispatching Priorities

- Do not assume goals and importance levels are correct
- Verify CPU dispatching priorities
- When a cap is enforced
  - Do the right workloads have first access to the CPU?

Imp4: DDFLO per 2

Imp5: BATLO per 3

Imp5: DDFBATLO per 1

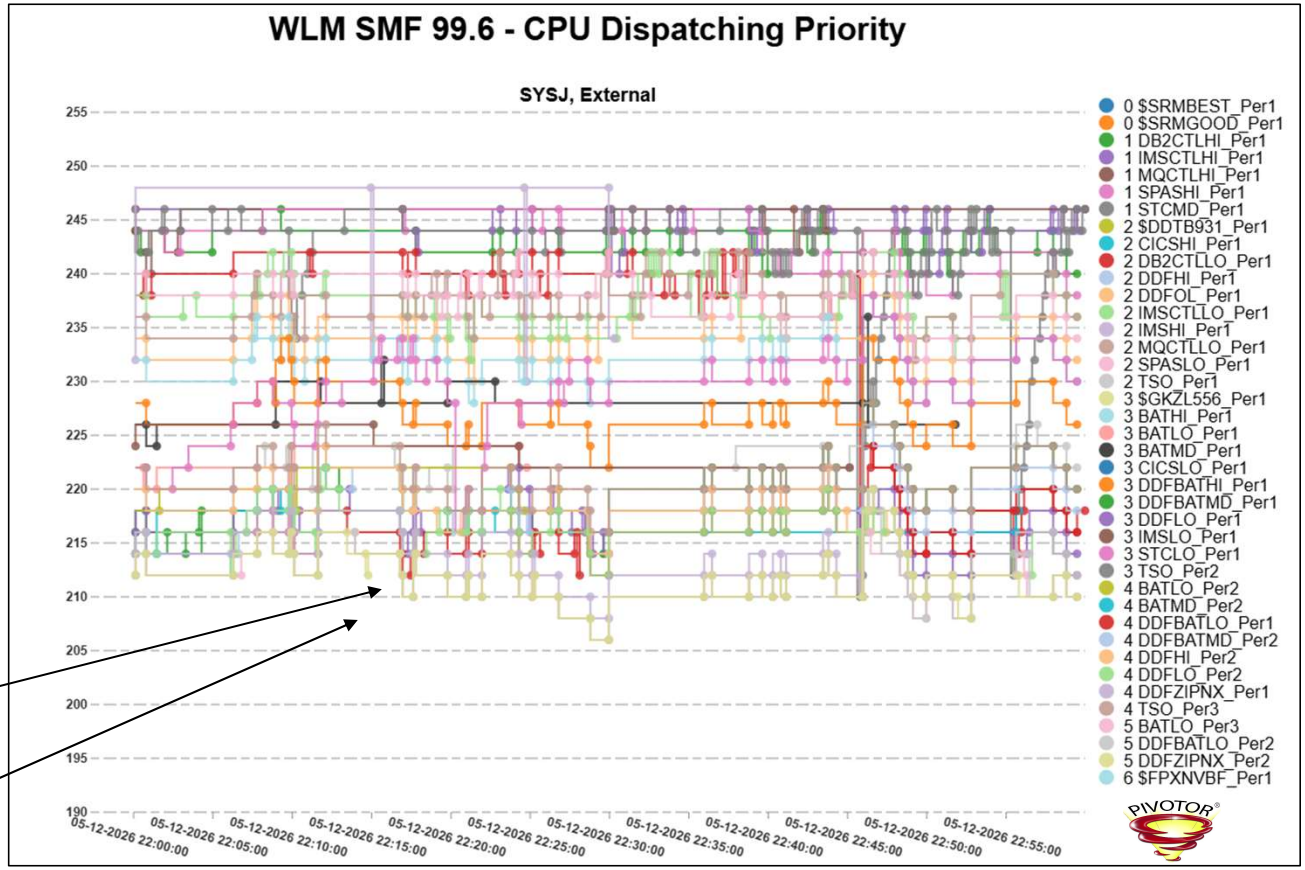

www.pivotor.com
chart

SMF 99
– CPU service consumed every 10 seconds

- Here is a chart of consumed by each WLM service class period and ordered by WLM importance level
  - Chart just shows 22:00 hour
- There are challenges in doing this, but assuming goals are correct, and CPU dispatching priorities are as hoped, then a somewhat straight forward exercise:
  - If you are thinking of imposing or lowering a maximum, one can roughly project which workloads will suffer
  - Must still look at how response times and velocities will change

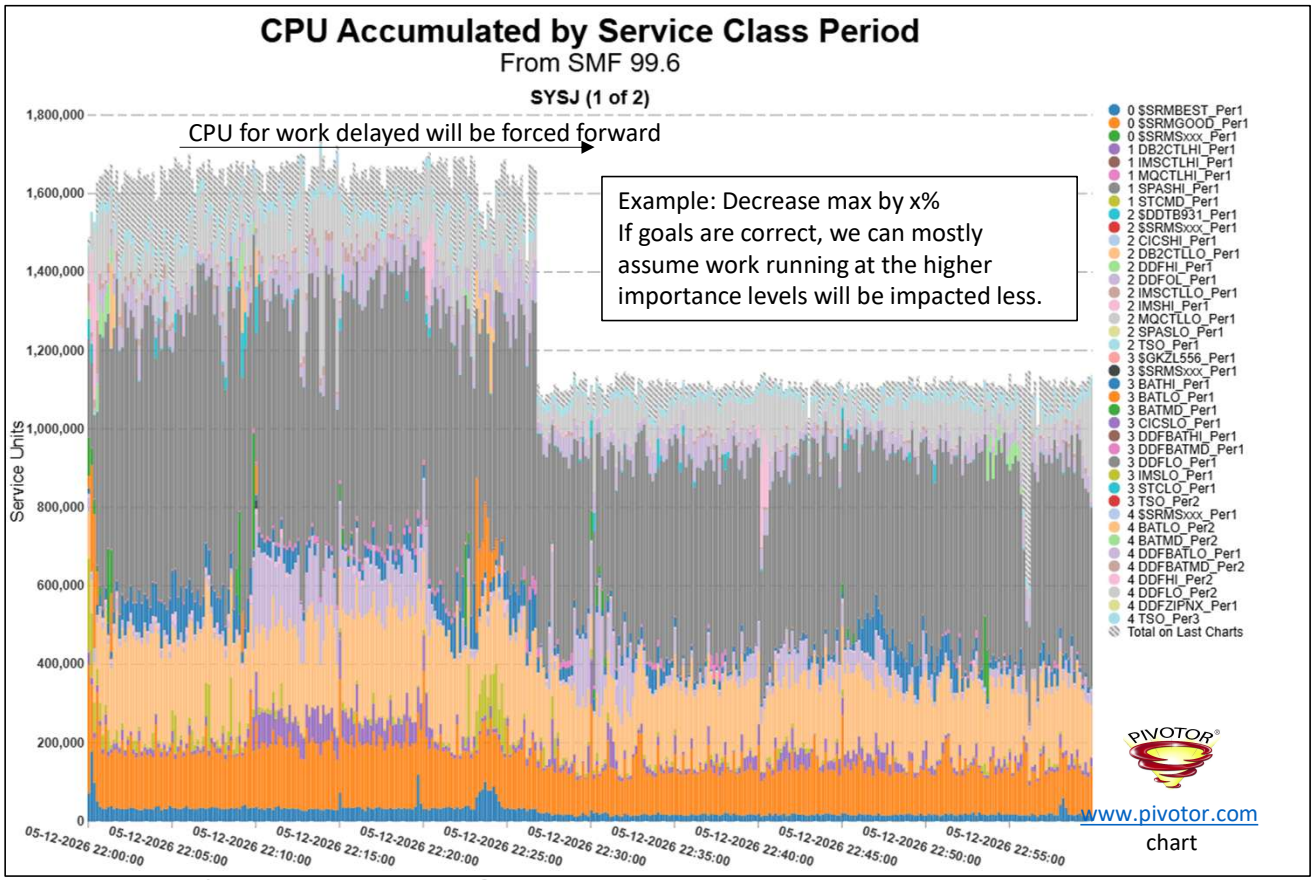

SMF 99
- Service consumed at CPU dispatching priorities

- Also look to see how much service is being made available above and below the large consuming workloads

- In this example, during capping DB2CTLHI does not use much CPU (dark blue)
  - But not much CPU used at higher priorities
  - And lots of CPU available to lower priority work

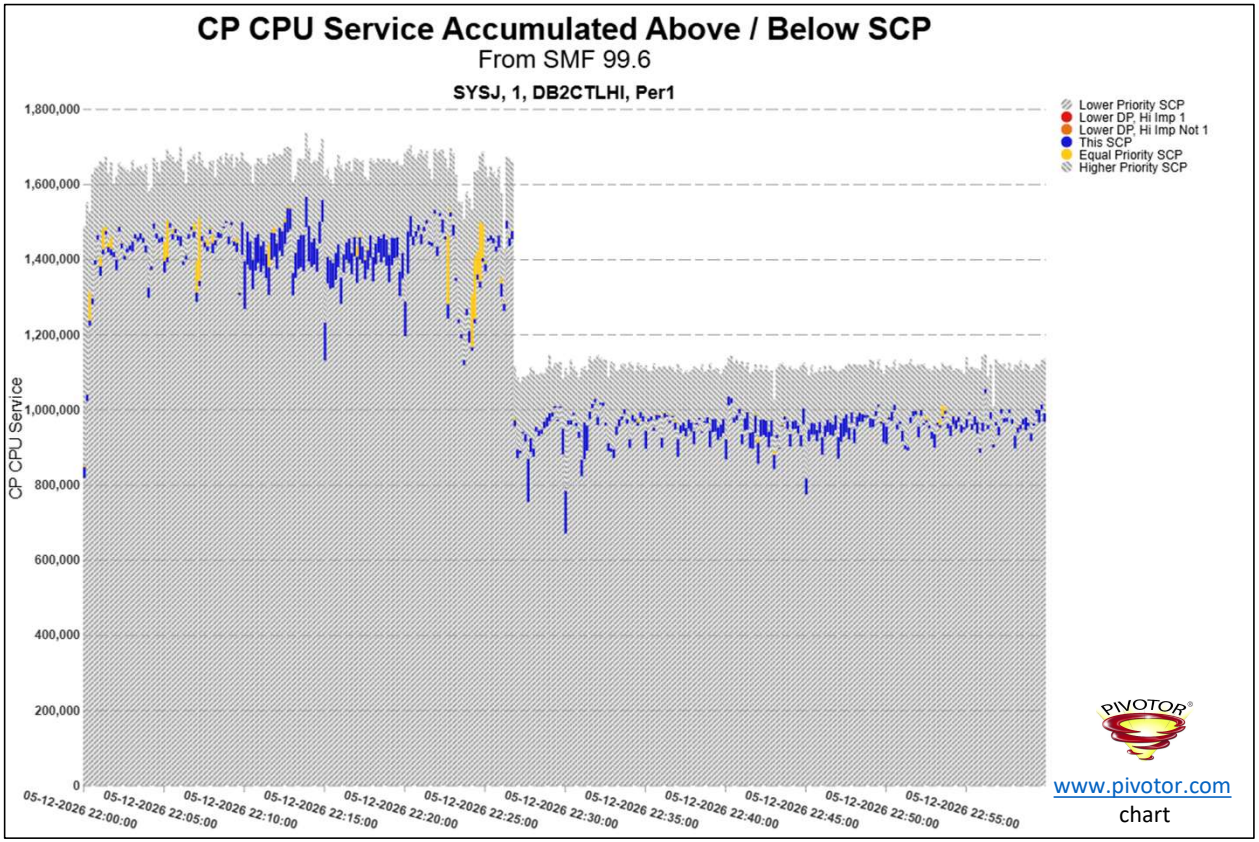

Transaction / Program level

Measuring online trans programs, database transactions

Transaction / Program Level – Application Level

- Many usages of z/OS transaction measurement
  - Measuring application and transaction CPU consumption

- SMF 101 – DB2 accounting
- SMF 110 – CICS transactions
- IMS performance log records
- SMF 116 – IBM MQ accounting
- SMF 120 – WAS transactions
- IDMs SMF record

Lots of Possible reports

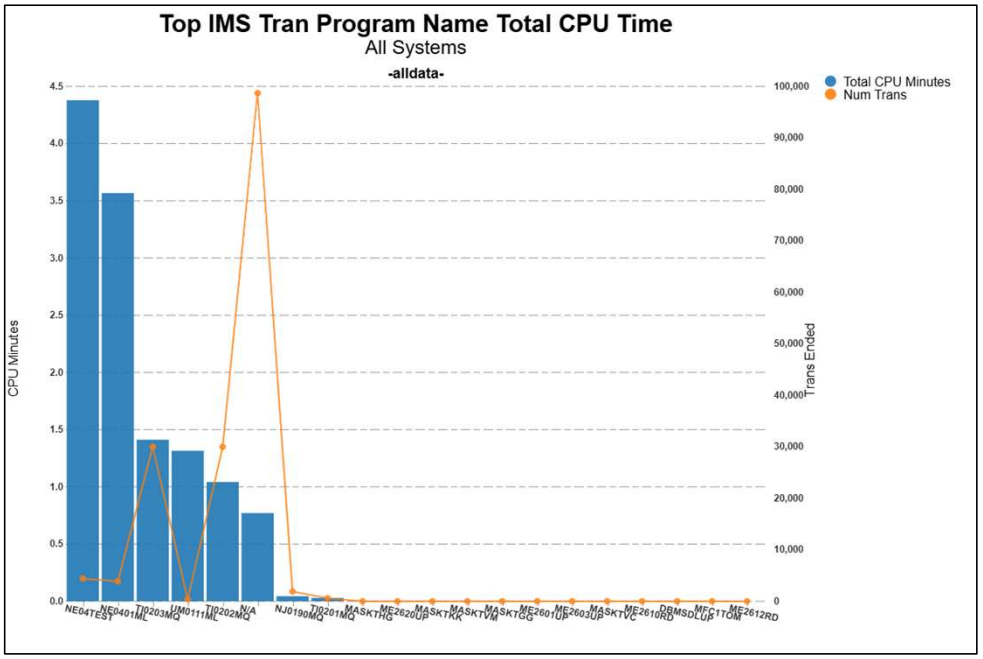

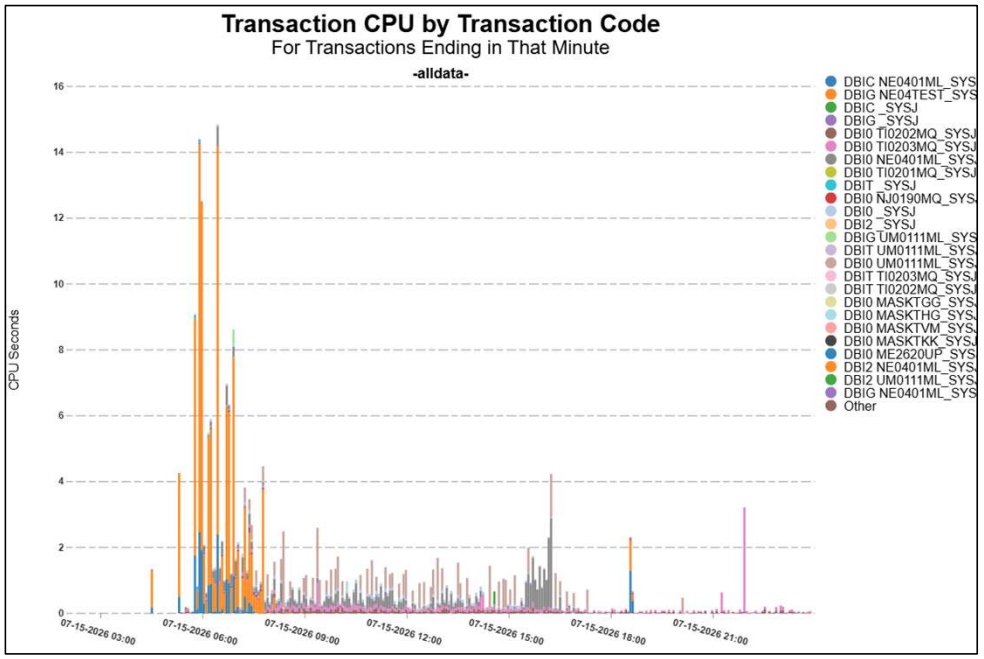

Pivotor Processor Analysis Reportsets

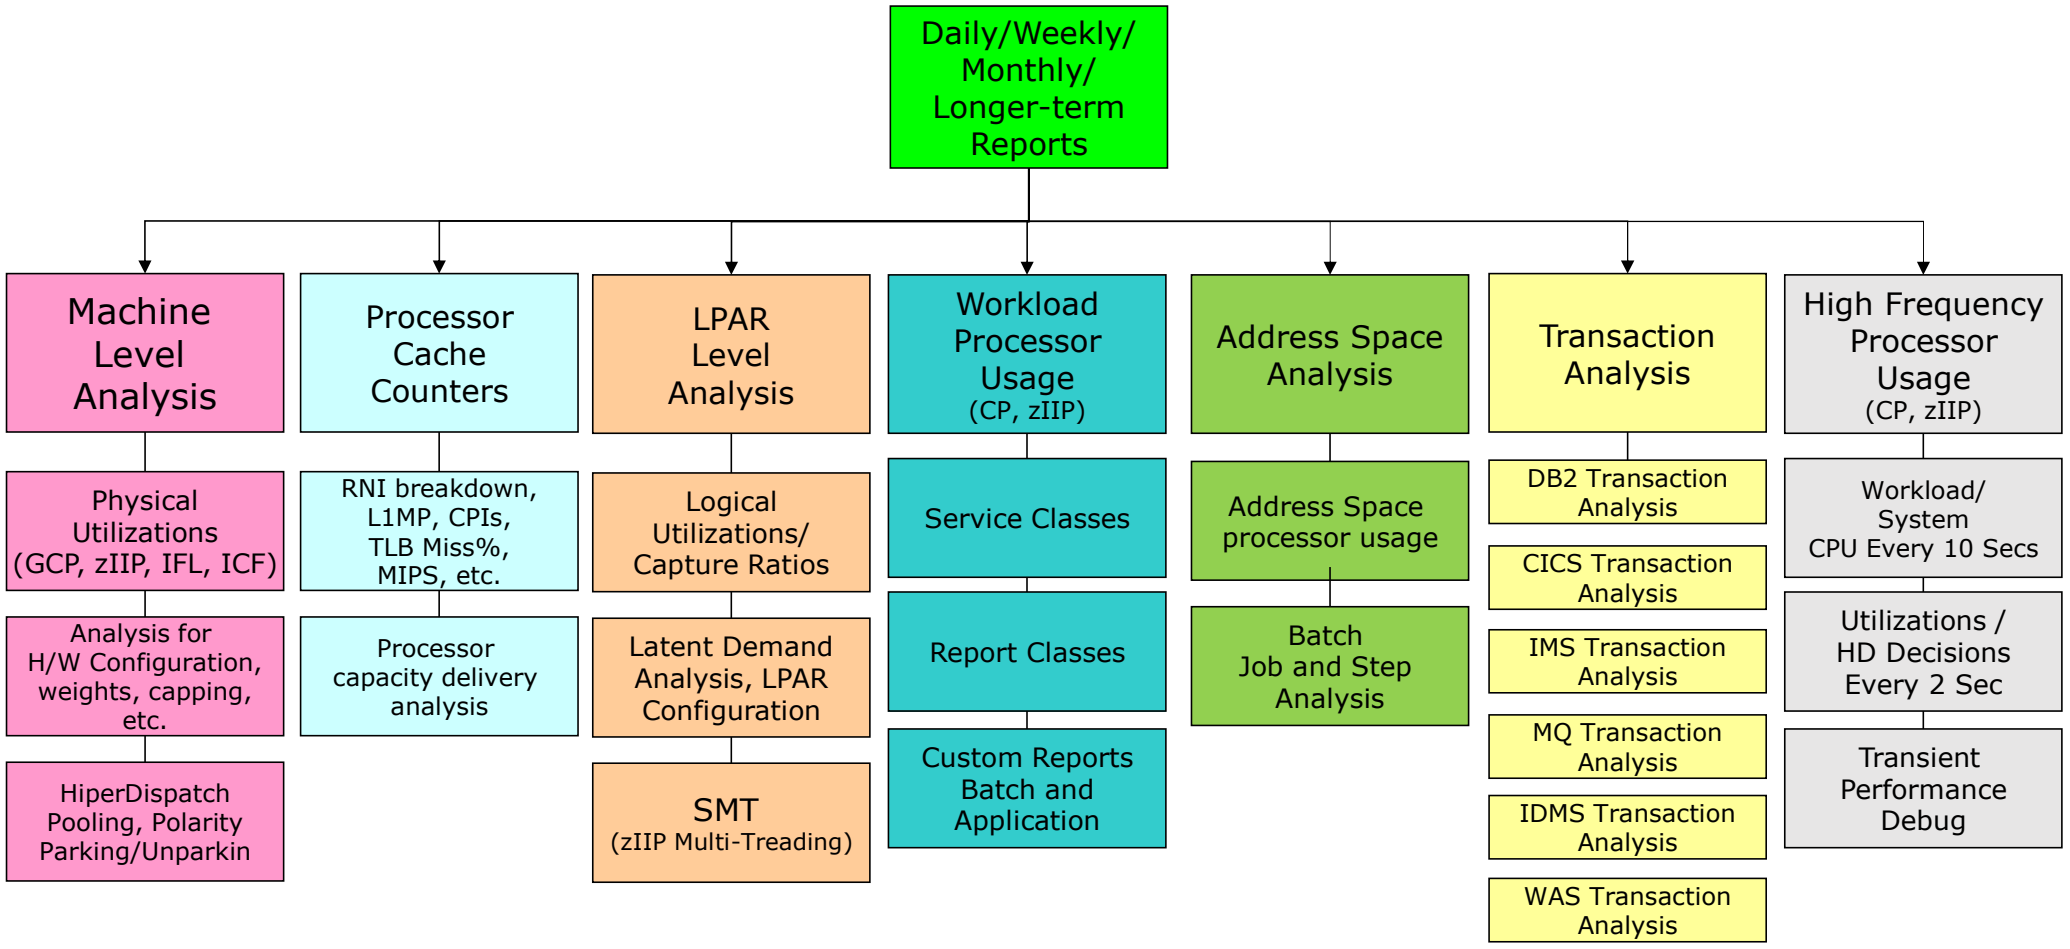

Thank You!

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