Pittsburgh, August 2019

CPU Measurements Demystified

Scott Chapman

Enterprise Performance Strategies

Scott.Chapman@epstrategies.com

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

EPS: We do z/OS performance…

- Pivotor - Reporting and analysis software and services
  - Not just reporting, but analysis based reporting based on our expertise
- Education and instruction
  - We have taught our z/OS performance workshops all over the world
- Consulting
  - Performance war rooms: concentrated, highly productive group discussions and analysis
- Information
  - We present around the world and participate in online forums

z/OS Performance workshops available

During these workshops you will be analyzing your own data!

- Essential z/OS Performance Tuning
  - Milwaukee WI, June 10-14, 2019
- Parallel Sysplex and z/OS Performance Tuning
  - Via the internet, TBD 2019
- WLM Performance and Re-evaluating Goals
  - Virginia Beach VA, October 21-25, 2019

Like what you see?

- The z/OS Performance Graphs you see here come from Pivotor™ but should be in most of the major reporting products

- If not, or you just 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 retention
  - That still encompasses over 100 reports!

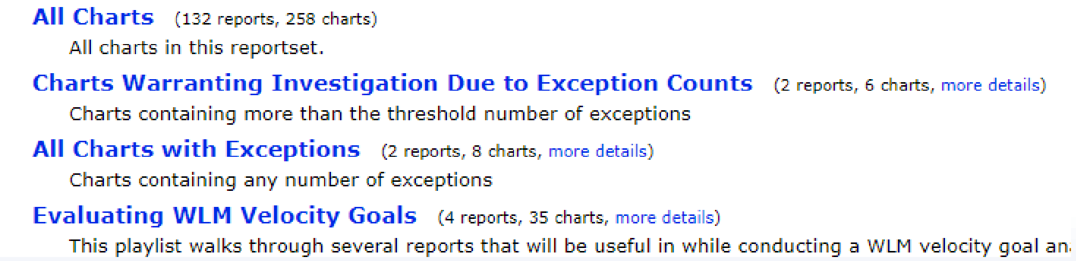

All Charts (132 reports, 258 charts)
All charts in this reportset.

Charts Warranting Investigation Due to Exception Counts (2 reports, 6 charts, more details)
Charts containing more than the threshold number of exceptions

All Charts with Exceptions (2 reports, 8 charts, more details)
Charts containing any number of exceptions

Evaluating WLM Velocity Goals (4 reports, 35 charts, more details)
This playlist walks through several reports that will be useful in while conducting a WLM velocity goal an

EPS presentations this week

What	Who	When	Where
Looking for Parallel Sysplex and Data Sharing Performance Opportunities	Peter Enrico	Tue 10:00	408
VSP: z/OS Performance Tuning – Some Top Things You May Not Know	Peter Enrico & Scott Chapman	Tue 2:15	304
z/OS CPU Measurements Demystified	Scott Chapman	Tue 3:30	408
z/OS Performance and Capacity Fundamentals for Mainframe Rookies	Scott Chapman	Tue 4:45	401-402
Revisiting WLM Options for the Modern Mainframe	Scott Chapman	Wed 10:00	408
Introduction to Parallel Sysplex and Data Sharing	Peter Enrico	Thu 3:30	317-318

You can also catch Scott at the COMUG presentation “Re-Framing the Mainframe…A User Group Research Project by COMUG” Thursday at 12:45 in 407

The Many Facets of CPU Measurements

- Mainframe CPU measurements are varied and sometimes confusing
- Yet understanding CPU measurements is very important
  - Software costs are usually the largest part of the mainframe budget
  - Software costs are usually driven by CPU capacity and/or consumption
- We need to understand CPU capacity
  - How much work can our CPU do?
- We need to understand CPU consumption
  - How much of our CPU capacity are we using?
- We also need to understand how our CPU capacity and consumption impacts our performance

CPU Nomenclature

- In this presentation, the terms CP, CPU, or engine are all meant generically and can apply to both specialty engines and general purpose engines
- Where the distinction matters:
    - GCP or GP = General Purpose Engines
    - zIIP = Unless otherwise noted, assume everything about zIIPs applies to zAAPs as well
        - zAAP on zIIP capability on current machines allow zAAP eligible work to be dispatched to zIIPs
        - zAAPs not available on z13 and later machines
- In this presentation, ignore SMT (available for zIIPs only) as that greatly complicates things
- Measurements for coupling facilities and IFLs will not be discussed

Primary Sources of CPU Measurements

- SMF 70 – LPAR
- SMF 72 – Workload (service class / report class)
- SMF 30 – Address Space
  - End of step (subtype 4) and end of job (subtype 5) records contain totals
    - Often used for chargeback and application attribution
  - Interval records (subtypes 2 and 3) record activity for just 1 interval
    - Can be used to find top consumers within a Service Class
  - Recommendations:
    - Set RMF to sync with SMF interval
    - Set SMF interval to no more than 15 minutes

CPU Time

- CPU time = total time that a CPU has spent performing work for task
  - Time that a workload is dispatched to a CPU
- z/Architecture provides CPU timers with nominal resolution of 1 µs
  - Yes, that's one millionth of a second, although the times aren't usually externalized to that precision
- Instruction EXTRACT CPU TIME (ECTG) can be used by problem-state programs to determine the amount of CPU time consumed by the current task
- When a CPU is interrupted to process something else, the CPU timer is readjusted once the interrupted task is dispatched again
- CPU timers are not dependent on the time-of-day clock because (e.g.) the time of day clock may be steered to remain in sync with a time source
  - One of the reasons why use of system time for performance analysis is problematic

CPU Time Variability

- Despite very precise CPU timers to measure CPU time, if you run the exact same workload multiple times, you will almost certainly observe variability in the CPU time
- There are a large number of things that can cause CPU time variation, a few significant ones include:
  - Processor cache contention
  - Lock availability / contention
  - Serialization of resources between CPs
  - Program load variability due to caching or directory changes
  - Changes in dataset allocation details such as the number of extents
- Variability can be significant: +/- 5% is not be out of the ordinary
  - Even larger variability is possible, even likely, depending on the configuration
  - When evaluating before/after performance, it's important to take multiple measurements: at least 3, preferably 5

CPU Time Availability

Again: Ignore SMT!

- Only one task can be dispatched to a CPU at time
  - So no more than one task will accumulate CPU time per CPU at any instant
- The max CPU time possible to accumulated in 1 second is 1 second
  - Across all LPARs on the machine
- In one 900 second RMF interval, JobA has accumulated 450 CPU seconds
  - If this ran on a CEC with 5 active CPUs, what portion of the CEC did JobA consume?
- How much CPU time would this job consume if it was moved to a machine whose CPUs were:
  - Twice as fast?
  - Half as fast?

CPU Time Answers

- 5 Active CPUs * 900 second interval = max 4500 CPU seconds / interval
- JobA's 450 CPU seconds = 450 / 4500 = 10% of available capacity
- If the CPUs were twice as fast, we'd expect them to get twice as much work done per unit of time so JobA would be expected to take 225 CPU seconds
- If the CPUs were half as fast, we'd expect them to get half as much work done per unit of time so JobA would be expected to take 900 CPU seconds

CPU Time Differences Between Machines

- CPU time will change between:
  - Generations of machines
  - Models of the same machine
  - Configurations of the same model
- So we'd like some sort of measurement that's comparable across machines
- We don't really have that, but we have metrics that try to *approximate* that:
  - SUs – Service Units
  - MSUs – Millions of Service Units
  - MIPS – Millions of Instructions Per Second
    (or Meaningless Indicators of Processor Speed)

Service Units

- Service Units intended to translate resource utilization to value that can be used transparently between machine configurations
  - WLM policy uses SUs to measure resource consumption
- Actual recorded values based on service coefficients defined in the WLM service definition (except when they're "unweighted")
  - Recommended values today:
    - CPU: 1, SRB: 1, I/O: 0, MSO: 0
  - CPU Service units = CPU time * SRM Constant * CPU Coefficient
  - SRM Constant recorded in SMF data, available in control blocks, online here: https://www-304.ibm.com/servers/resourcelink/lib03060.nsf/pages/srmindex
  - SRM constant is based on number of GCPs online to the LPAR and the machine model
  - So on a Z05, an LPAR with 3 online GCPs gets the SRM constant for a Z03
    - Even if one of those 3 is parked

(Ignoring details about types of CPU time)

CEC Average CP CPUs Online

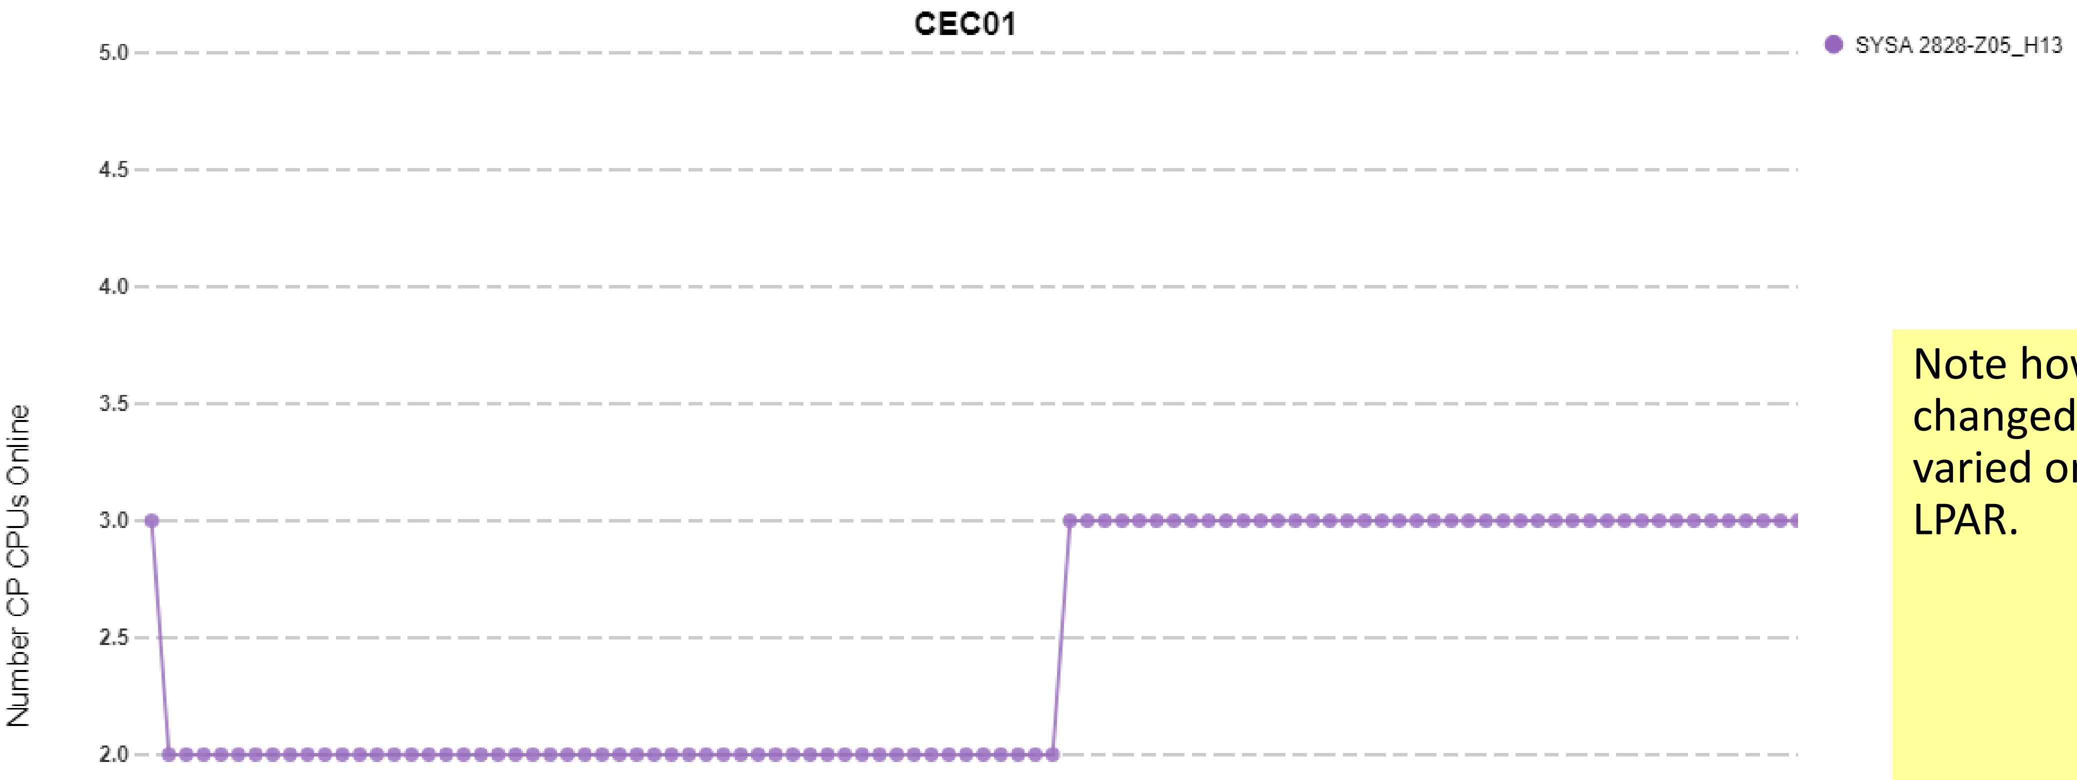

SU/Sec Constant - For Each System Over Time

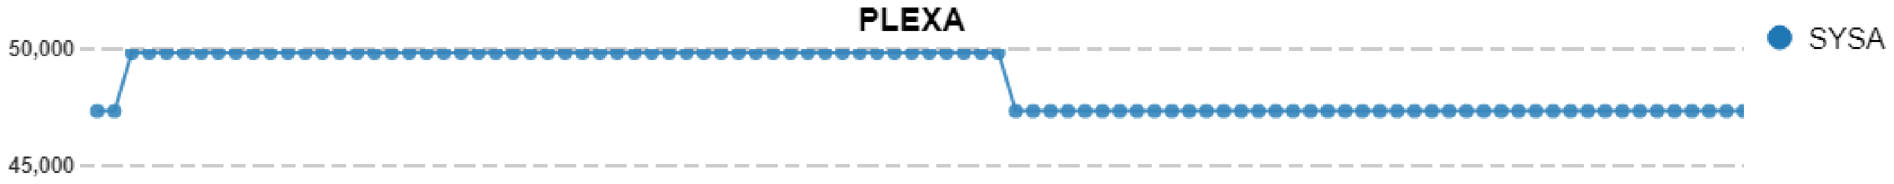

Note how the SU/Sec changed when a CP was varied on/off from the LPAR.

SU/Sec for some zBC12 Machines

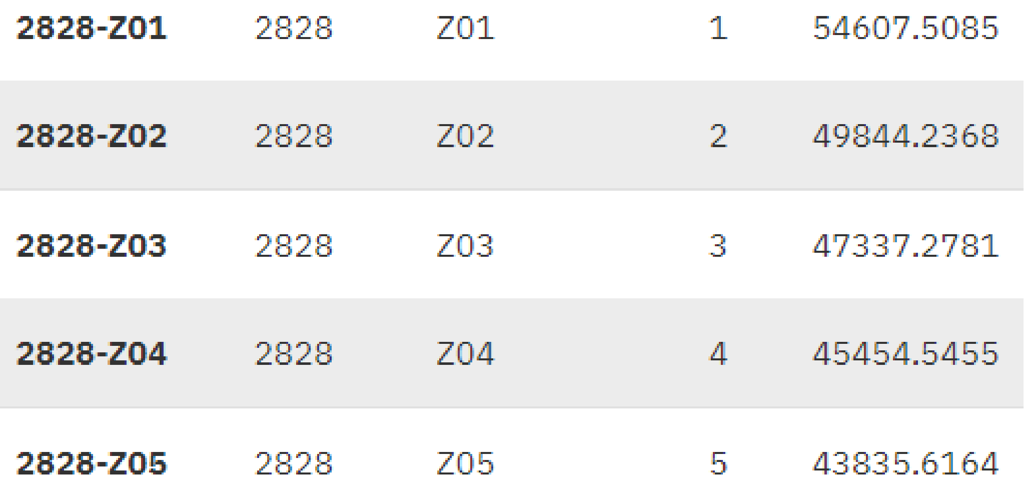

2828-Z01	2828	Z01	1	54607.5085
2828-Z02	2828	Z02	2	49844.2368
2828-Z03	2828	Z03	3	47337.2781
2828-Z04	2828	Z04	4	45454.5455
2828-Z05	2828	Z05	5	43835.6164

- Notice how LPARs with smaller number of CPs get higher SU/sec constants
- This makes sense for a dedicated machine due to the MP-effect
  - Today, in some configurations, HiperDispatch may make this less accurate than it was in the past
    - In previous example, the third CP was parked virtually the entire time
  - But this is how z/OS does it: nothing you can do about it
  - And SU/sec is not so accurate that it would really matter the SRM controls that use it
  - But it shows how SU/sec isn’t necessarily any more accurate than MIPS or MSUs

MSUs

- Originally, really was approximately Millions of SUs/hour
  - SU/sec * CPUs * 3600 / 1 Million ≈ MSUs
- With the z990 in 2003, IBM introduced the “technology dividend” and reduced the MSU ratings for the processors 10% relative to the “hardware” MSU ratings
  - Reduce software costs as an incentive to upgrade to a new machine
  - This continued for the z9 and z10 machines
  - Starting with the z196, IBM switched to delivering software cost reductions with direct software pricing mechanisms
- So the rated MSUs is no longer directly tied to the SU/sec rating
  - Although it is quite close to: (SU/sec * CPUs * 3600 / 1000000) * 0.664

MIPS

- Millions of Instructions Per Second
  - Only had a chance of being somewhat valid definition around the dawn of computing
    - Different instructions take different amounts of time to execute + other complications
    - In 1972, S/360 158, actually had a count of instructions, on a particular workload
    - But… almost immediately realized that they couldn't get 1 MIPS for TSO on the 1 MIPS machine
  - Meaningless Indicators of Processor Speed
  - Meaningless Indicator Provided by a Salesman
- MIPS ratings are provided from multiple sources
  - Gartner (most commonly referenced by ISVs)
  - Cheryl Watson
  - Other consultants / products
  - You can generate your own from the LSPR data
    - That's how the consultants do it
    - Although in reality, you really want to use zPCR to compare machine configurations
- IBM provides "PCI" which is essentially MIPS, without calling it MIPS
  - Possibly because "MIPS" has connotations for software pricing
  - PCI = Processor Capacity Index
  - IBM Speakers will often refer to PCI as MIPS, I try to use PCI so the origin is understood

When asked "How many MIPS do we have?" my first answer was always "Whose MIPS chart is referenced in the software contract you're concerned about?"

LSPR and “Single-number” Metrics

- LSPR = Large System Performance Reference
  - Large System = Mainframe
  - Provides capacity comparison between machines based on particular measured workloads at high utilizations
  - Same workloads run in same environment on old and new machines to compare results
  - Workloads and environments periodically updated
- SRM Constants, MSUs, MIPS all derived on the “average” LSPR workload for a “median” customer LPAR configuration
  - Machine B can process 2x as many transactions/second as Machine A
  - A was said to be 200 MIPS so we now say B is 400 MIPS
  - ***Everything is relative!***
- Not every configuration tested: results interpolated

Relating SU/sec, MSUs, PCI

- Because they're all sourced from the same ratios, SU/sec, MSUs, and PCI are all in fact related and should track each other in consistent ratios
- In other words:
  - Multiply MSUs by ~8.1 to approximate PCI
  - Divide SU/Sec by about 52.2 (or maybe 50.8) to approximate PCI
- Some of the variability no doubt due to rounding intermediate results

These charts break one of my own data visualization rules!

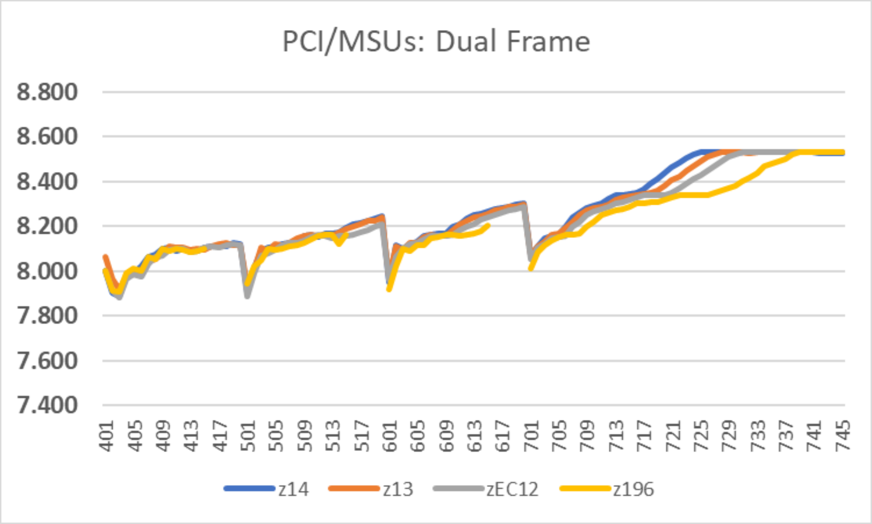
PCI/MSUs: Dual Frame

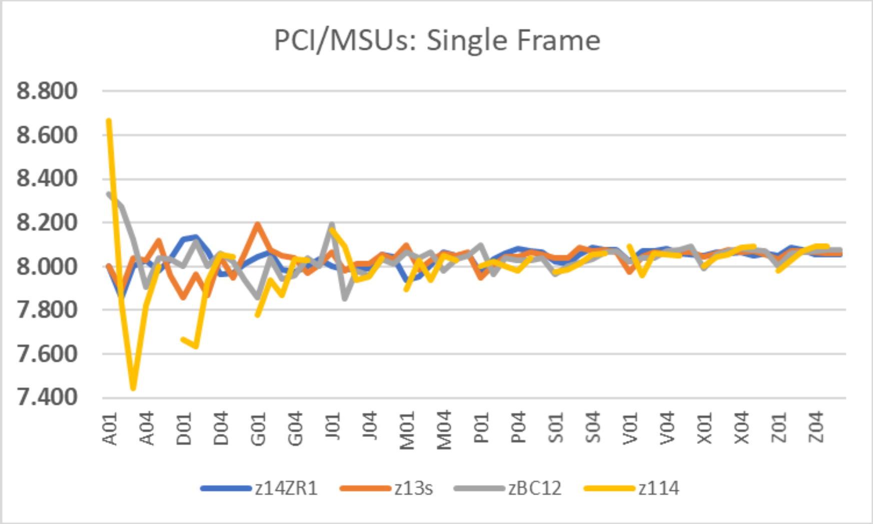
PCI/MSUs: Single Frame

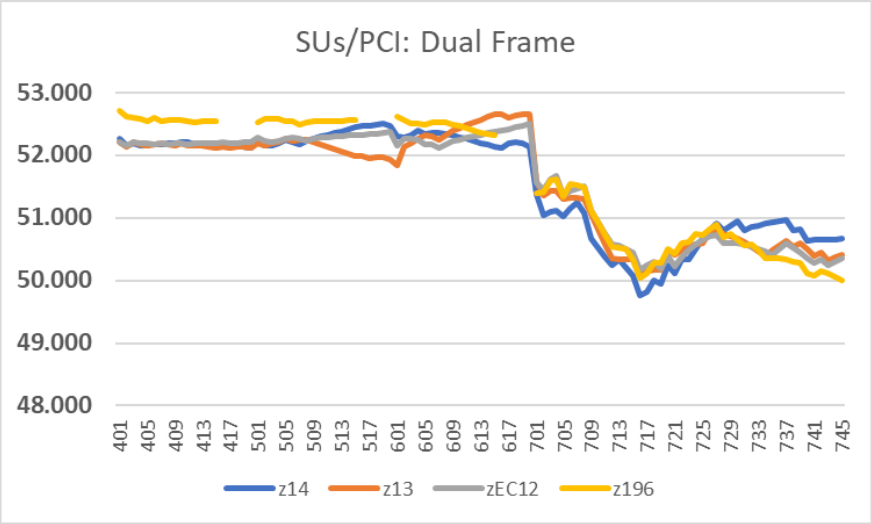
SUs/PCI: Dual Frame

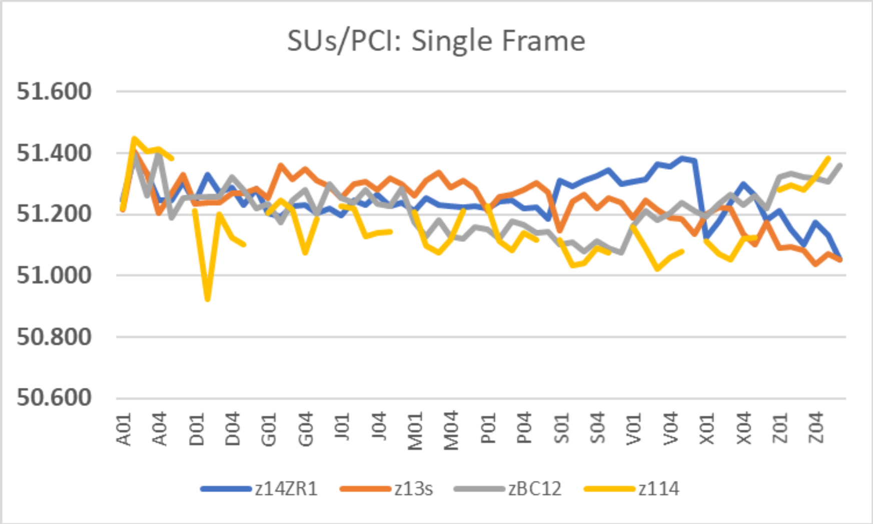
SUs/PCI: Single Frame

Relating SU/sec, MSUs, PCI

- When we make the charts zero-based it makes it easier to understand that the variability in the ratios is really pretty small percentage-wise

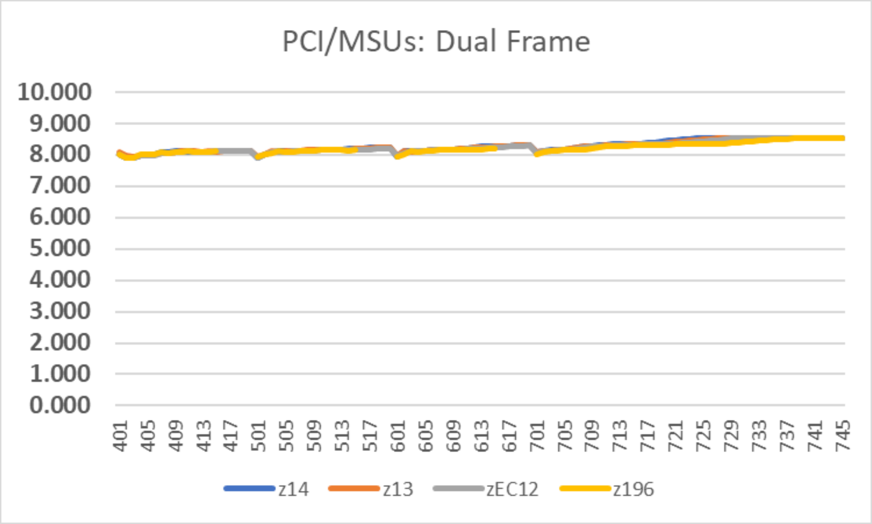

PCI/MSUs: Dual Frame

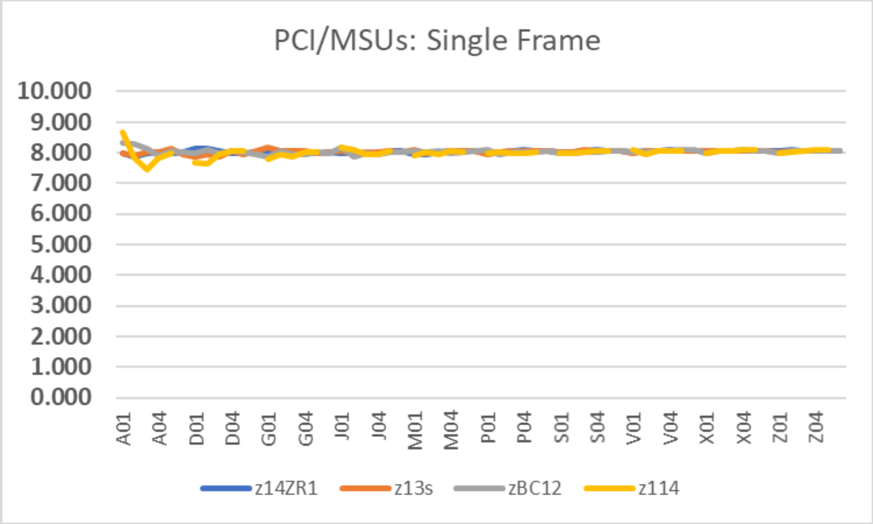

PCI/MSUs: Single Frame

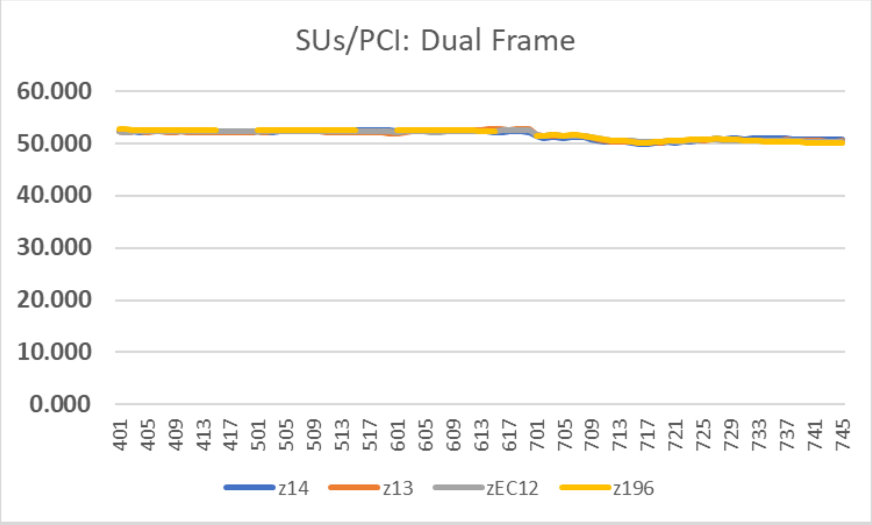

SUs/PCI: Dual Frame

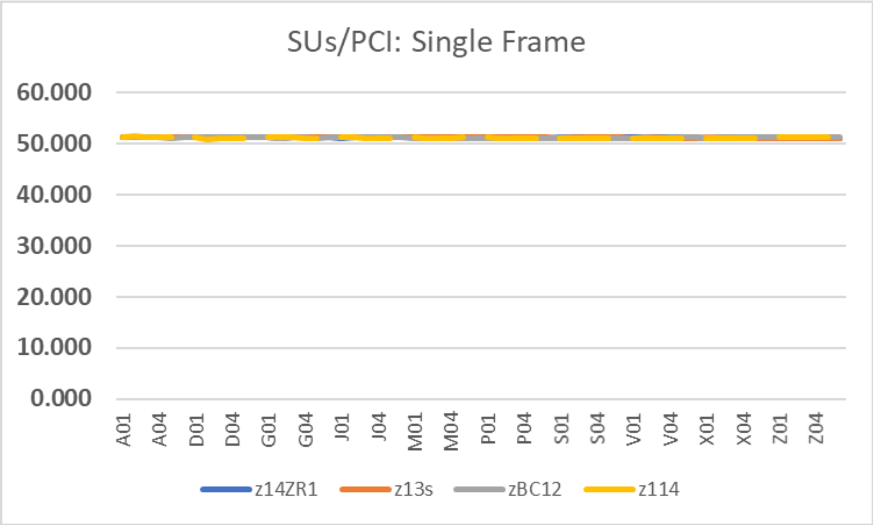

SUs/PCI: Single Frame

Unless there is a very good reason for it, charts should be zero-based!

What are these good for?

- SUs
  - WLM & SRM measurement of workloads (recorded in some SMF records)
- MSUs
  - Primarily used for IBM MLC pricing
- MIPS
  - Often used for ISV pricing
- All
  - Because your manager insists on seeing a report in MIPS or MSUs
  - Approximate estimate of machine capacity relative to other machines
  - But you really need to use zPCR to evaluate your particular configuration instead of the “average” workload in a “median” configuration

Previous calculations are interesting, but when you really need these values, you need to look them up

More Concrete Measurements

- CPU Time
  - Time spent using a processor
- Physical Utilization Percent
  - CEC Utilization
  - APPL% - Percent of one CP a workload is using
  - Derived from the CPU time
- CPs Used
  - Basically utilization percent as well, but scaled such that 100% = 1 CPU capacity
  - So if an LPAR used 1350 seconds in a 900 second interval, it used 1.5 CPs of capacity
  - AKA, Appl %

- All of these let me intuitively understand how much of my capacity I'm using
  - Aren't susceptible to issues such as changing measurements when varying CPs online

Deriving CPU Measurements

- Everything derives from the CPU time, but some measurements in SMF are recorded in SUs
  - CPU time = (SU / CPU coefficient) / SU per sec constant
- Physical Utilization is just CPU time / interval time
  - 300 CPU seconds in a 900 second interval = 33.3 APPL%
  - Which is 0.333 of a single CP
- For percent of total CEC: CPU / (interval time * CPs)
  - If an LPAR used 1200 CPU seconds 1200 / (900 * 3) = 44.4% of the total machine
- To convert to MIPS or MSUs, multiply the percent of the total CEC by the MIPS or MSU rating of the machine (3907-L03: 929 PCI, 116 MSUs)
  - 929 * .444 = 412.5 PCI
  - 116 * .444 = 51.5 MSUs

Breaking Down CPU Consumption

- Total utilization of the CEC = total dispatch time of all the LPARs + time spent to manage LPARs (*PHYSICAL*)
- Note Uncaptured time within an LPAR
  - This is time that cannot be directly attributed to a workload

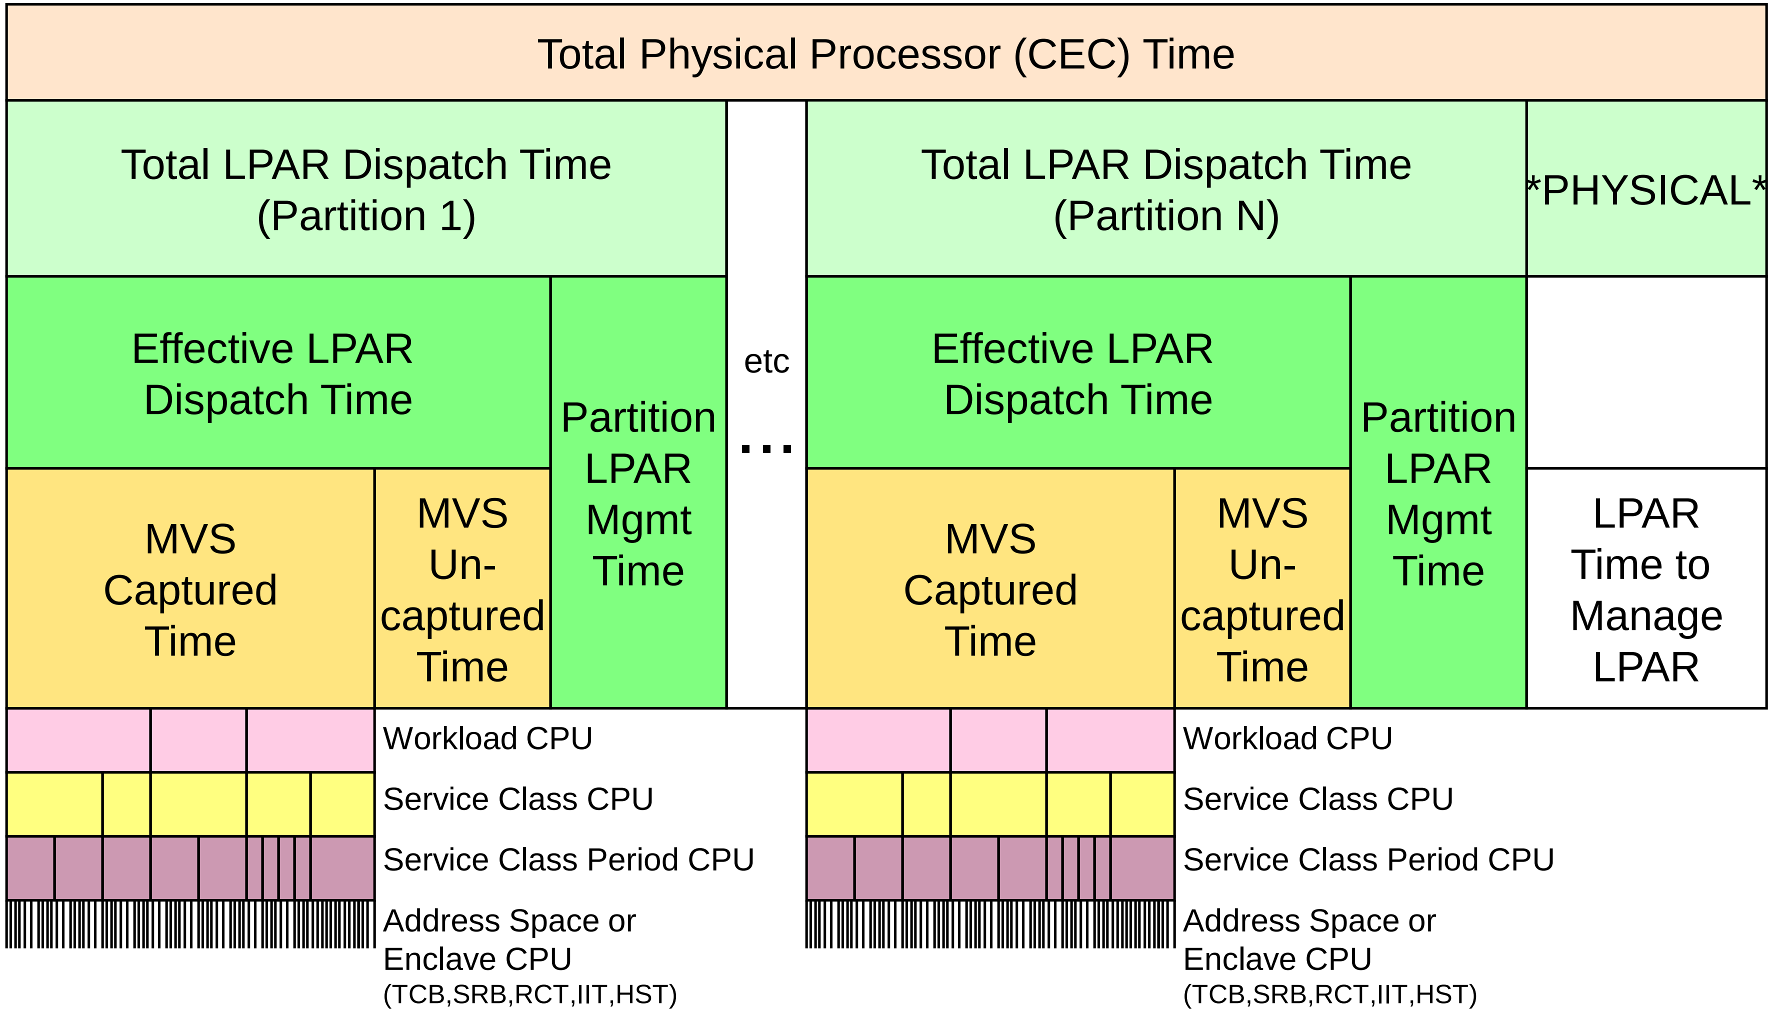

Total Physical Processor (CEC) Time

Total LPAR Dispatch Time (Partition 1)

Effective LPAR Dispatch Time

MVS Captured Time

MVS Un-captured Time

Partition LPAR Mgmt Time

Workload CPU

Service Class CPU

Service Class Period CPU

Address Space or Enclave CPU (TCB,SRB,RCT,IIT,HST)

etc …

Total LPAR Dispatch Time (Partition N)

Effective LPAR Dispatch Time

MVS Captured Time

MVS Un-captured Time

Partition LPAR Mgmt Time

Workload CPU

Service Class CPU

Service Class Period CPU

Address Space or Enclave CPU (TCB,SRB,RCT,IIT,HST)

PHYSICAL

LPAR Time to Manage LPAR

Breaking Down CPU Consumption

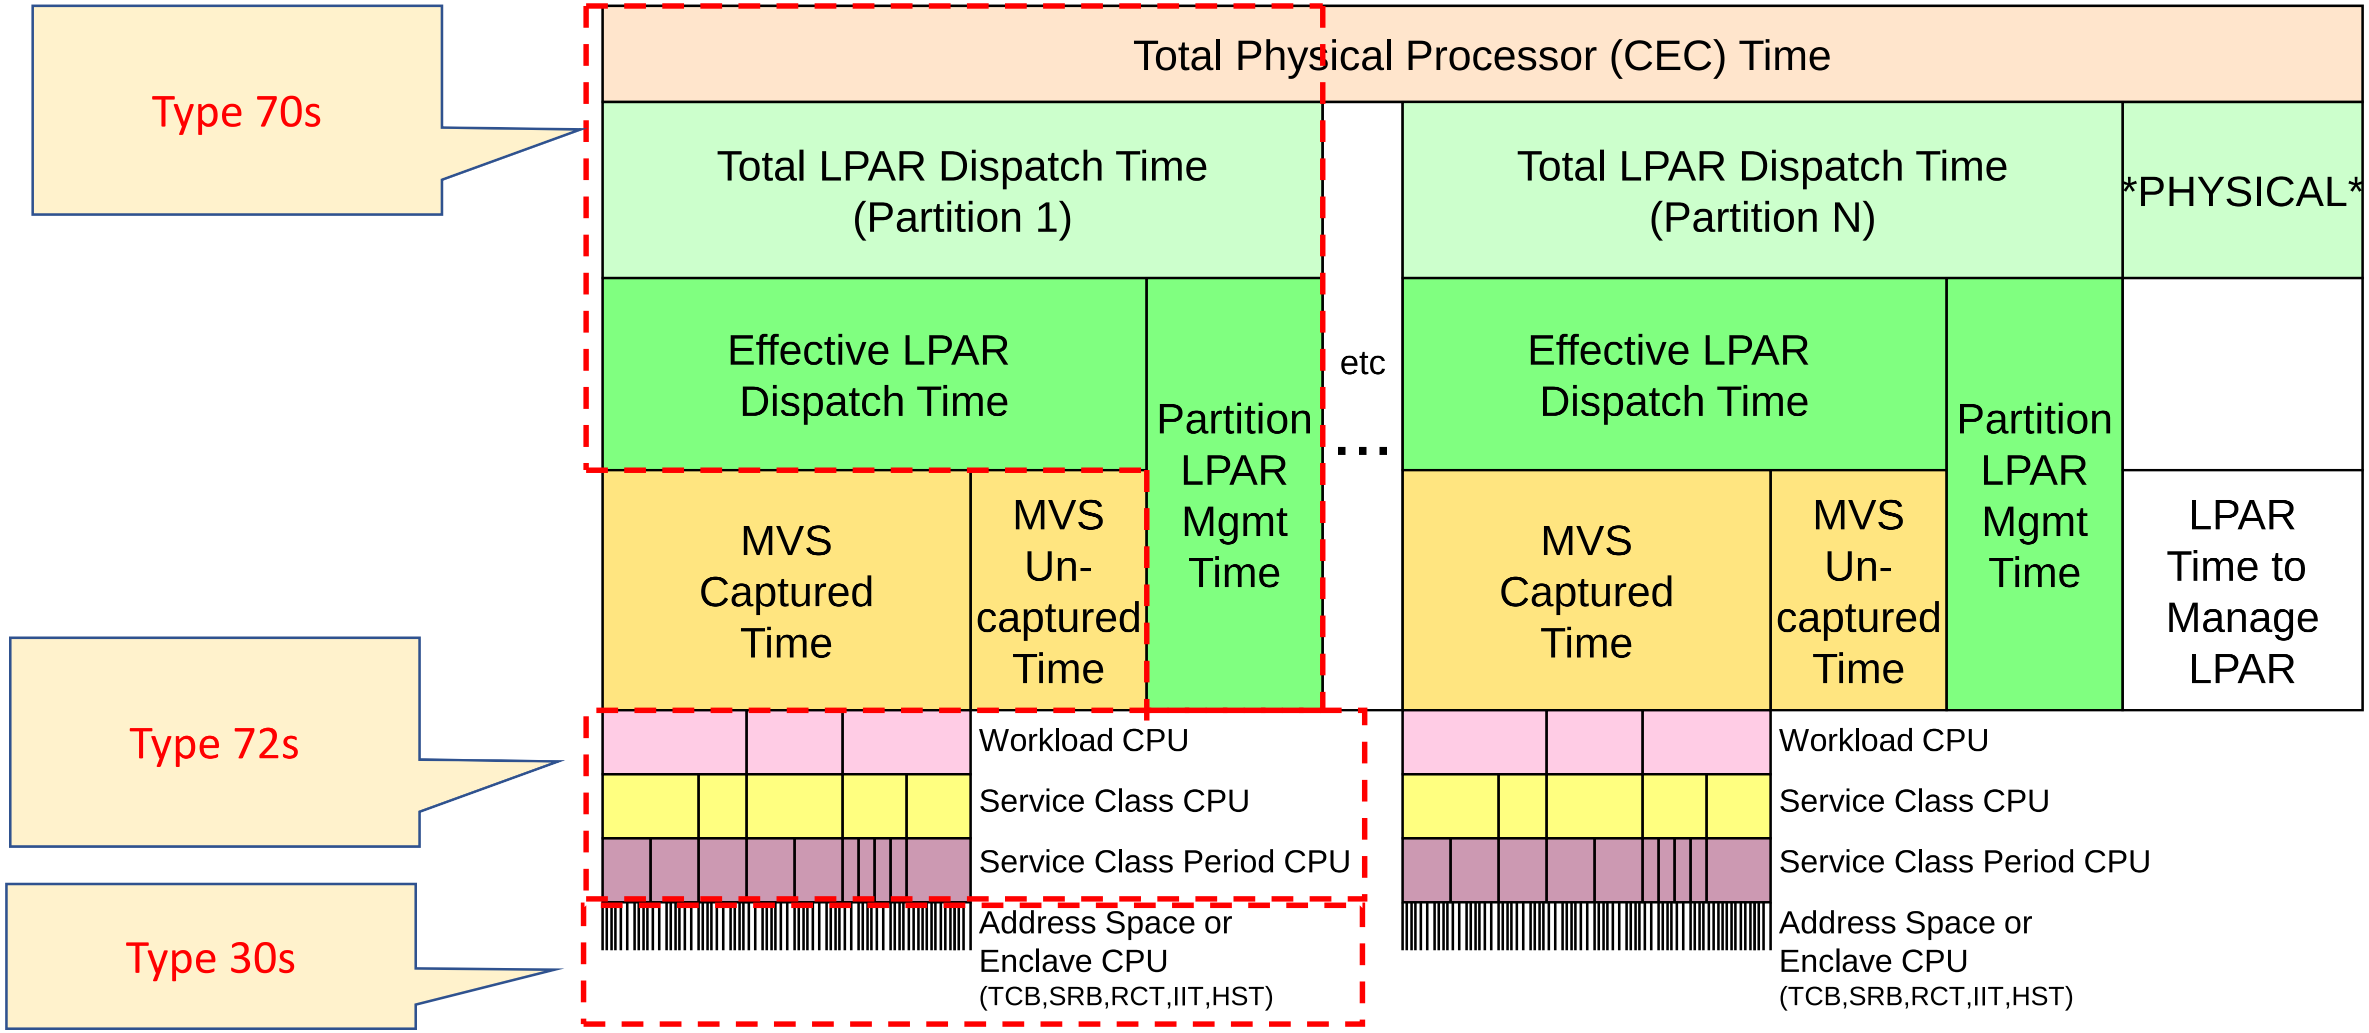

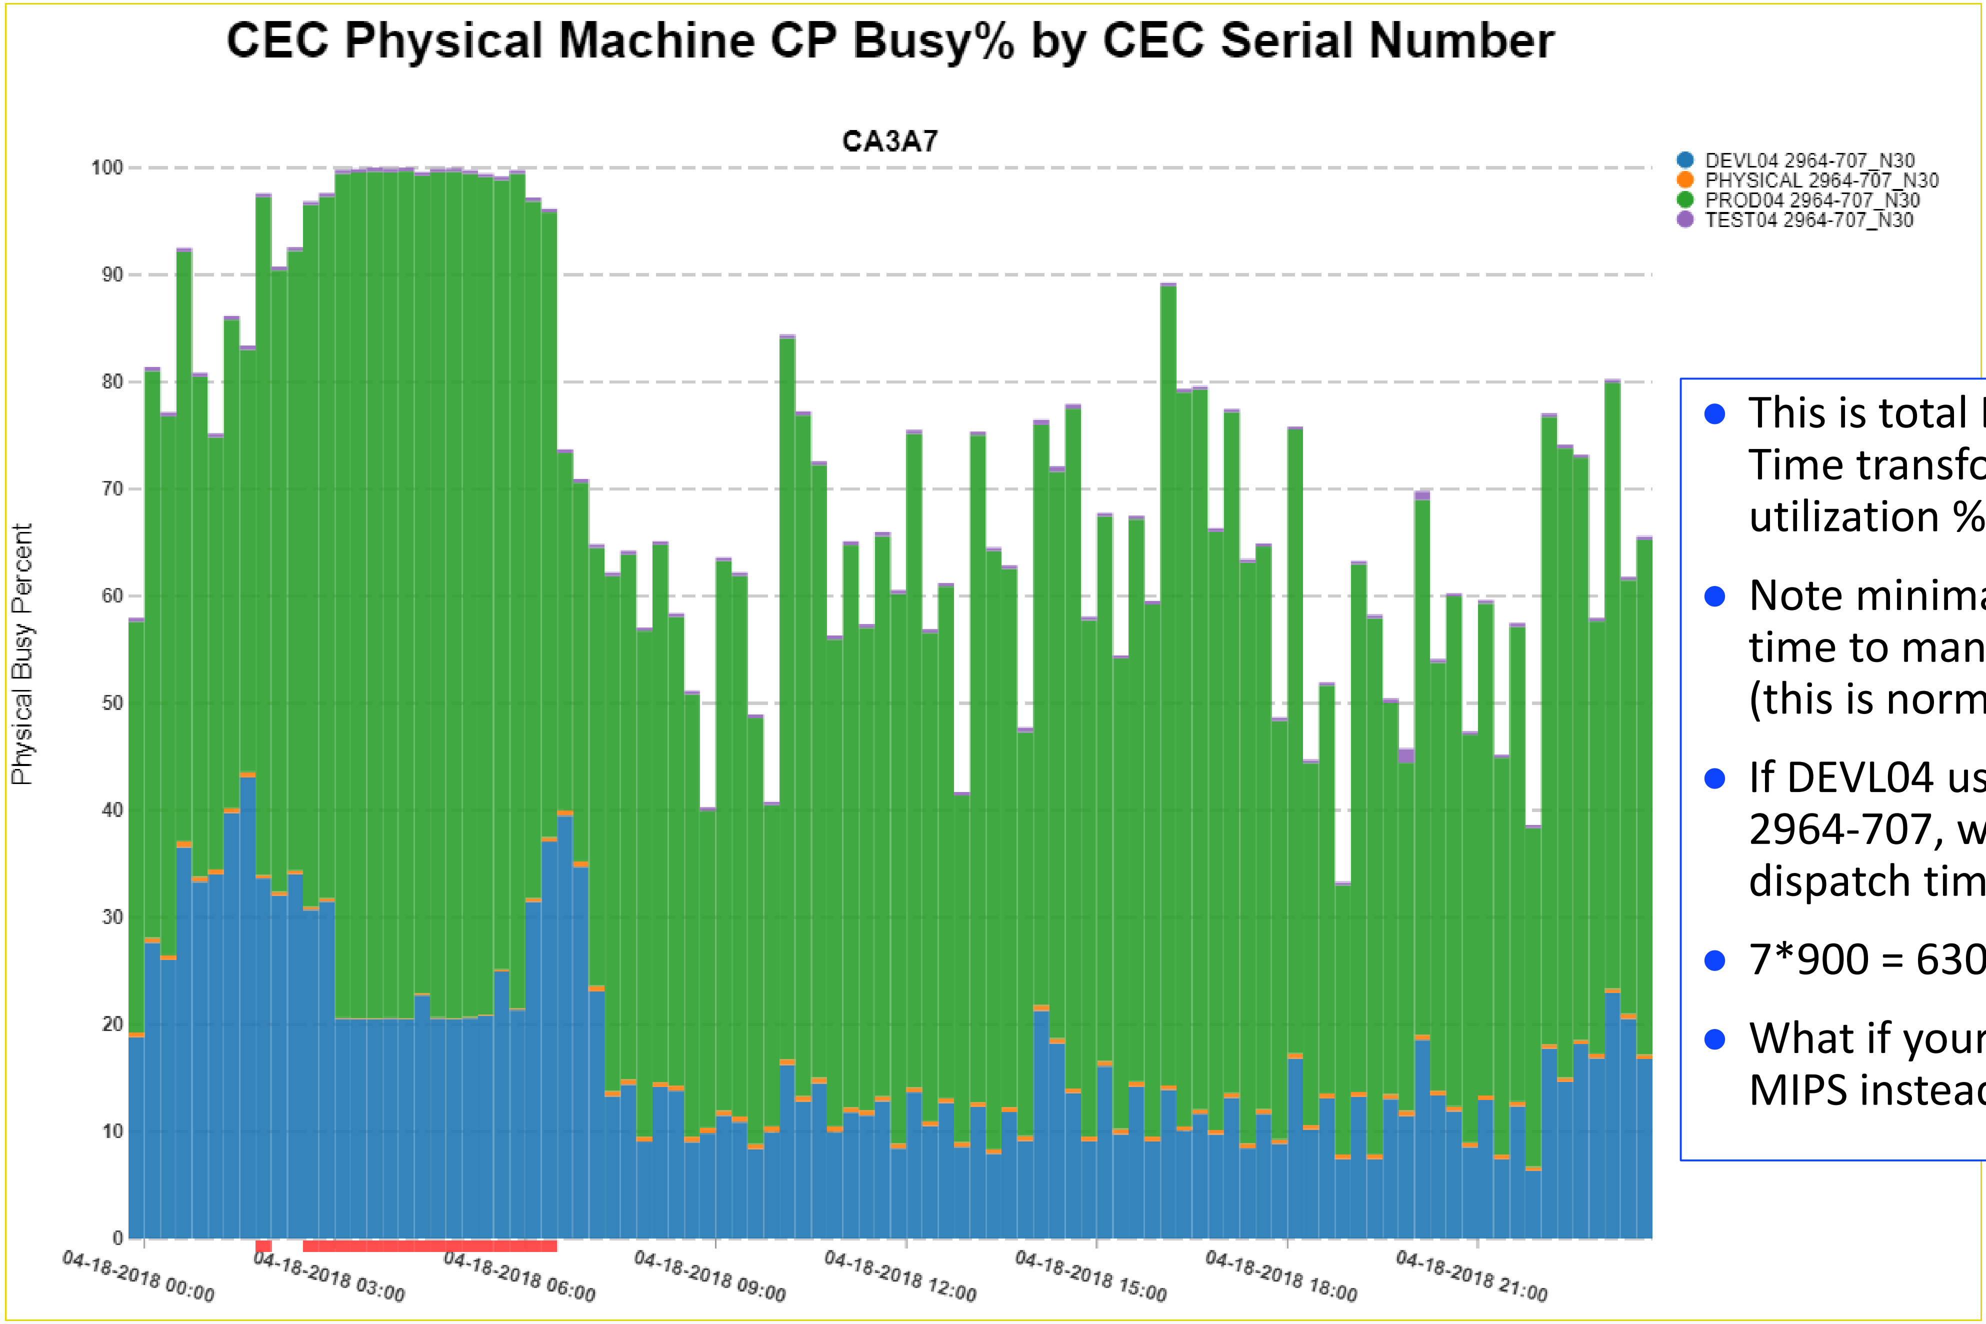

- This is total LPAR Dispatch Time transformed to physical utilization %.
- Note minimal "Physical" time to manage the LPARs (this is normal).
- If DEVL04 used 20% of the 2964-707, what was it's dispatch time?
- 7*900 = 6300 * .2 = 1260
- What if your manager wants MIPS instead?

CEC Physical Machine CP MIPS by CEC Serial Number

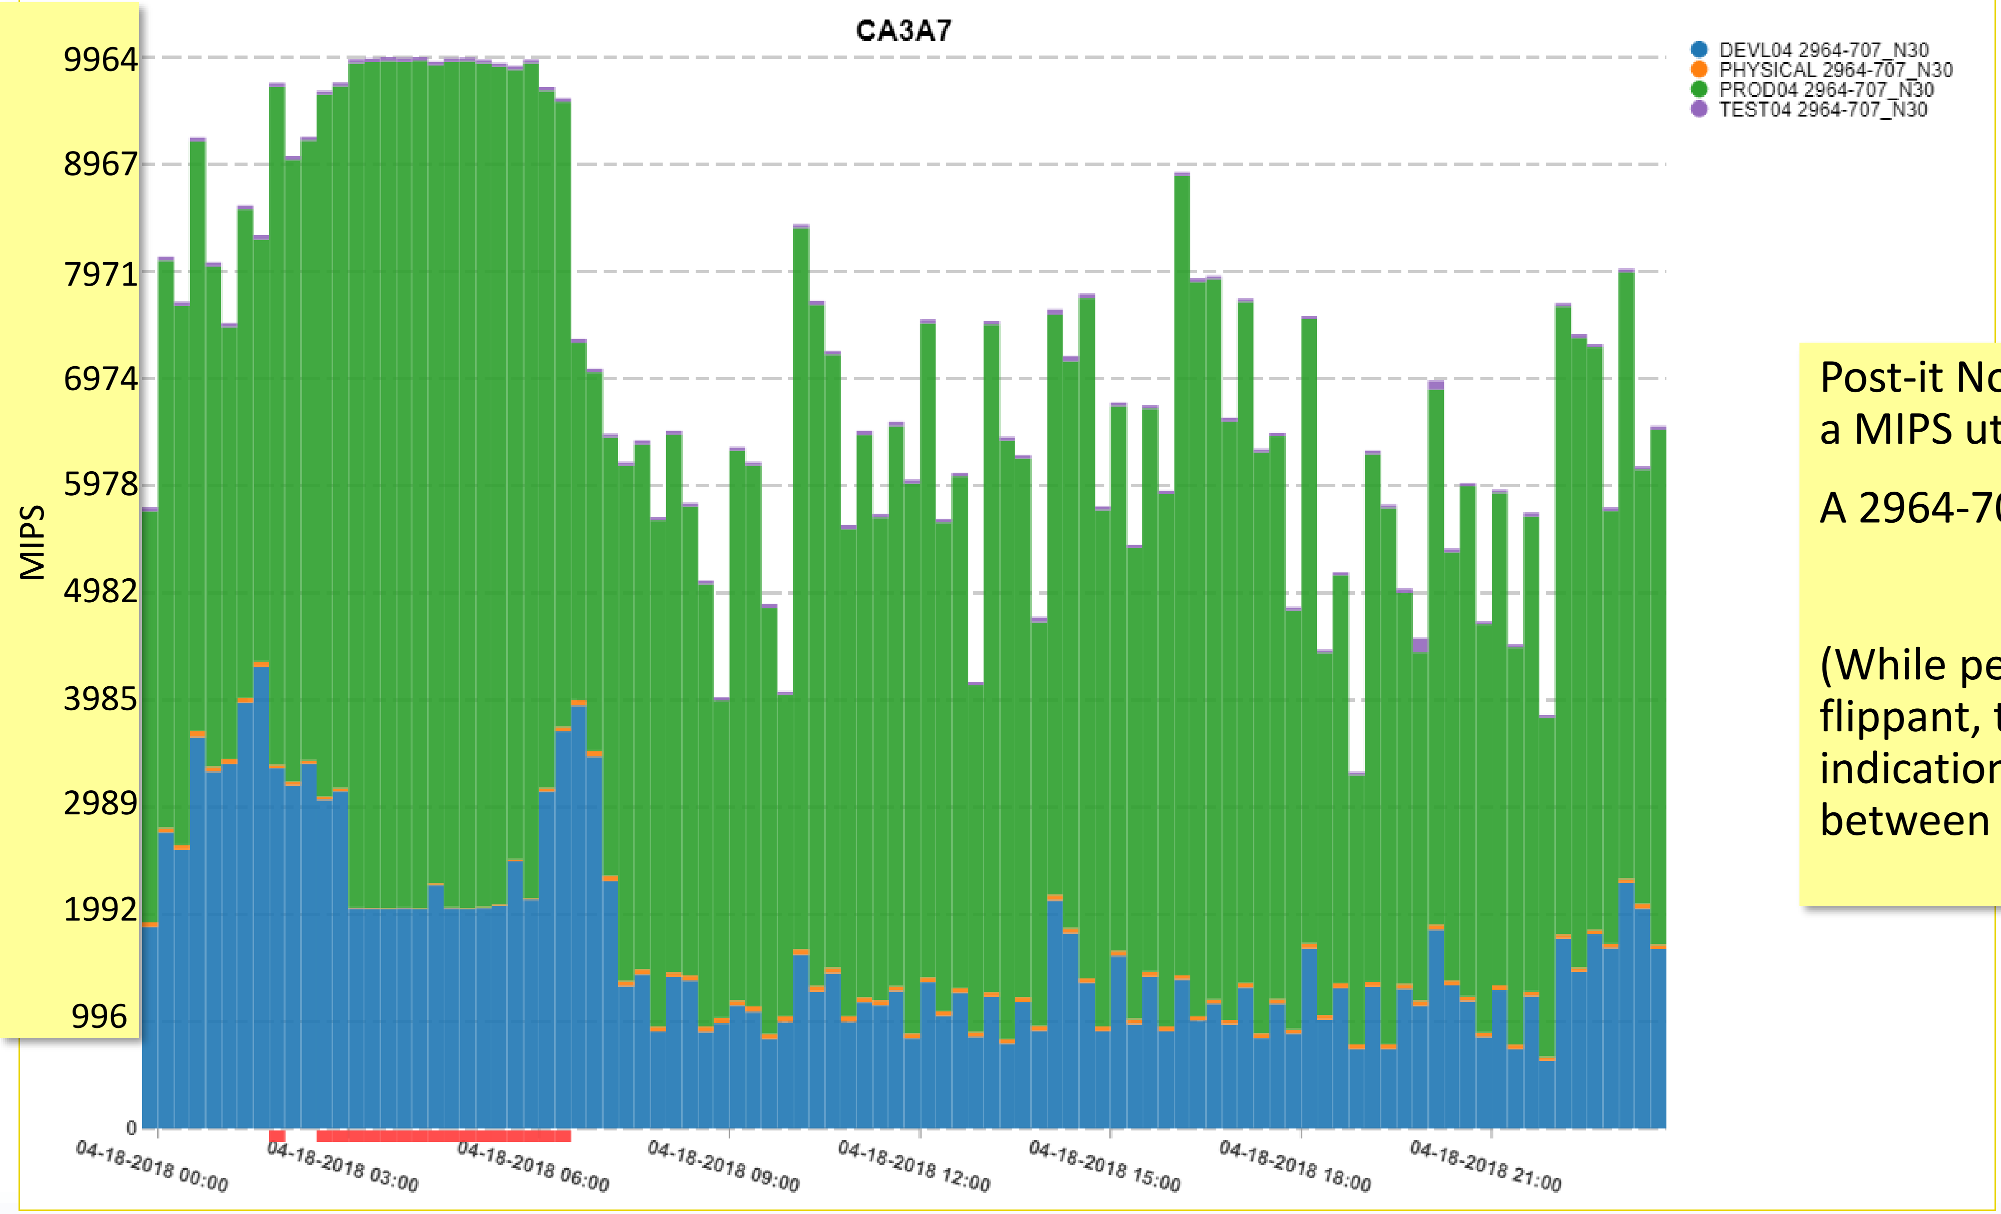

Post-it Note conversion to a MIPS utilization chart!

A 2964-707 is 9964 PCI

(While perhaps a bit flippant, this is an indication of conversion between CPU utilization.)

CEC Physical Machine CP Busy%

General Purpose CPU

470BC, 2964-710, N63

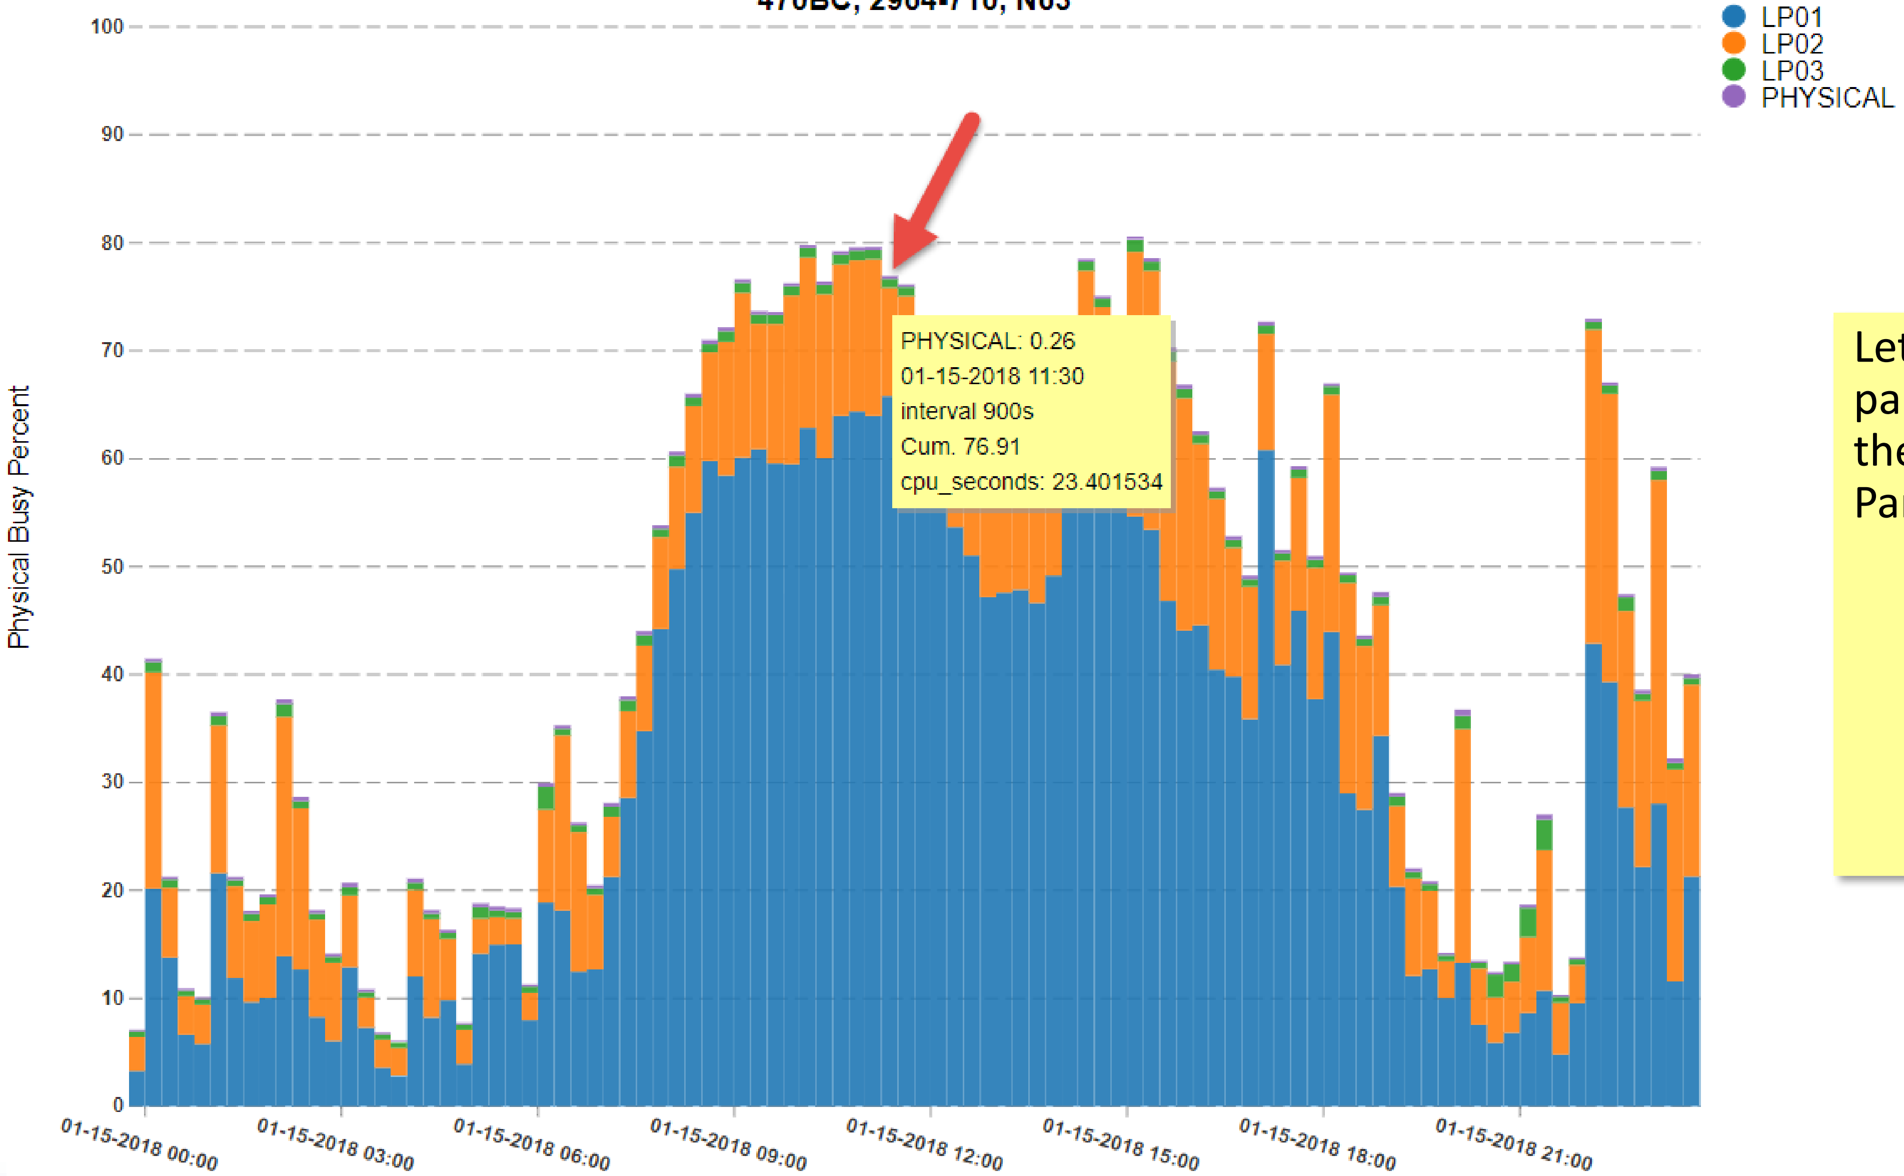

Let’s look at this one particular interval from the RMF PostProcessor Partition Data Report

RMF Partition Data Report

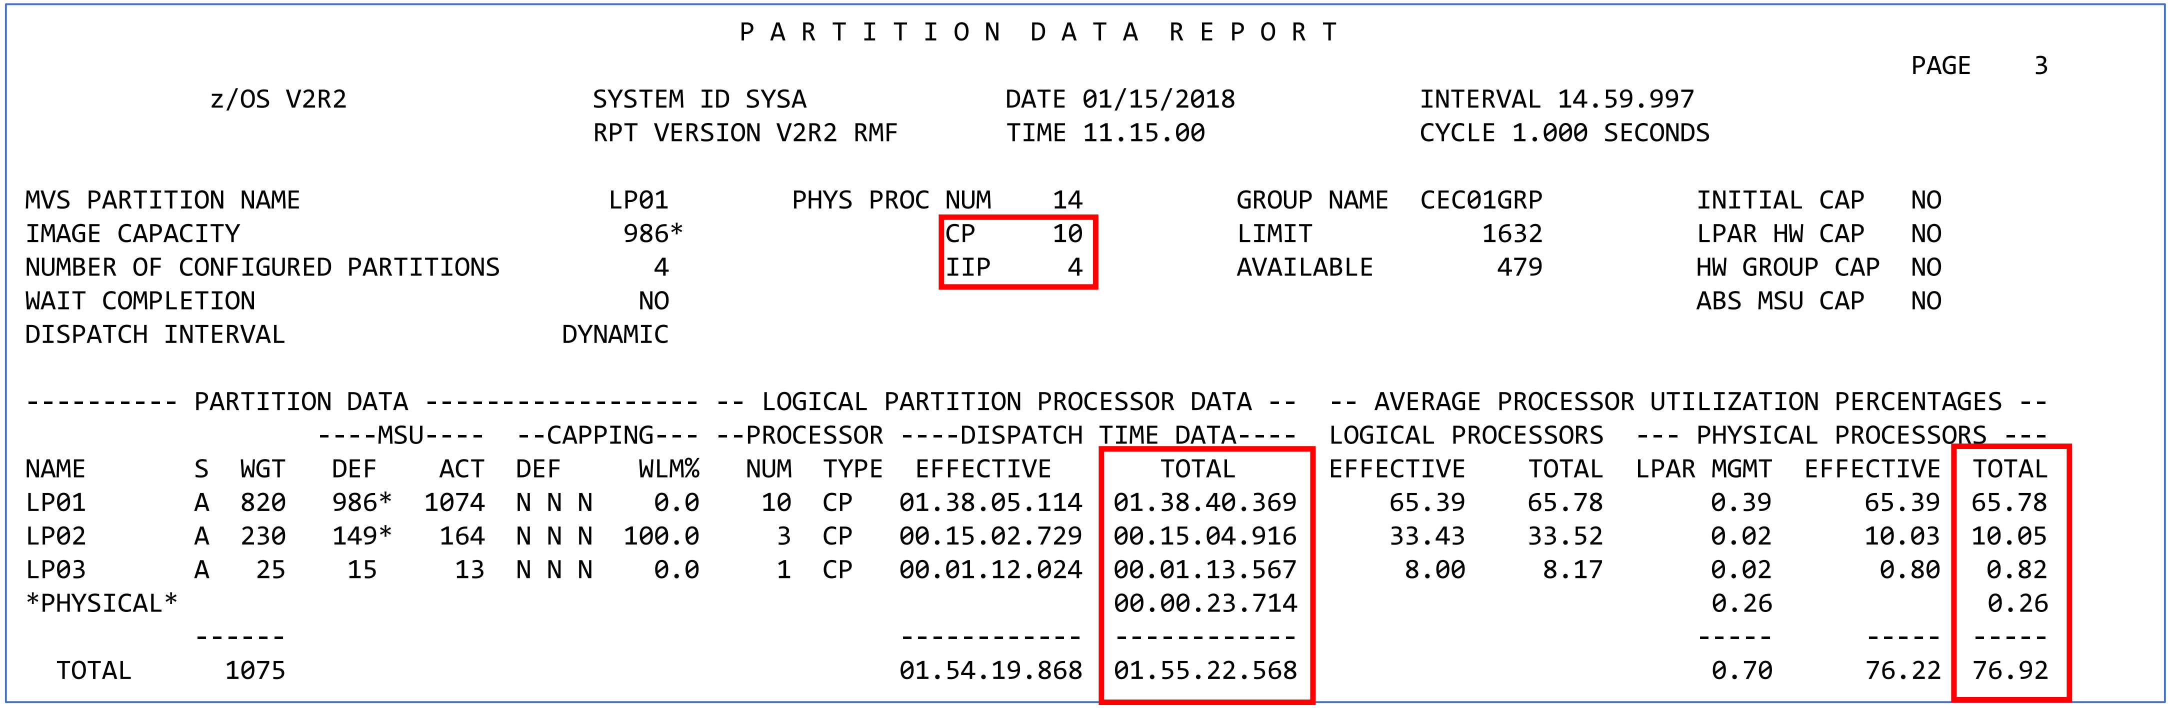

```
                                               P A R T I T I O N  D A T A  R E P O R T
                                                                                                                           PAGE    3
            z/OS V2R2                SYSTEM ID SYSA             DATE 01/15/2018            INTERVAL 14.59.997
                                     RPT VERSION V2R2 RMF       TIME 11.15.00              CYCLE 1.000 SECONDS

MVS PARTITION NAME                    LP01        PHYS PROC NUM    14          GROUP NAME  CEC01GRP          INITIAL CAP   NO
IMAGE CAPACITY                         986*                 CP     10          LIMIT           1632          LPAR HW CAP   NO
NUMBER OF CONFIGURED PARTITIONS          4                  IIP     4          AVAILABLE        479          HW GROUP CAP  NO
WAIT COMPLETION                         NO                                                                   ABS MSU CAP   NO
DISPATCH INTERVAL                  DYNAMIC

---------- PARTITION DATA ------------------ -- LOGICAL PARTITION PROCESSOR DATA --  -- AVERAGE PROCESSOR UTILIZATION PERCENTAGES --
                   ----MSU----  --CAPPING--- --PROCESSOR ----DISPATCH TIME DATA----  LOGICAL PROCESSORS  --- PHYSICAL PROCESSORS ---
NAME       S  WGT   DEF    ACT  DEF     WLM%   NUM  TYPE  EFFECTIVE       TOTAL      EFFECTIVE    TOTAL  LPAR MGMT  EFFECTIVE  TOTAL
LP01       A  820   986*  1074  N N N    0.0    10  CP   01.38.05.114  01.38.40.369      65.39    65.78       0.39      65.39  65.78
LP02       A  230   149*   164  N N N  100.0     3  CP   00.15.02.729  00.15.04.916      33.43    33.52       0.02      10.03  10.05
LP03       A   25    15     13  N N N    0.0     1  CP   00.01.12.024  00.01.13.567       8.00     8.17       0.02       0.80   0.82
*PHYSICAL*                                                             00.00.23.714                           0.26              0.26
           ------                                        ------------  ------------                          -----      -----  -----
  TOTAL      1075                                        01.54.19.868  01.55.22.568                           0.70      76.22  76.92
```

- Official interval duration is 899.997 seconds * 10 CPs = up to 8999.97 CPU seconds available in the interval
- Total of 1*3600+55*60+22.568 seconds consumed = 6922.568
- 6922.568 / 8999.97 = 76.917%
- This is a 2964-710, rated at 1632 MSUs, so 1632 * .76917 = 1255.29 MSUs used

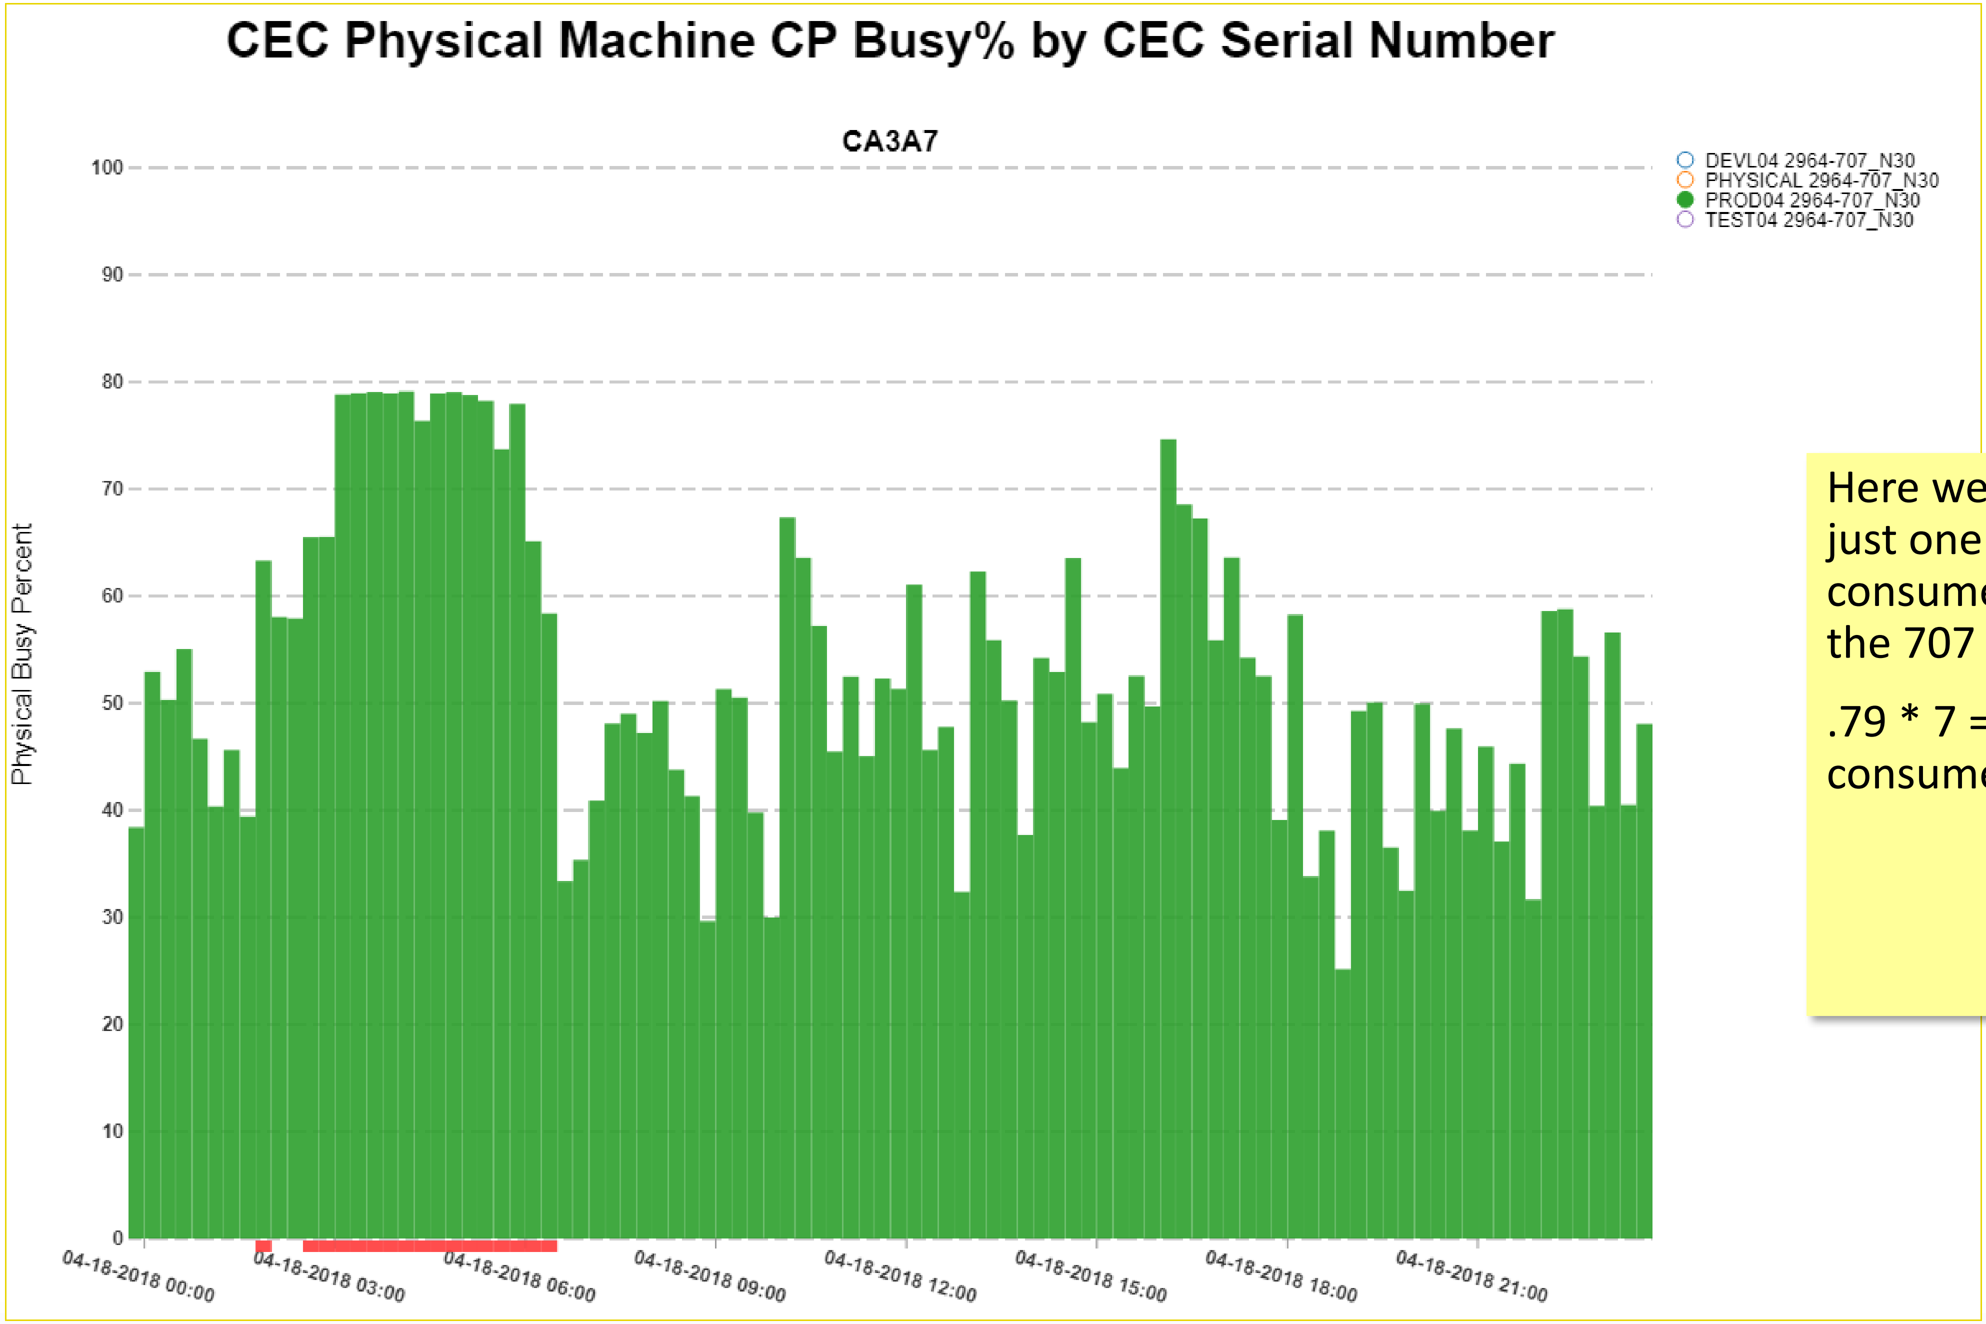

Here we've focused on just one LPAR, which consumed about 79% of the 707 at its peak.

.79 * 7 = 5.53 CPs consumed at it's peak

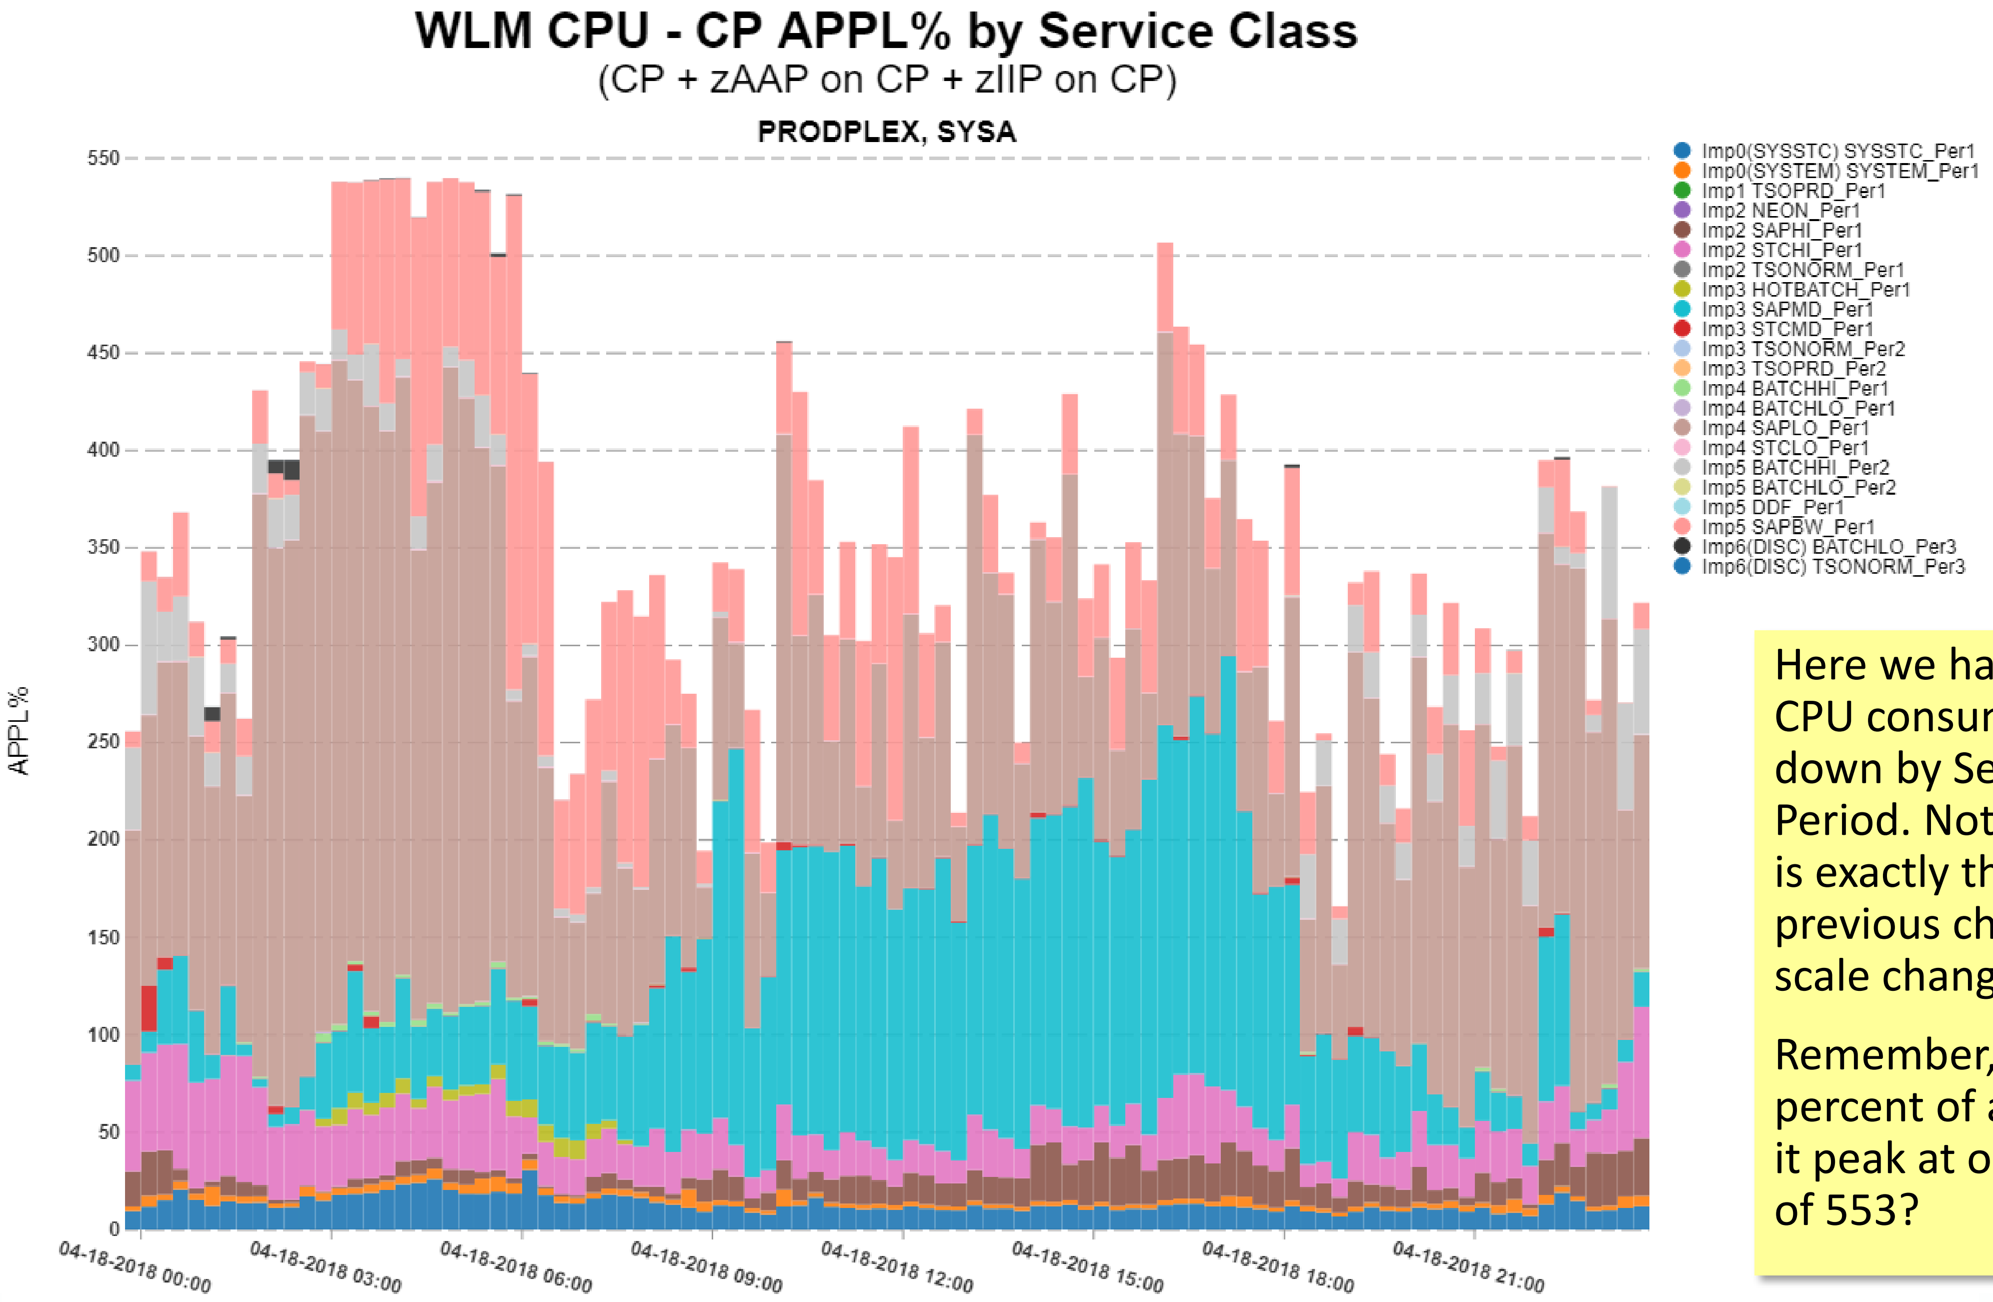

Here we have that LPAR's CPU consumption broken down by Service Class Period. Notice the profile is exactly the same as the previous chart just the scale changed.

Remember, APPL% is percent of a CP. Why does it peak at only 540 instead of 553?

Capture Ratio

- Captured time = CPU time attributed to particular service class periods in the type 72 records
- Capture Ratio = $\frac{Captured\ time}{Effective\ dispatch\ time}$
- Uncaptured time = $(Captured\ time) - (Effective\ dispatch\ time)$
- Some uncaptured time is normal and is because z/OS can't directly attribute some system-related work to a specific workload
  - Interrupts
  - Contention
  - Paging
  - Periodic system management
  - SLIP traps and system traces

Target Capture Ratio: >90%

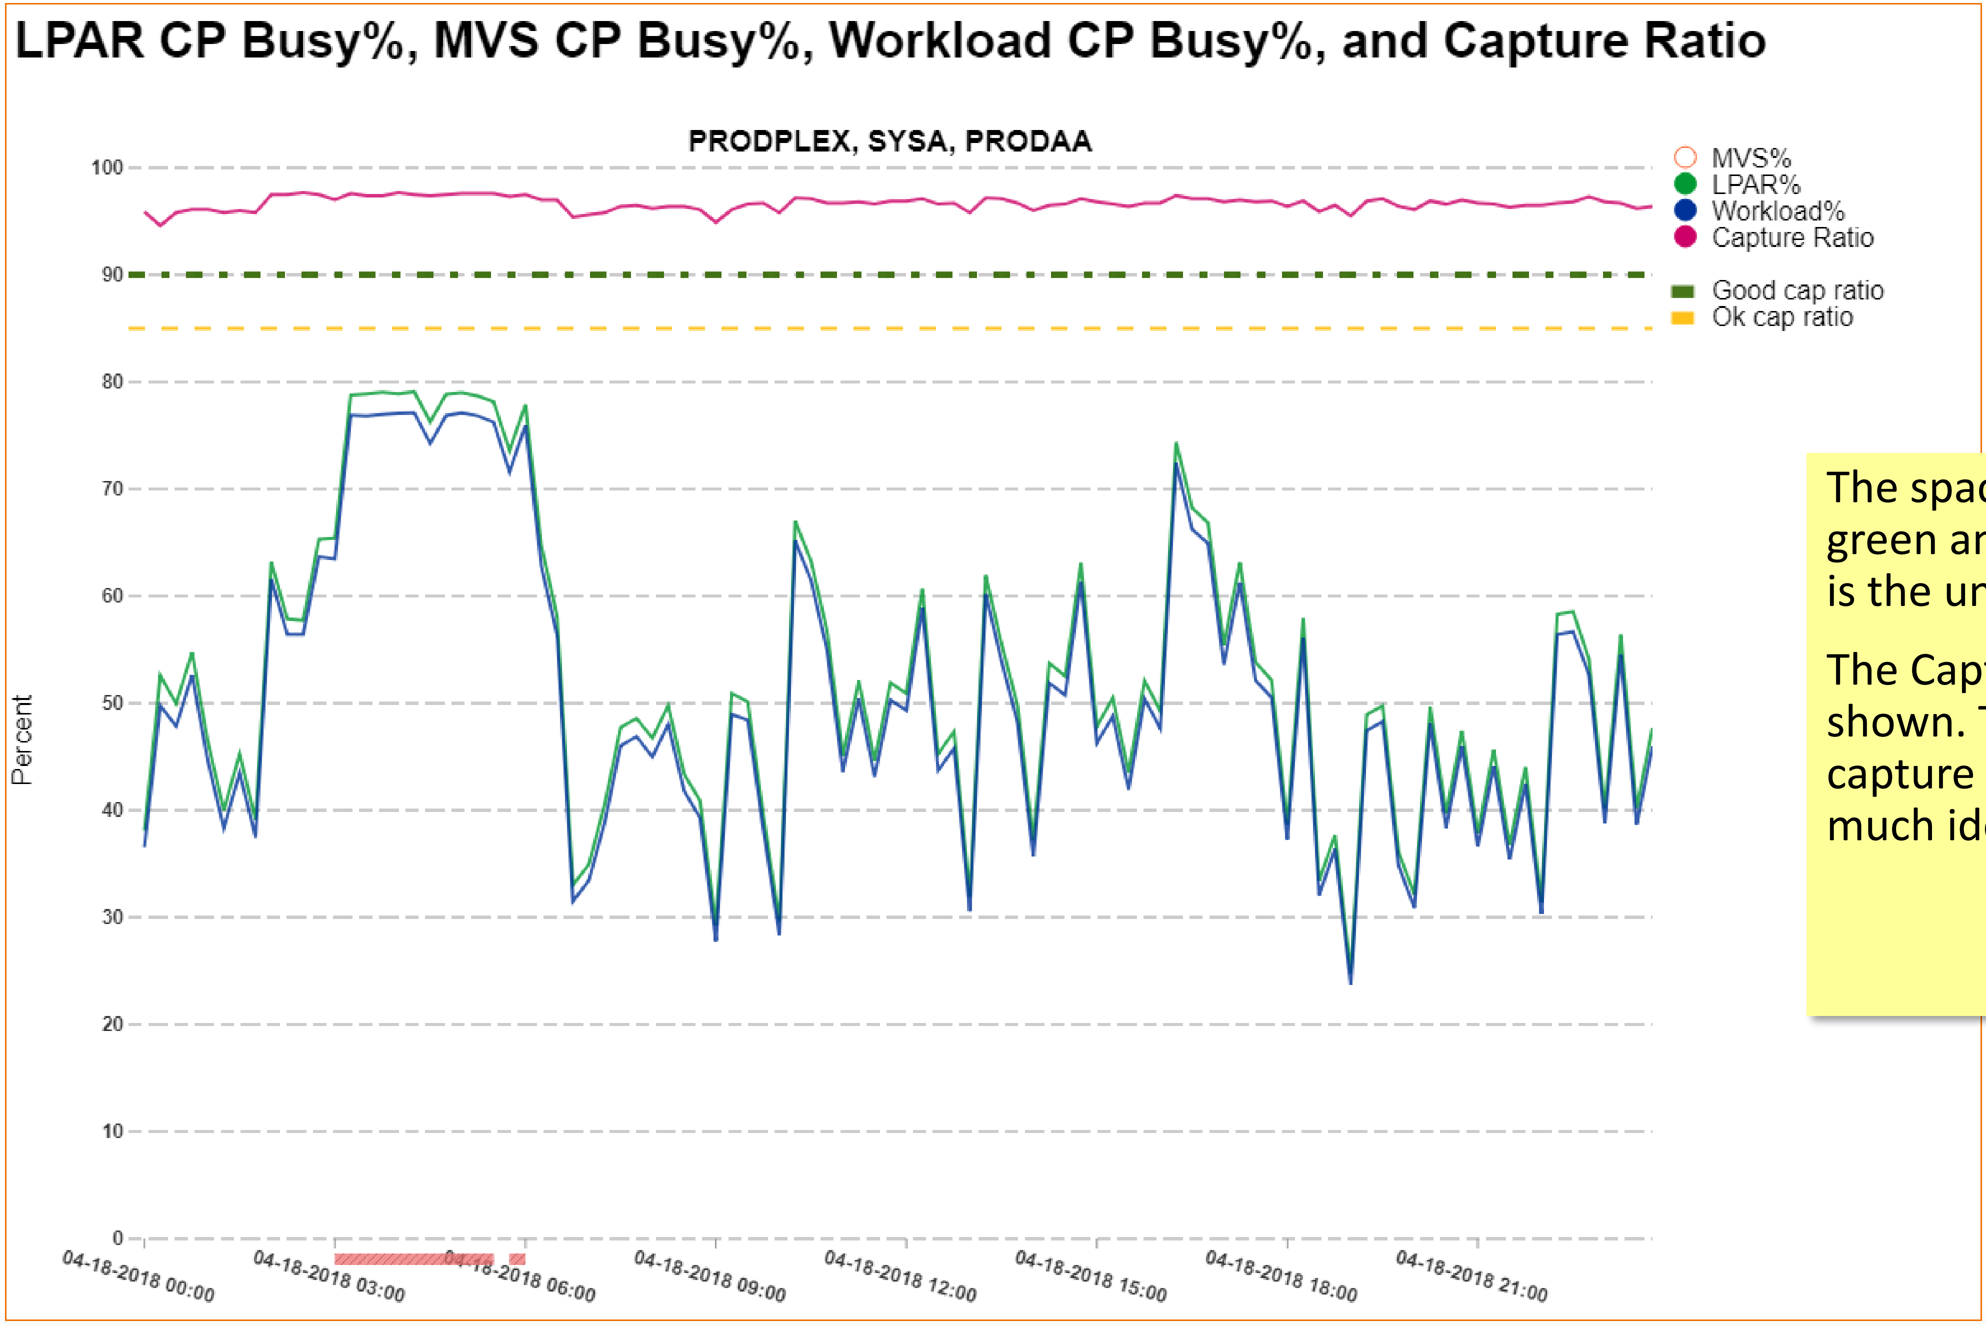

The space between the green and the blue lines is the uncaptured time.

The Capture Ratio is also shown. This system's capture ratio is pretty much ideal.

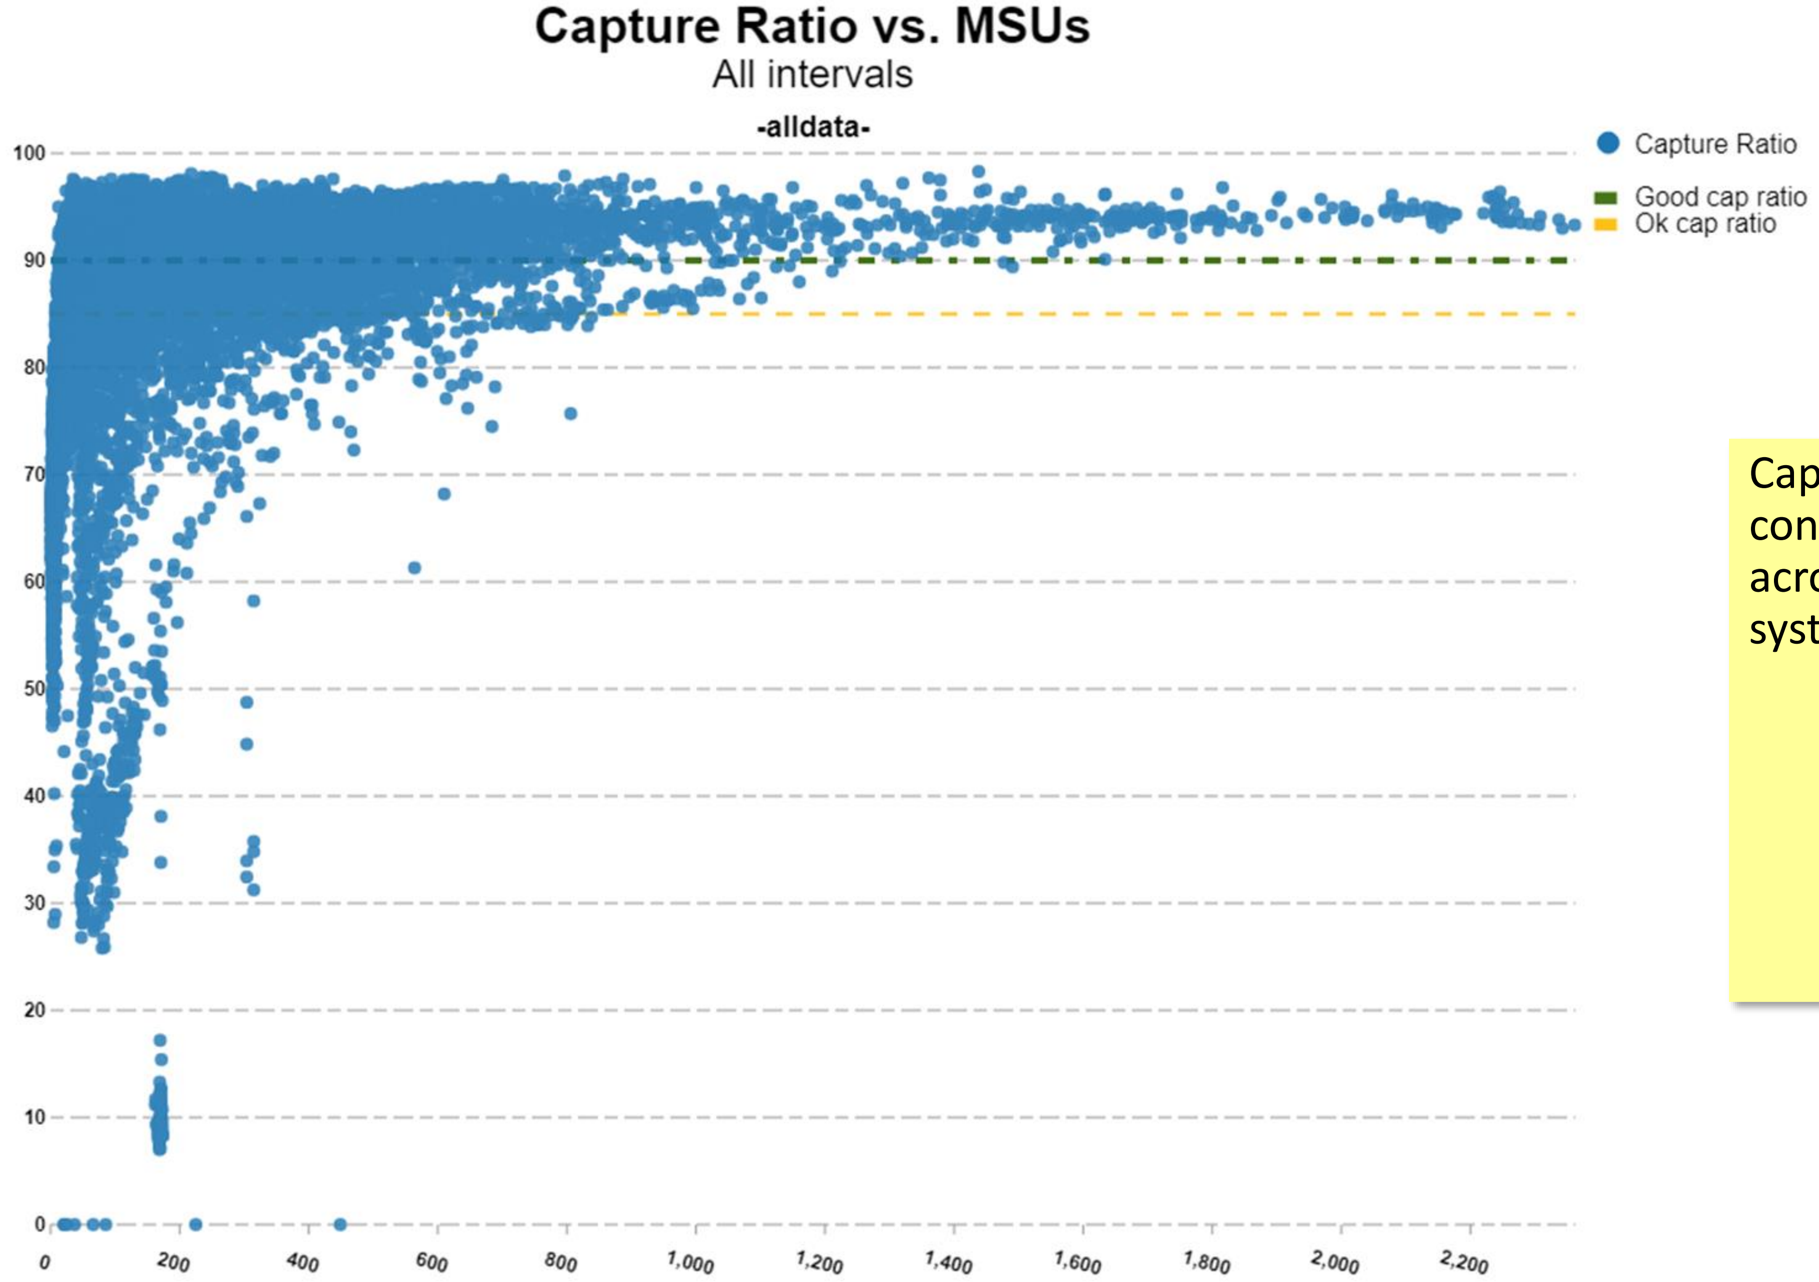

Capture ratio vs. MSUs consumed in interval across many different systems and intervals

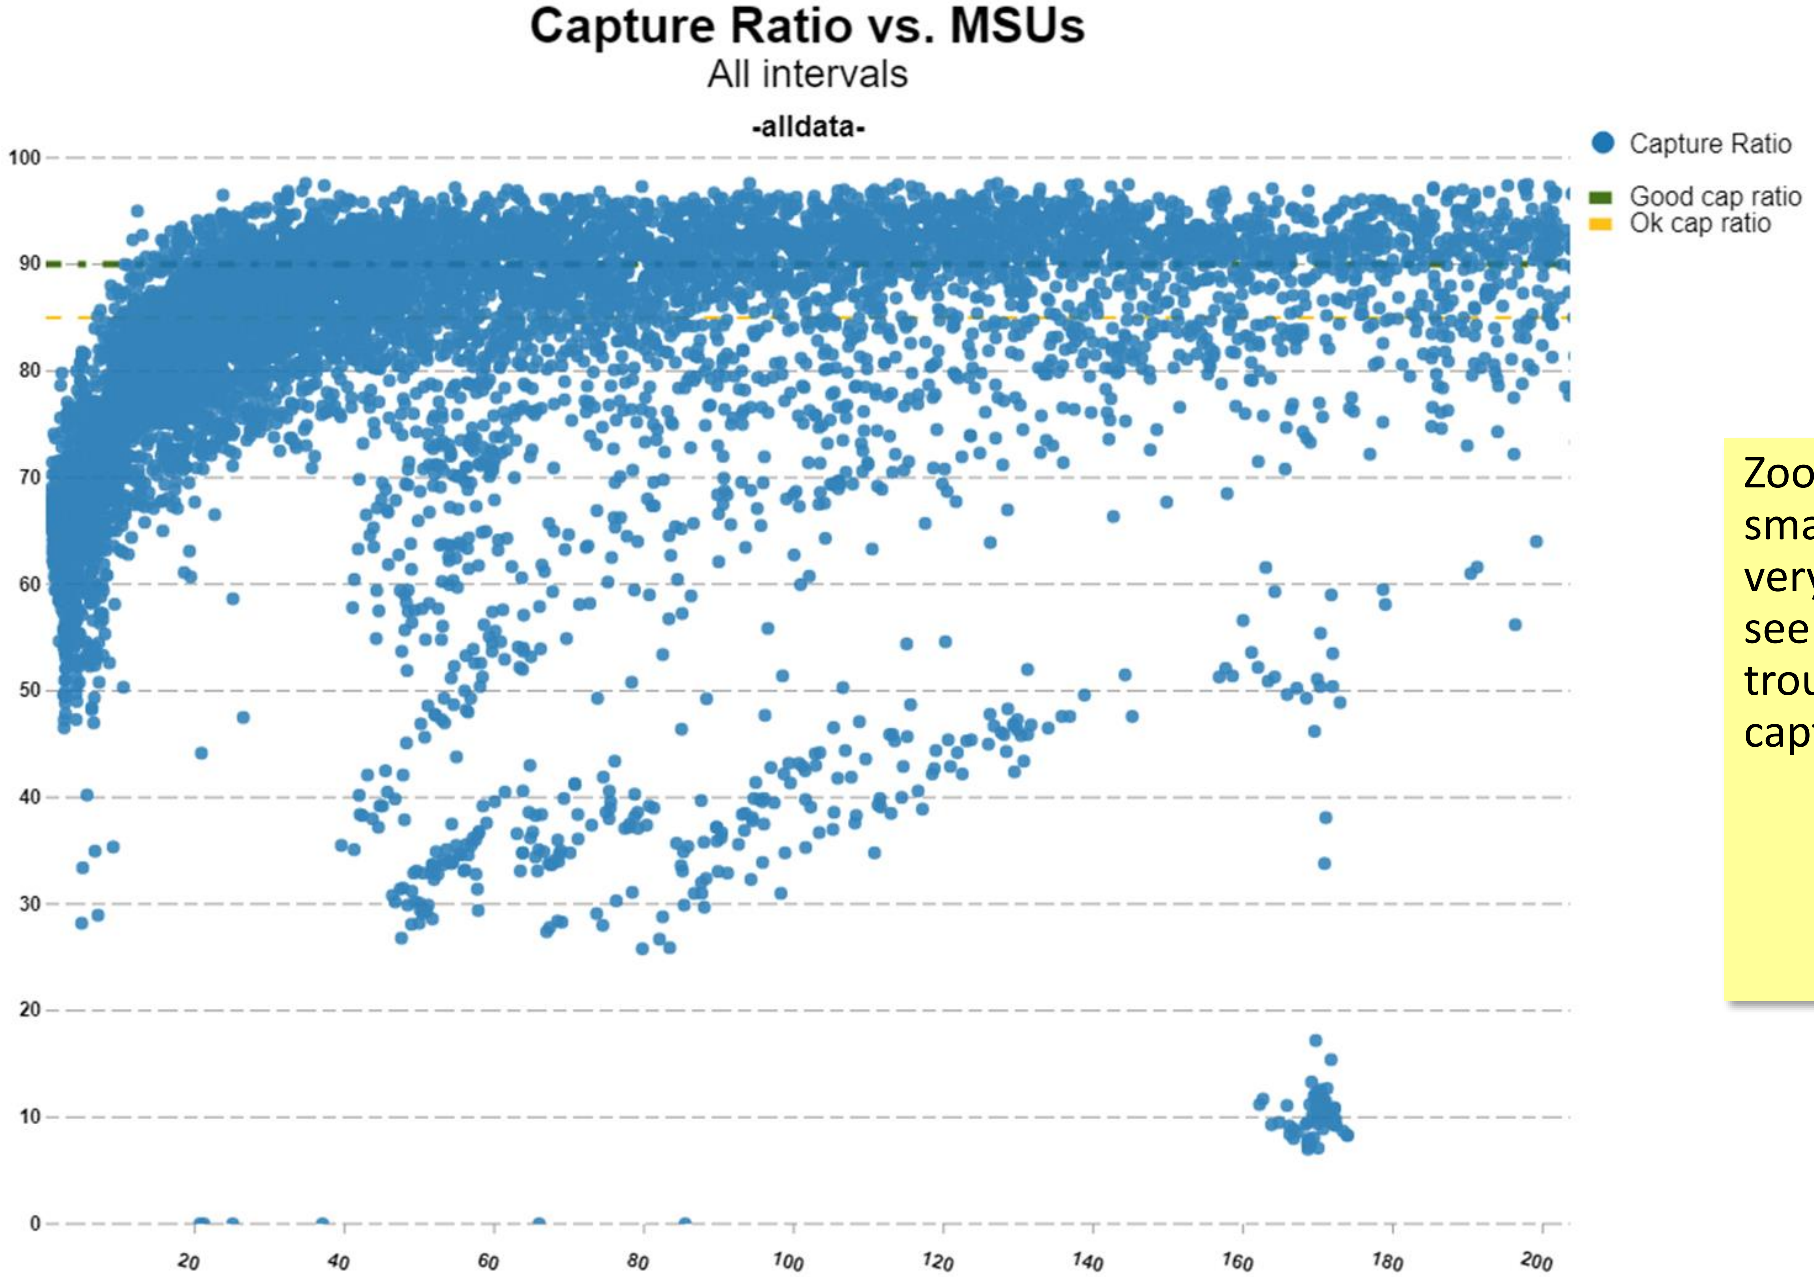

Zoomed in to see smaller systems better: very small systems to do seem to have more trouble achieving high capture ratios

Evaluating Capture Ratio

- If capture ratio is above 90%: no concern, probably little tuning opportunity
- Some systems clearly run relatively low capture ratios all the time
  - These most often are lightly loaded systems
- If a capture ratio is regularly below 85% and the system is a significant contributor to your software costs, I would look into it
  - Look for old traces or SLIP traps that may have been left enabled
  - Do you have RMPTTOM set too low in IEAOPTxx?
    - Also may be able to increase it above default (usually 3000) for lightly used systems
  - Look for paging activity
  - Look for contention
  - Unfortunately, tracking down uncaptured time is usually difficult
    - Because it is “uncaptured”
  - Do you have large SMS pools?

zIIPs (also zAAPs on older machines)

- zIIPs always run at “full” speed, i.e. the equivalent as a 7xx or Zxx model
- Sometimes we’d like to compare zIIP consumption to the GCPs
- “Normalization factor” in (e.g.) R723NFFS provide a conversion from zIIP to GCP
  - Multiply by NFFS and divide by 256 to get GCP equivalent
  - In this case (2965-P02) 766: x * 766 / 256
- Normalization factor is ratio of SU/sec of the single-CP sub-capacity and full-capacity models
  - E.G. P01 vs Z01
  - Not expected to perfectly match the cycle time ratio from the SMF 113s

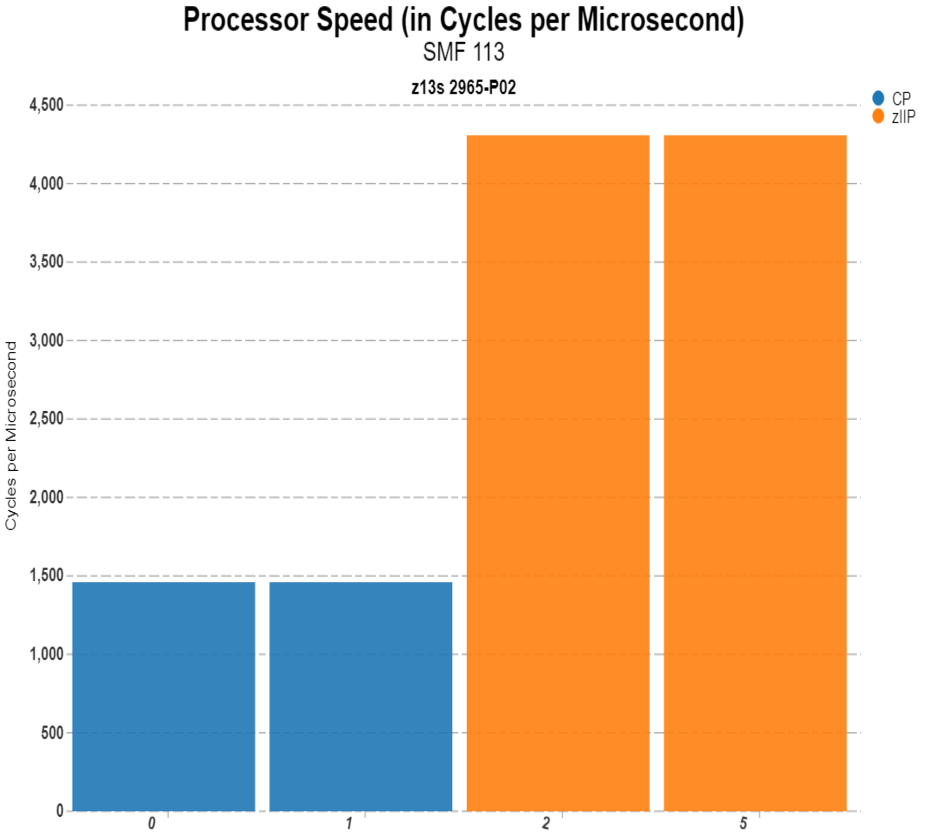

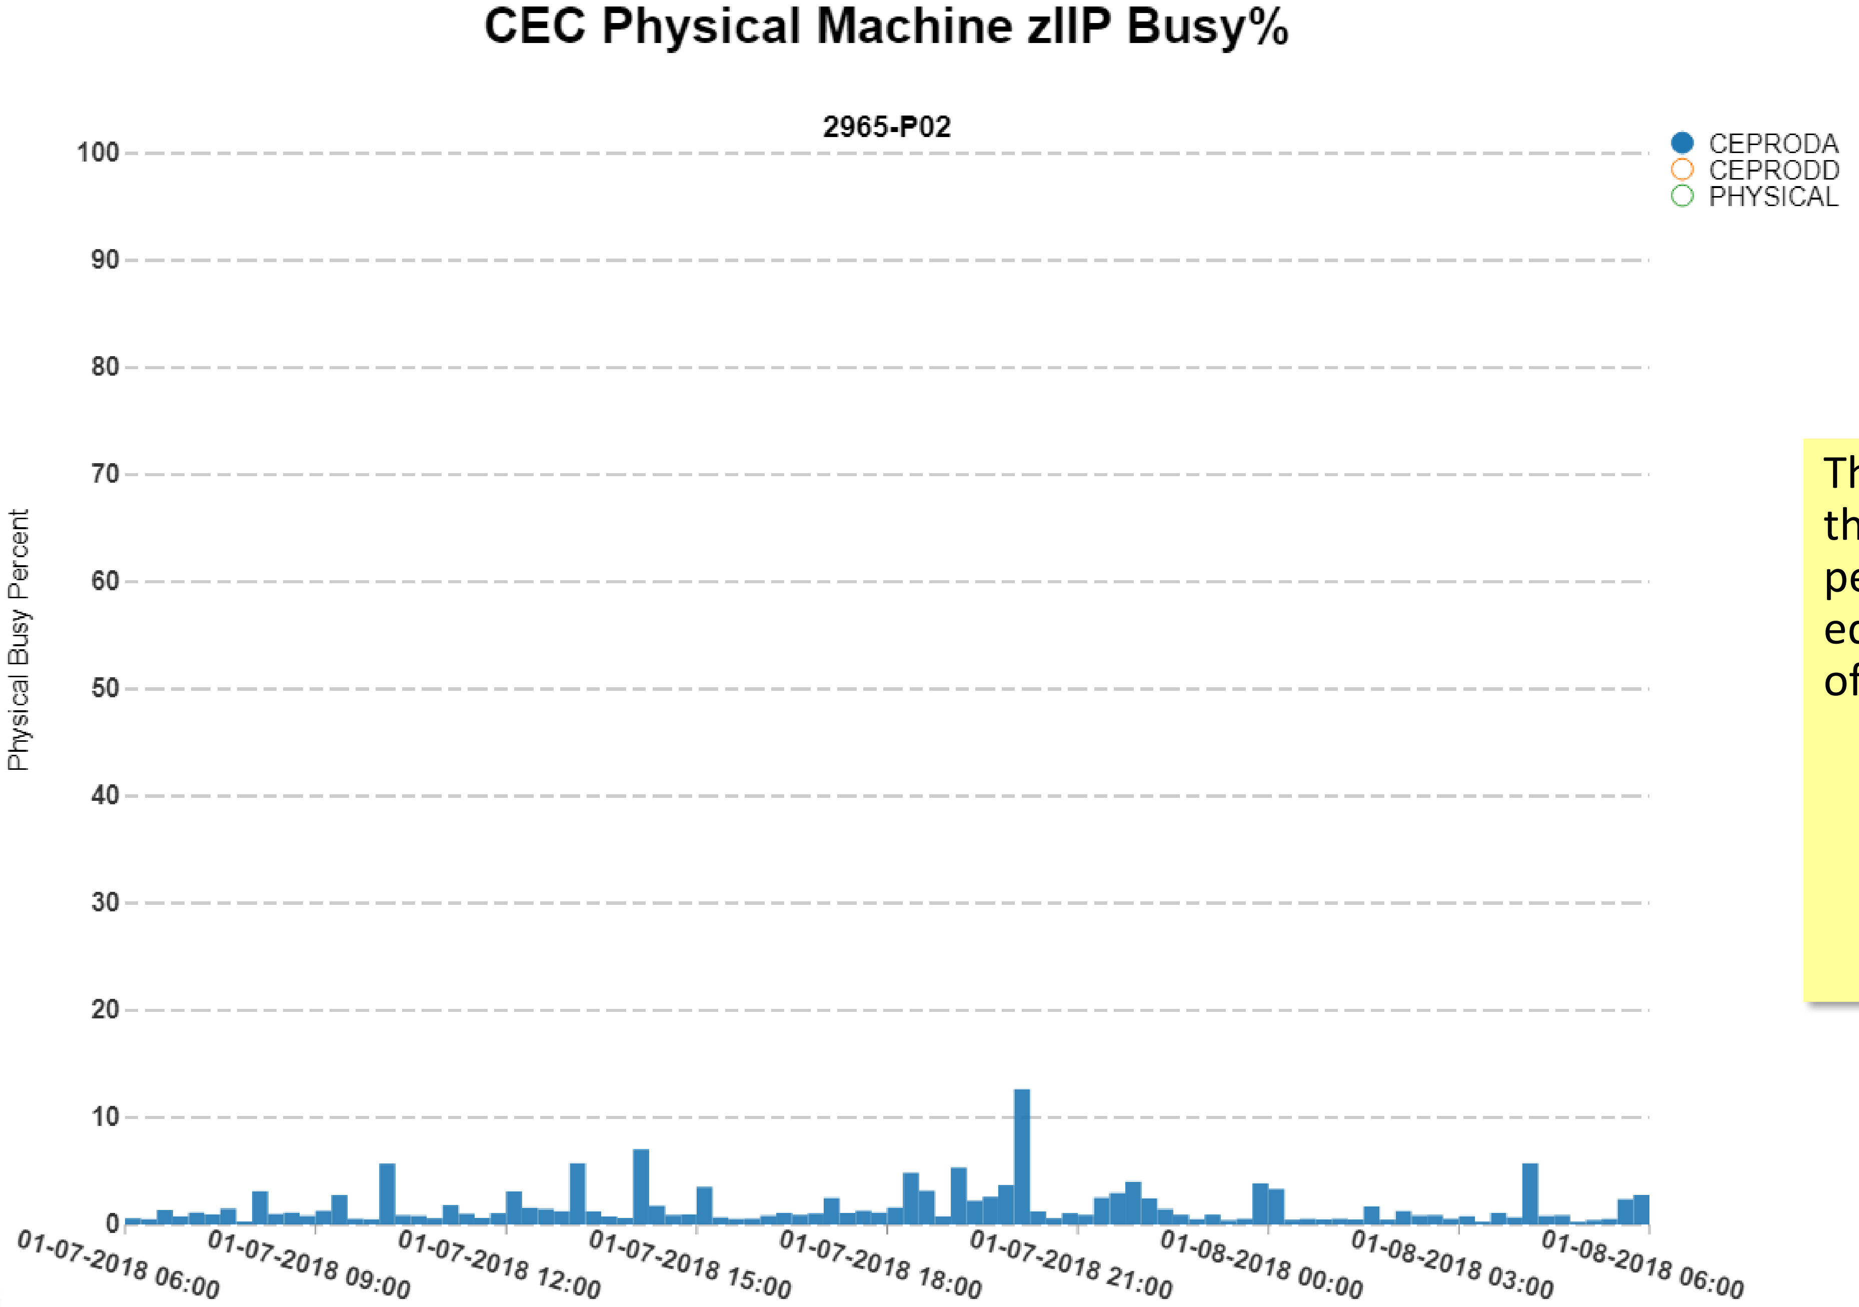

There are two zIIPs on this machine. So that peak utilization of 12.6% equates to 2*.126 = .252 of a single zIIP.

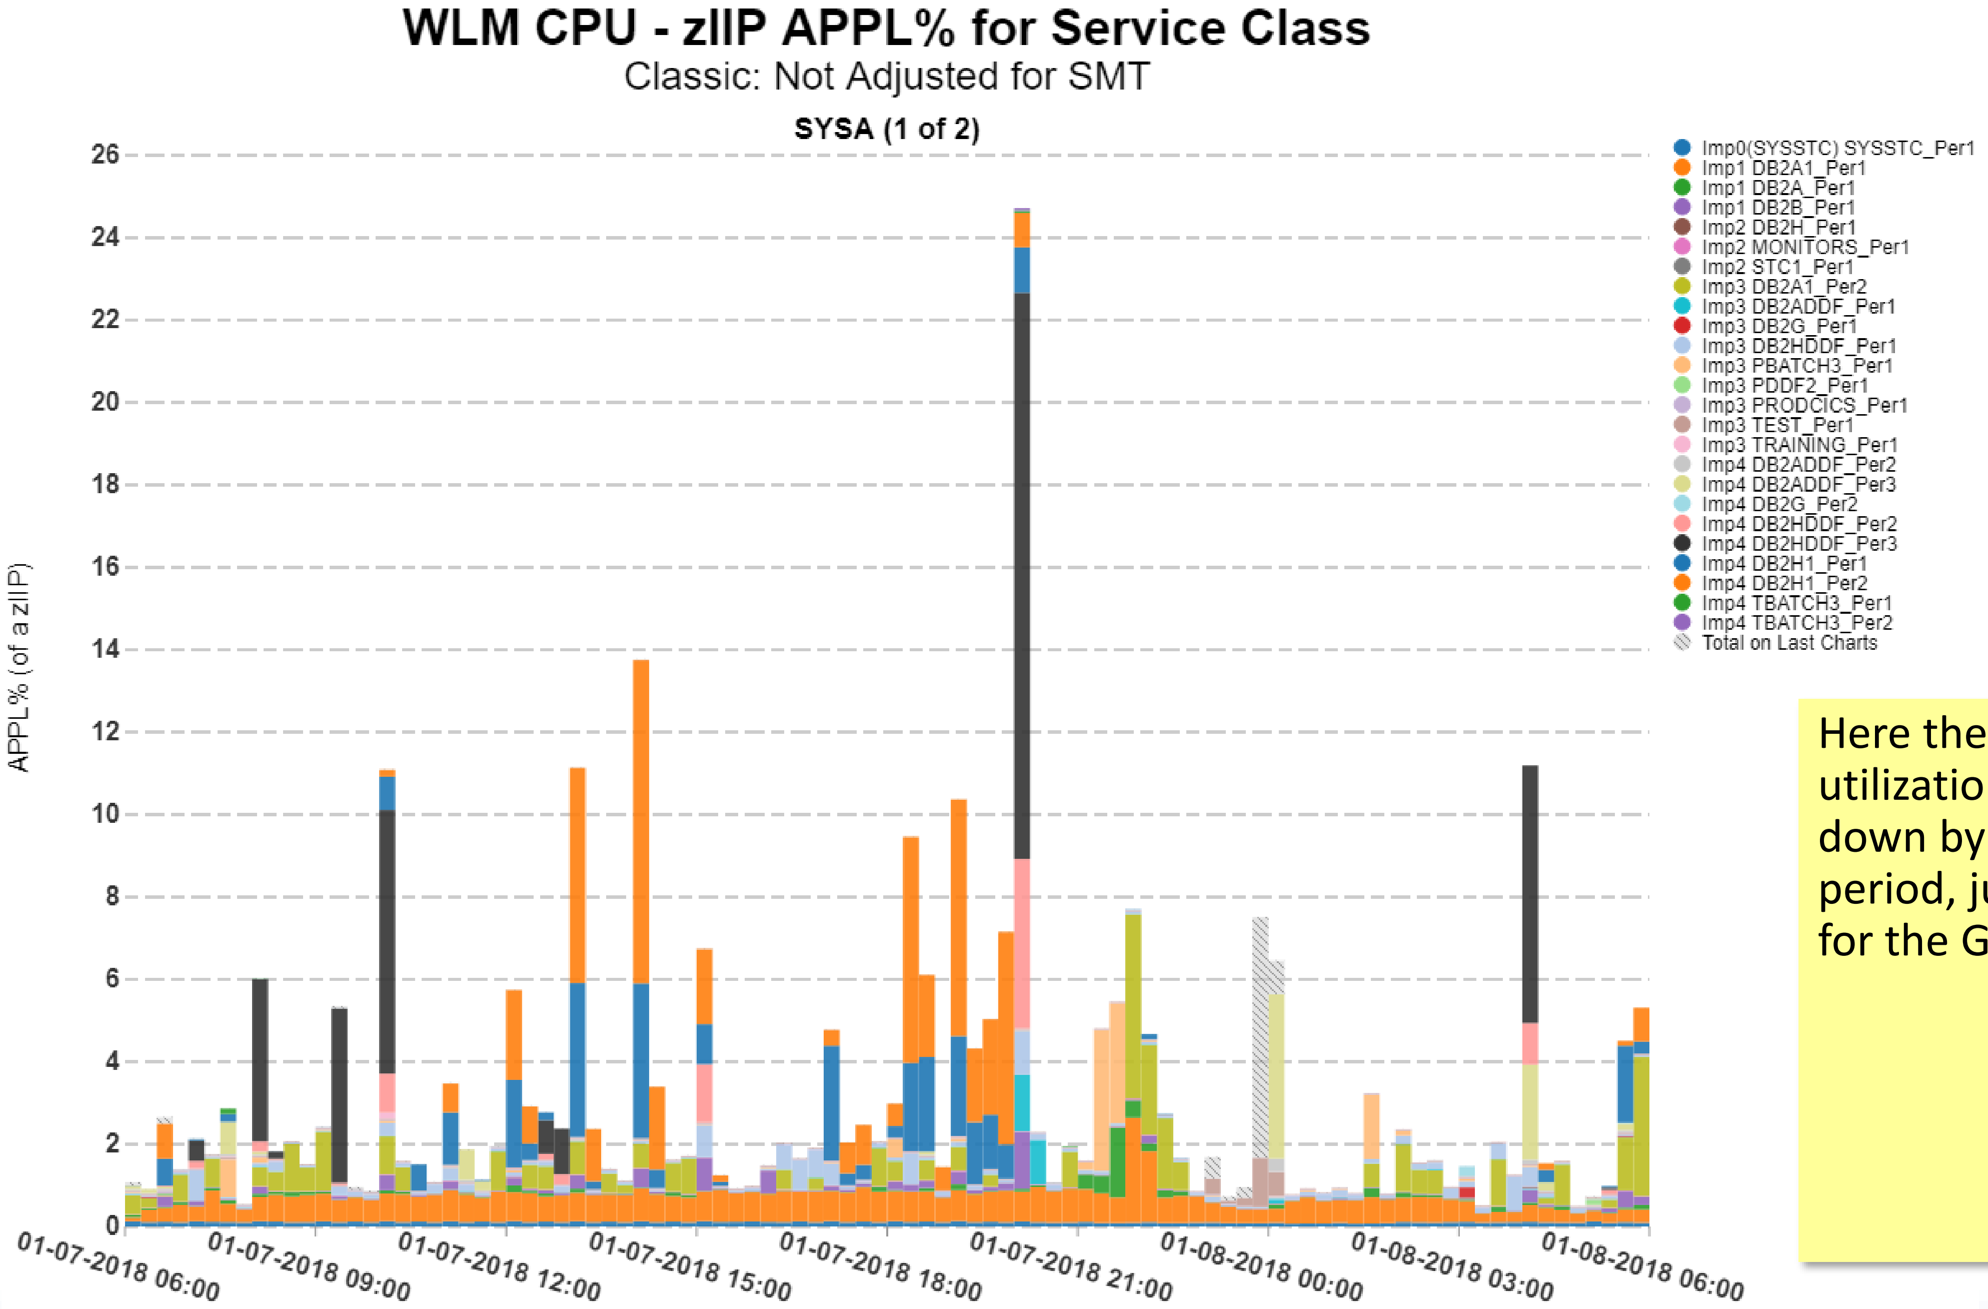

WLM CPU - zIIP APPL% for Service Class

Classic: Not Adjusted for SMT

SYSA (1 of 2)

Here the physical utilization is broken down by service class period, just like we did for the GCPs.

WLM CPU - zIIP APPL% GCP for Service Class

Classic: Not Adjusted for SMT

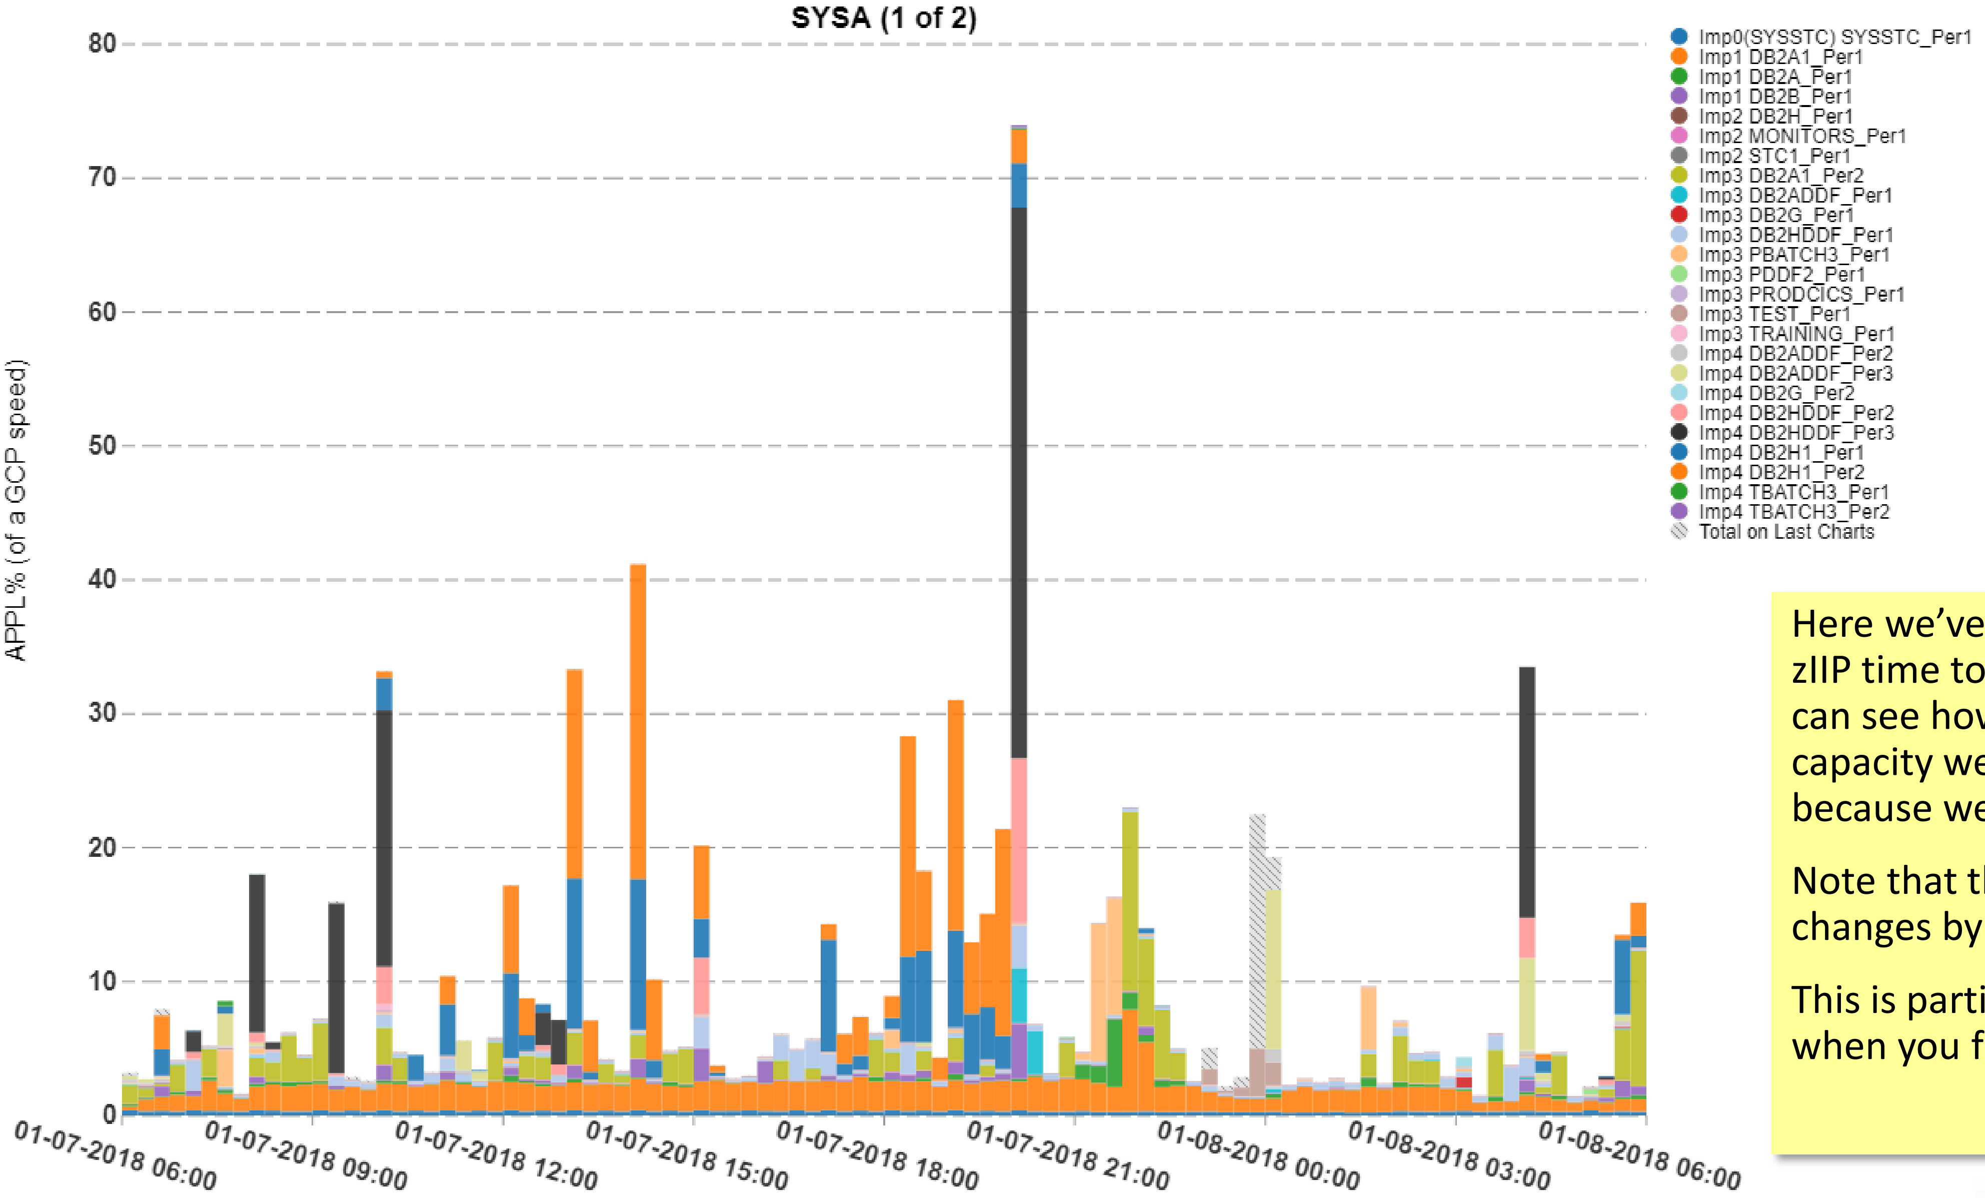

Here we've normalized the zIIP time to GCP time so we can see how much GCP capacity we're saving because we have zIIPs.

Note that the scale just changes by 766/256.

This is particularly interesting when you first add zIIP.

Logical vs. Physical Measurements

- What we've discussed so far have been physical measurements
  - That is, they come from a direct measurement of CPU time on a physical CP
  - Even though all the work is running on logical CPs, the CPU time measurements are physical measurements because only one LPAR (and one workload, barring SMT) can use a physical CP at a time
  - PR/SM time slicing is responsible for dispatching logical CPs to physical CPs
  - z/OS is responsible for dispatching logical CPs to workloads
- Sometimes, we might want to look at certain logical measurements
  - Primary reason is to understand how our logical capacity (CPs and weight) is being consumed
    - Generally less necessary today with most LPARs having 1 or 2 more logical CPs than minimally necessary and allowing HiperDispatch to manage them
  - Logical measurements are based on the number of logical CPs instead of physical CPs, but for time reasons, I'm not going to get into them

How many MIPS is that batch job?

- Essentially this question is: how much of my capacity is a particular job using?
- The answer is easy to calculate but depends on the time frame you want to consider:
  - Over the duration of the batch job?
  - Over the duration of the RMF interval? SMF Type 30 interval?
  - Over the duration of the hour?
  - Over the duration of the day?

Example Batch Job MIPS

- Example:
  - Job Z starts at 11:10 and runs for 14 minutes. The first interval record cut at 11:15 shows 3 minutes of CPU time. The second (and final interval) shows 4 minutes of CPU time, over that final 9 minutes of elapsed time. Total CPU time was 7 minutes. This is a 5 way machine rated at 840 MIPS.

CPUs: 5	MIPS: 840

Description	Duration	Max Possible CPU Time	Recorded CPU Time	Recorded / Max = Pct Total Capacity	Consumed MIPS
Course of Batch Job	14	70	7	10%	84.0
SMF 30 Interval 1	5	25	3	12%	100.8
SMF 30 Interval 2	9	45	4	9%	74.7
RMF Interval 1	15	75	3	4%	33.6
RMF Interval 2	15	75	4	5%	44.8
Hour	60	300	7	2%	19.6
Day	1440	7200	7	0%	0.8

Over the course of the batch job seems to make the most sense, but don't add those consumed MIPS up for all batch jobs over the course of an hour or day because you'll quite possibly end up with more than your total installed MIPS!

Summary

- CPU time consumed is the basis for understanding how much of your CPU capacity is consumed
- CPU capacity for a given machine can be described by various CPU metrics such as SU/sec, MSUs, MIPS, and PCI
- Convert from CPU time to those other CPU metrics by multiplying the percent of the total capacity consumed by the specific capacity rating for the machine
- Capacity consumed depends on the CPU time consumed during the specific interval you're concerned with
  - Same CPU time over larger interval = lower percentage of total capacity
- All the math is pretty simple