Understanding MVS Busy % versus LPAR Busy % versus Physical Busy %

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

- **Understanding MVS Busy % versus LPAR Busy % versus Physical Busy %**

- During this webinar, Peter Enrico will explore the SMF 70 processor dispatch measurements and the formulas that these measurements are used with to calculate physical and logical and MVS processor utilizations. Also discussed will be the conceptual difference and usage between physical and logical utilizations.

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
  - February 19-23, 2024
- Parallel Sysplex and z/OS Performance Tuning
  - August 20-21, 2024
- Essential z/OS Performance Tuning
  - October 7-11, 2024
- Also… please make sure you are signed up for our free monthly z/OS educational webinars! (email contact@epstrategies.com)

Like what you see?

- Free z/OS Performance Educational webinars!
  - The titles for our 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*

- If you want a free cursory review of your environment, let us know!
  - We're always happy to process a day's worth of data and show you the results
  - See also: http://pivotor.com/cursoryReview.html

Like what you see?

- The z/OS Performance Graphs you see here come from Pivotor
- If you don't see them in your performance reporting tool, or you just want a free cursory performance 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!

Physical Busy % vs LPAR Busy % vs MVS Busy %

- PR/SM Physical Busy Utilization
  - Helps us gain insights into the physical constraints of the machine / CEC
  - Based on the number of physical processors and dispatch times

- PR/SM LPAR (Logical) Busy Utilization
  - Helps us gain insights into the logical constraints of the LPAR / z/OS system
  - Based on number of logical processors and dispatch times

- MVS Busy Utilization
  - Helps us gain insights into the demand for CPU by the LPAR / z/OS system
  - Based on the number of logical processors and wait times

Breakdown of General-Purpose Processor

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

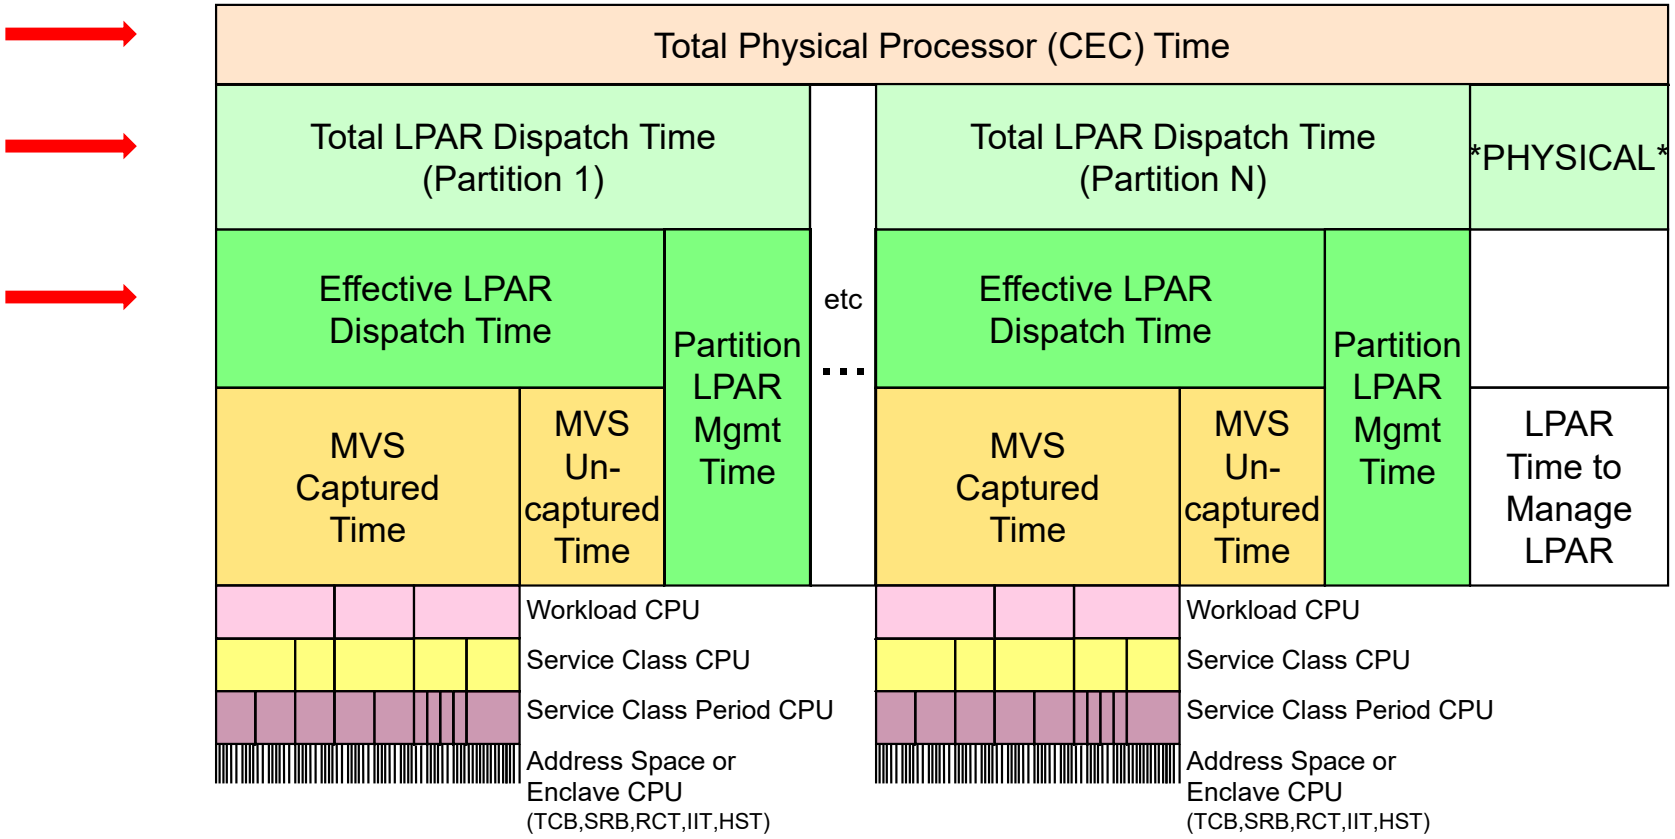

Breakdown of zIIP Engine Time

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

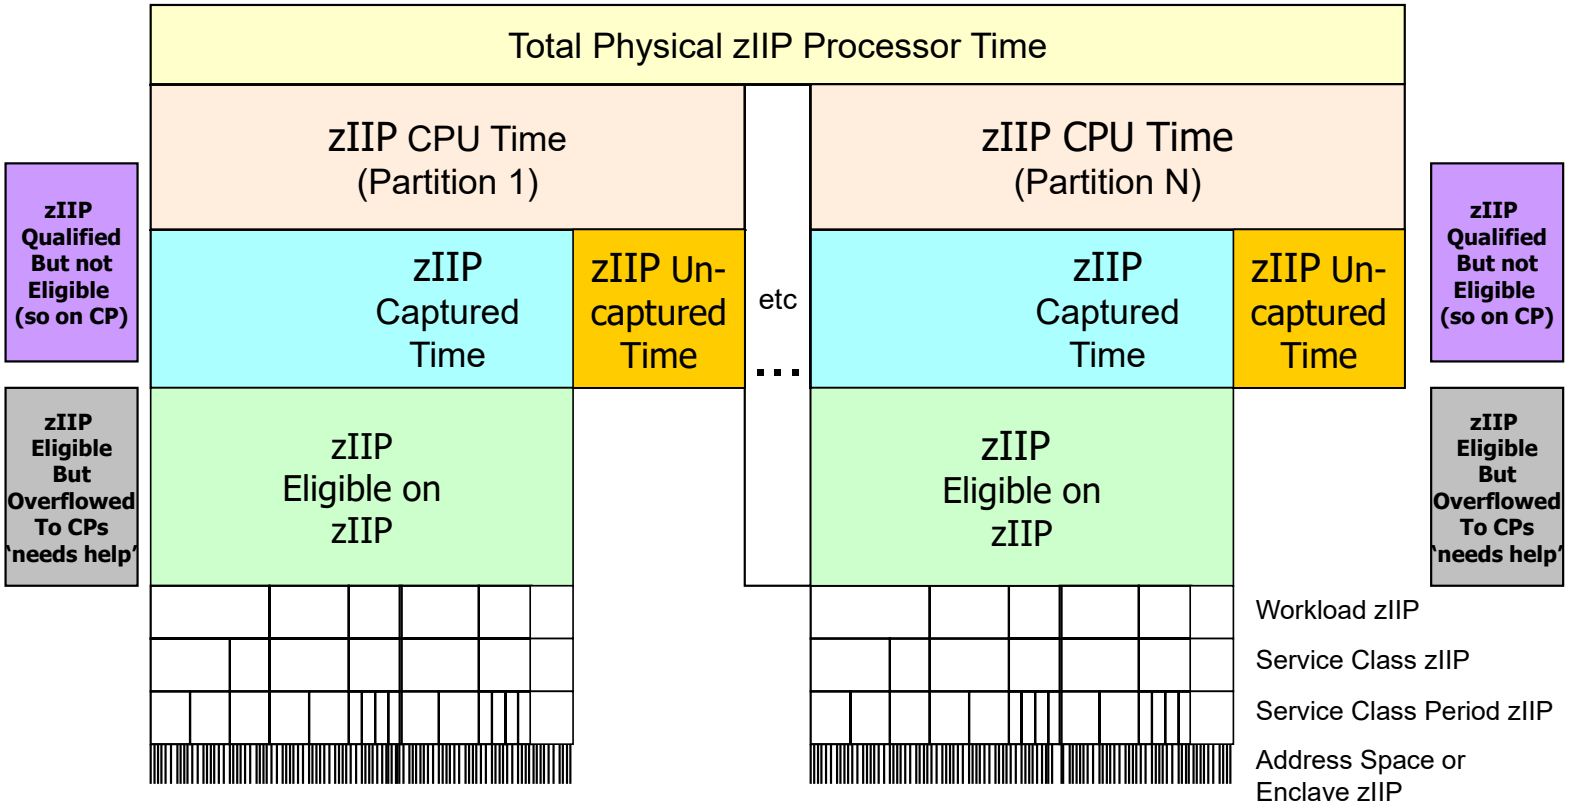

LPAR Terminology Review

- Physical Processors
  - Physical CEC processors that are used by the different partitions
  - Processors can be
    - GCP – General CPU Processor
    - ICF – Integrated Coupling Facilities
    - IFL – Integrated Facilities for Linux
    - zIIP – zArchitecture Integrated Information Processor

Physical Utilization helps us understand how busy this physical processor pool is

Same exercise for each pool of processors

- GCPs
- ICFs
- IFLs
- zIIPs

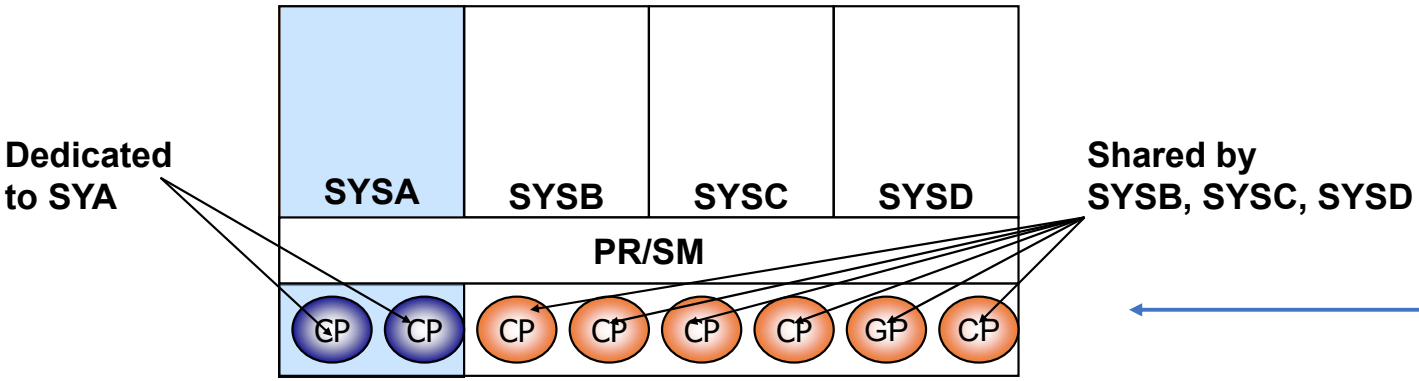

LPAR Terminology Review

- Logical Processor
  - Each system image as some number of logical processors assigned
    - System image thinks it has 100% of its number of processors
  - Dedicated processors
    - Physical processor dedicated to a partition 100% of the time
    - Accumulates both CPU using and wait/idle time
  - Shared logical processors
    - Physical processor that can be share among one or more partitions
    - Physical processors not dedicated to a particular partition

Logical Utilization helps us understand how busy each LPARs logical pool of processors is.

- How busy is SYSA's 2 logical CPs
- How busy is SYSB's 6 logical CPs
- How busy is SYSC's 2 logical CPs
- How busy is SYSD's 1 logical CP

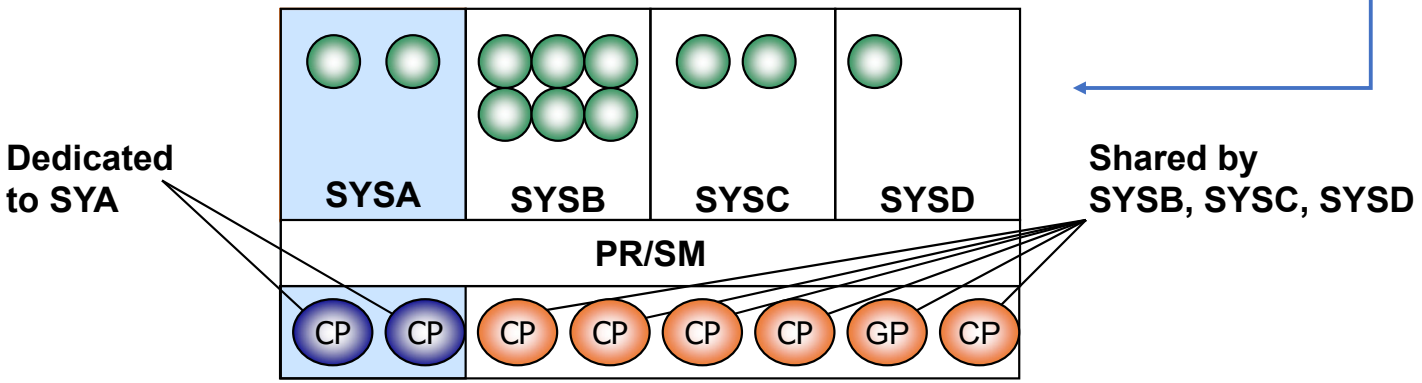

CPU Time in PR/SM

- There are a variety of time values to understand
  - From a single CP point of view, some of these times are as follows:

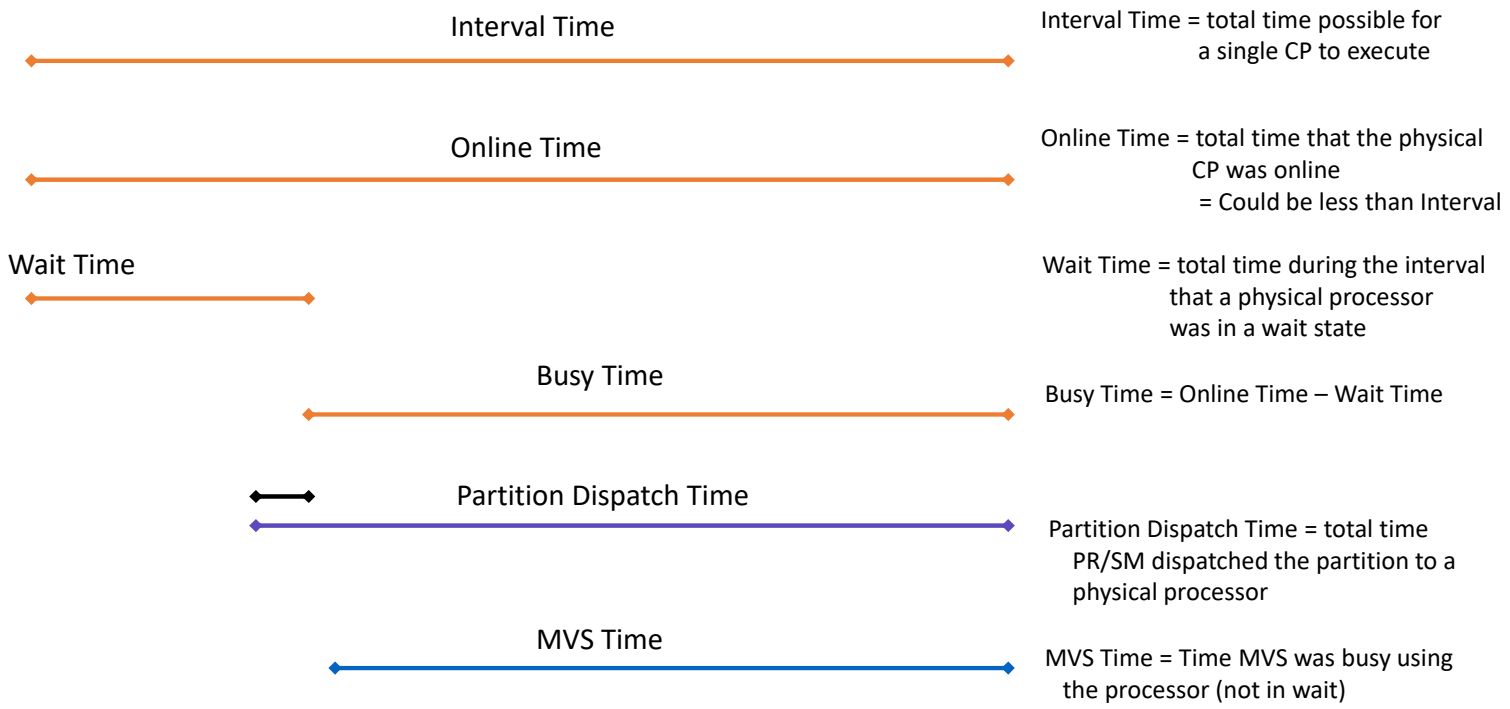

LPAR Terminology Review

- Say SYSC is using 100% of its logical capacity
- Would adding another physical CPU to the CEC help?

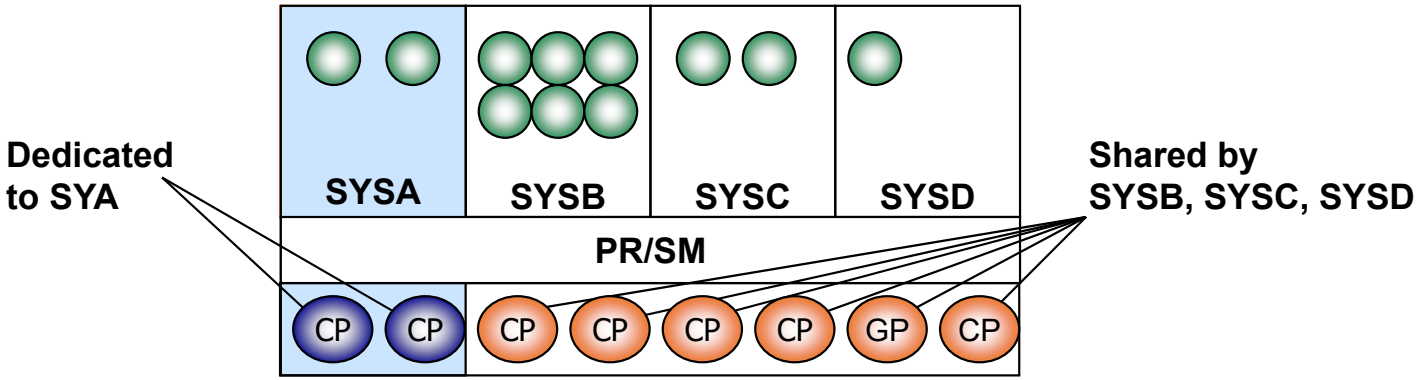

Dispatch Time

- Partition Effective Dispatch Time
  - Time a physical CPU was dispatched to a logical CPU during measurement interval
    - Think of this as time that z/OS system and the workloads got to use the physical CPU
  - For *PHYSICAL* this value is blank

- Partition Total Dispatch Time
  - Includes Effective Dispatch Time plus LPAR Management time
  - For *PHYSICAL* this value is includes the processor time that cannot be attributed to any one partition
  - Time that LPAR spent managing itself

- LPAR Management Time = Delta between Total Time and Effective Time
  - Time PR/SM spent managing a particular partition

Example of 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

MVS PARTITION NAME	PRD1	PHYS PROC NUM	7
IMAGE CAPACITY	292	CP	4
NUMBER OF CONFIGURED PARTITIONS	5	ICF	1
WAIT COMPLETION	NO	IIP	2
DISPATCH INTERVAL	DYNAMIC		

GROUP NAME	PLEX01	INITIAL CAP	NO
LIMIT	292	LPAR HW CAP	NO
AVAILABLE	159	HW GROUP CAP	NO
		ABS MSU CAP	NO

---------- PARTITION DATA ---------- -- LOGICAL PARTITION PROCESSOR DATA -- -- AVERAGE PROCESSOR UTILIZATION PERCENTAGES -

				----MSU----		--CAPPING---		--PROCESSOR		----DISPATCH TIME DATA----		LOGICAL PROCESSORS		--- PHYSICAL PROCESSORS ---		
NAME	S	BT	WGT	DEF	ACT	DEF	WLM%	NUM	TYPE	EFFECTIVE	TOTAL	EFFECTIVE	TOTAL	LPAR MGMT	EFFECTIVE	TOTAL
PRD1	A	N	869	0	92	N N N	0.0	4.0	CP	00.15.41.497	00.15.44.687	26.15	26.24	0.09	26.15	26.24
DEV1	A	N	105	0	33	N N N	0.0	2.0	CP	00.05.38.443	00.05.40.398	18.80	18.91	0.05	9.40	9.46
TEST	A	N	26	0	2	N N N	0.0	2.0	CP	00.00.21.356	00.00.21.830	1.19	1.21	0.01	0.59	0.61
PHYSICAL											00.00.04.494			0.12		0.12
TOTAL			1000							00.21.41.296	00.21.51.411			0.28	36.15	36.43
CF01	A		DED					1	ICF	00.14.59.993	00.14.59.993	100.0	100.0	0.00	100.0	100.0
PHYSICAL											00.00.00.005			0.00		0.00
TOTAL			0							00.14.59.993	00.14.59.998			0.00	100.0	100.0
PRD1	A	N	145			N N N		2	IIP	00.00.49.347	00.00.50.651	2.74	2.81	0.07	2.74	2.81
DEV1	A	N	10			N N N		1	IIP	00.00.29.036	00.00.30.031	3.23	3.34	0.06	1.61	1.67
TEST	A	N	6			N N N		1	IIP	00.00.02.197	00.00.02.282	0.24	0.25	0.00	0.12	0.13
PHYSICAL											00.00.01.954			0.11		0.11
TOTAL			161							00.01.20.580	00.01.24.920			0.24	4.48	4.72

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

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

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

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

LPAR Management Times and Utilizations

- LPAR Management Time
  - Time PR/SM spent managing a partition
  - For the partition '*PHYSICAL*", this is the amount of time PR/SM spent managing itself. It is time that could not be attributed to any single partition.

- Physical Processor Utilization LPAR Management
  - Percentage of the measurement interval that PR/SM spent managing the partition. It is reported as a percentage of total physical time possible

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

- *PHYSICAL* Partition

$$\frac{\sum \text{Partition Total Dispatch Time for partition PHYSICAL}}{\text{No of Physical Processors} * \text{Interval Time}} * 100$$

RMF Partition Data Report – Machine Utilization Values

- Total lines
  - Reports the total utilization of the physical processors on the machine
  - Provides a view of total utilization by all LPARs of physical processor resource

```
-- PARTITION DATA -- -- LOGICAL PARTITION PROCESSOR DATA -- -- AVERAGE PROCESSOR UTILIZATION PERCENTAGES -

                    --PROCESSOR ----DISPATCH TIME DATA---- LOGICAL PROCESSORS --- PHYSICAL PROCESSORS ---
NAME       S BT WGT   NUM  TYPE  EFFECTIVE       TOTAL     EFFECTIVE    TOTAL LPAR MGMT  EFFECTIVE  TOTAL

PRD1       A  N 869   4.0  CP   00.15.41.497  00.15.44.687     26.15    26.24      0.09      26.15  26.24
DEV1       A  N 105   2.0  CP   00.05.38.443  00.05.40.398     18.80    18.91      0.05       9.40   9.46
TEST       A  N  26   2.0  CP   00.00.21.356  00.00.21.830      1.19     1.21      0.01       0.59   0.61
*PHYSICAL*                                   00.00.04.494                          0.12                0.12
           ------              ------------  ------------                       -----      -----  -----
  TOTAL      1000              00.21.41.296  00.21.51.411                       0.28      36.15  36.43


CF01       A    DED     1  ICF  00.14.59.993  00.14.59.993     100.0    100.0      0.00      100.0  100.0
*PHYSICAL*                                   00.00.00.005                          0.00                0.00
           ------              ------------  ------------                       -----      -----  -----
  TOTAL         0              00.14.59.993  00.14.59.998                       0.00      100.0  100.0


PRD1       A  N 145     2  IIP  00.00.49.347  00.00.50.651      2.74     2.81      0.07       2.74   2.81
DEV1       A  N  10     1  IIP  00.00.29.036  00.00.30.031      3.23     3.34      0.06       1.61   1.67
TEST       A  N   6     1  IIP  00.00.02.197  00.00.02.282      0.24     0.25      0.00       0.12   0.13
*PHYSICAL*                                   00.00.01.954                          0.11                0.11
           ------              ------------  ------------                       -----      -----  -----
  TOTAL       161              00.01.20.580  00.01.24.920                       0.24       4.48   4.72
```

Example of Physical Processor Busy

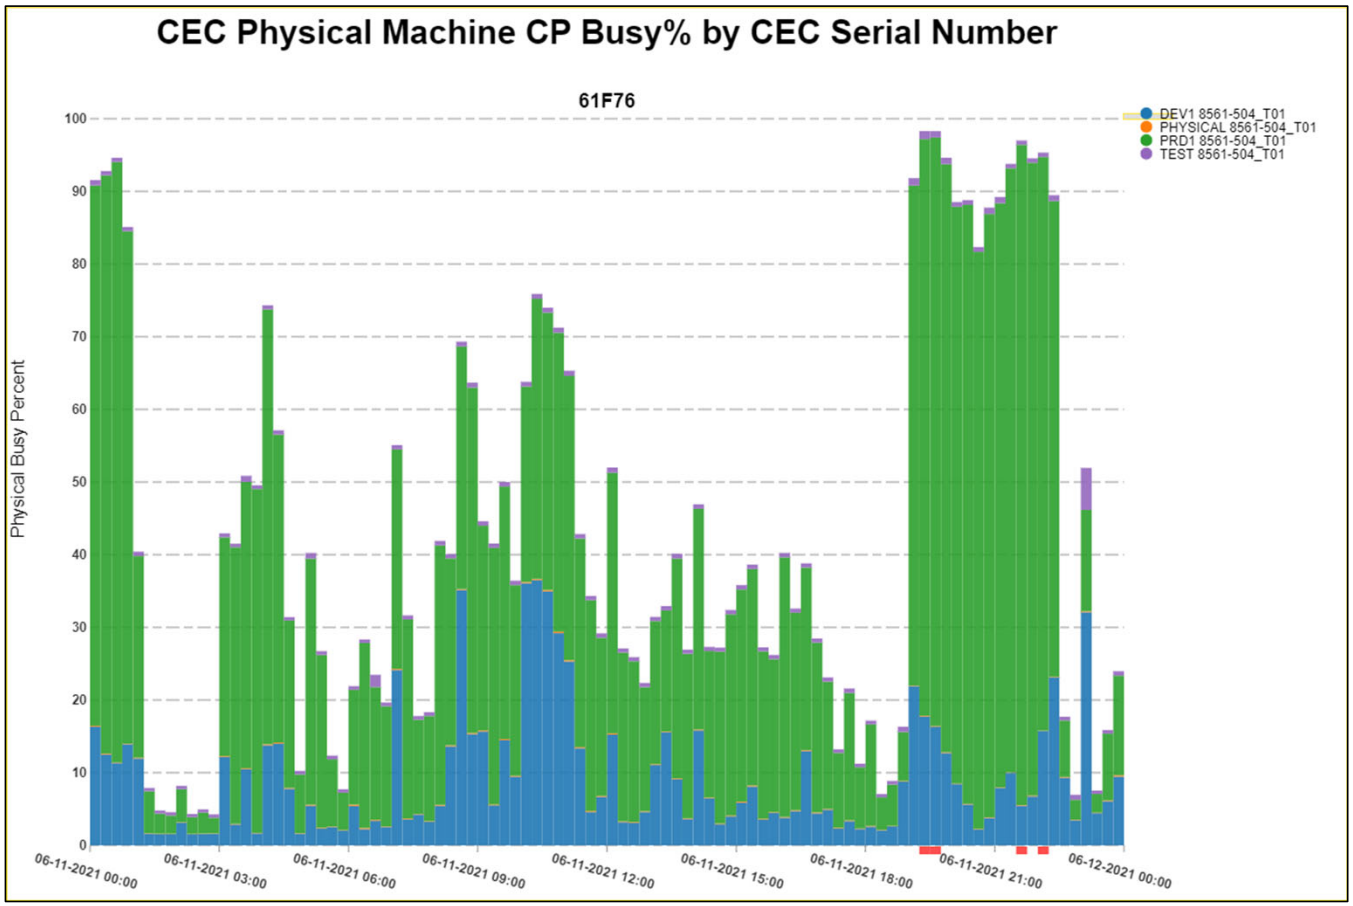

For each logical partition, the physical utilization of the machine is a function of the number of physical processors.

The partitions of all partitions is based on the number of physical processors of the machine, and not the number of logical processors assigned to the partition.

Physical Machine Utilization as Line Chart

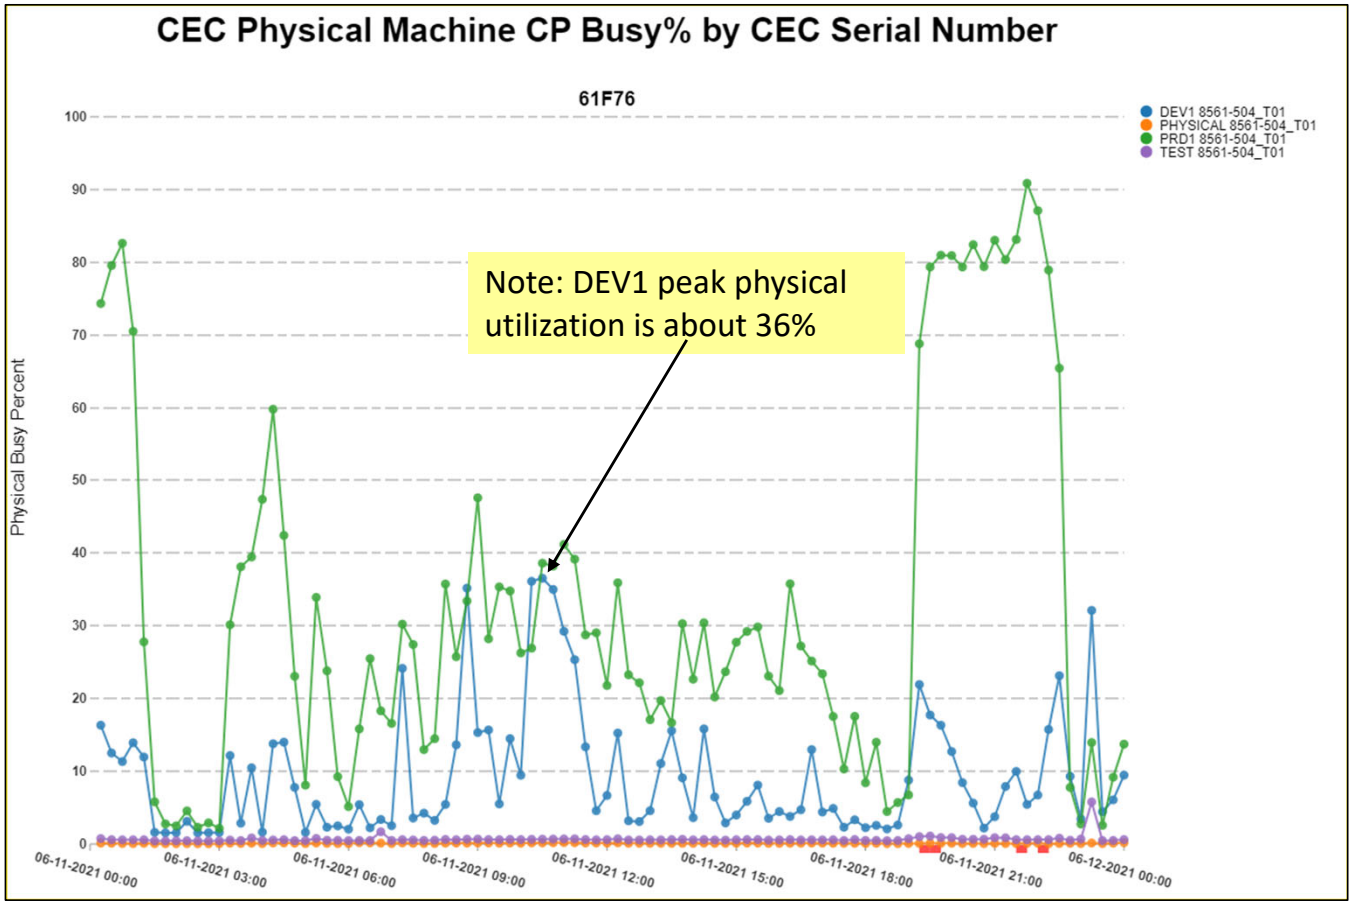

For each logical partition, the physical utilization of the machine is a function of the number of physical processors.

The utilization of each partition is based on the number of physical processors of the machine, and not the number of logical processors assigned to the partition.

Example of Logical Processor Busy

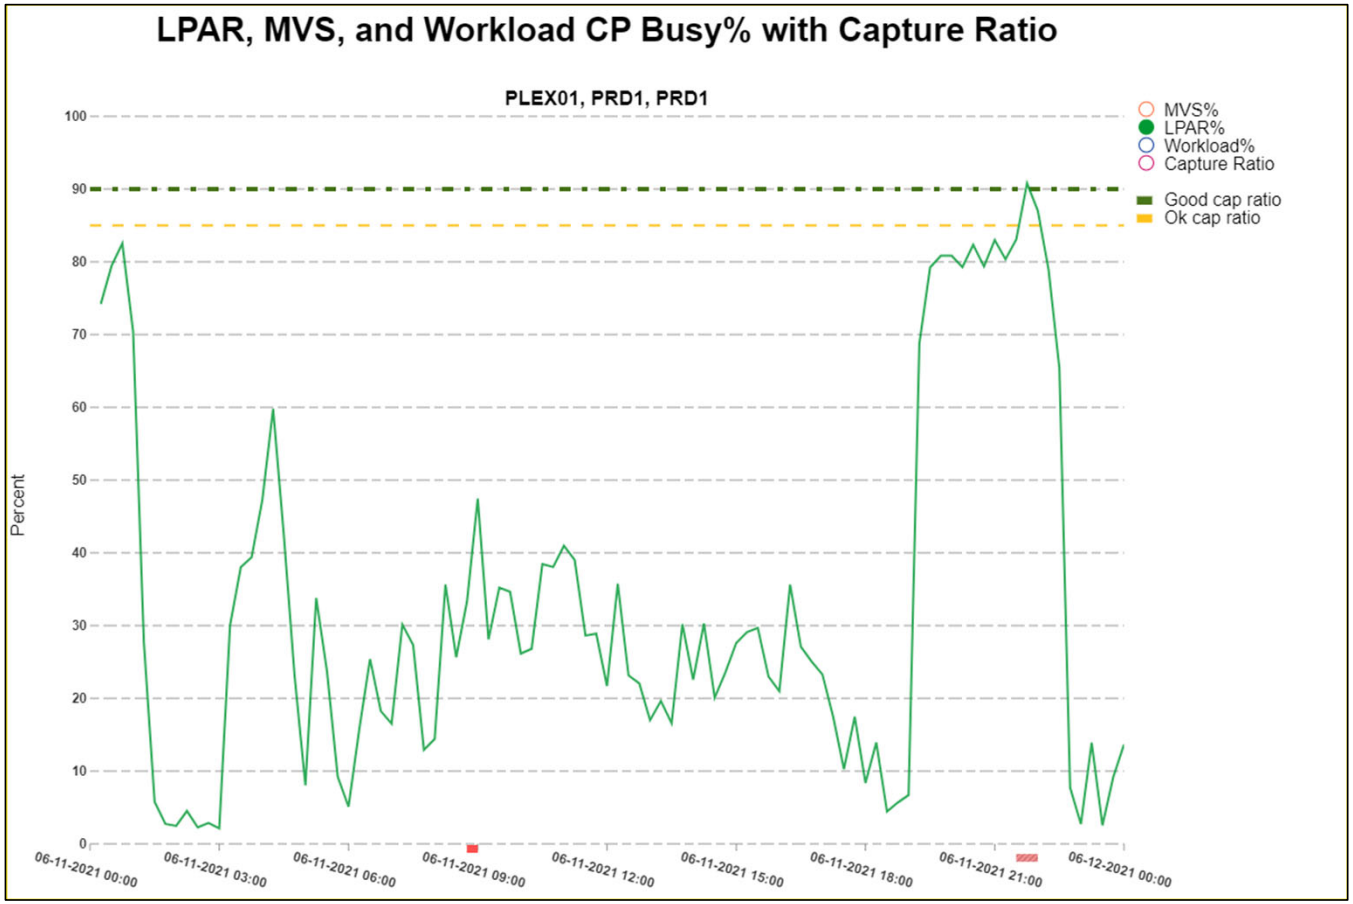

When the number of logical processors assigned to a partition is equal to the number of physical processors on the CEC, the LPAR utilization for that LPAR will equal the physical utilization for that LPAR.

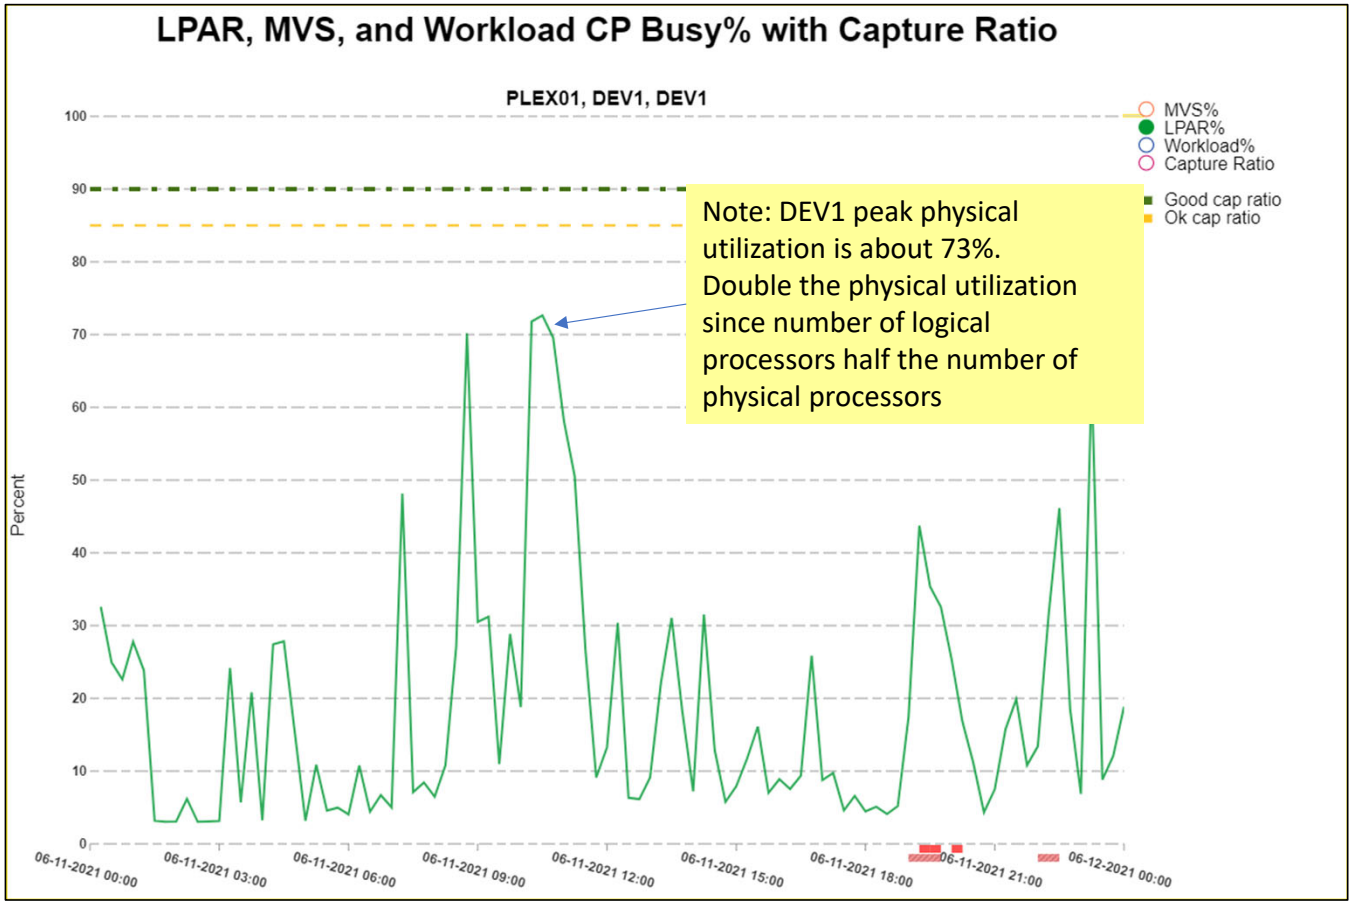

When the number of logical processors assigned to a partition is NOT equal to the number of physical processors on the CEC, the LPAR utilization for that LPAR will be higher since we are looking at the constraint of the number of logical processors.

So, what is MVS Busy %?

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

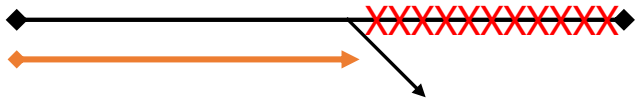

Voluntary Wait
- z/OS voluntarily gives up the processor
- MVS time equals dispatch time

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

LPAR Busy % with Config CPs and only Unparked CPs

- LPAR Busy % based on configured number of logical processors
  - Reports logical constraint of the LPAR
- LPAR Busy % based on unparked number of logical processors
  - Reports the HiperDispatch constraint

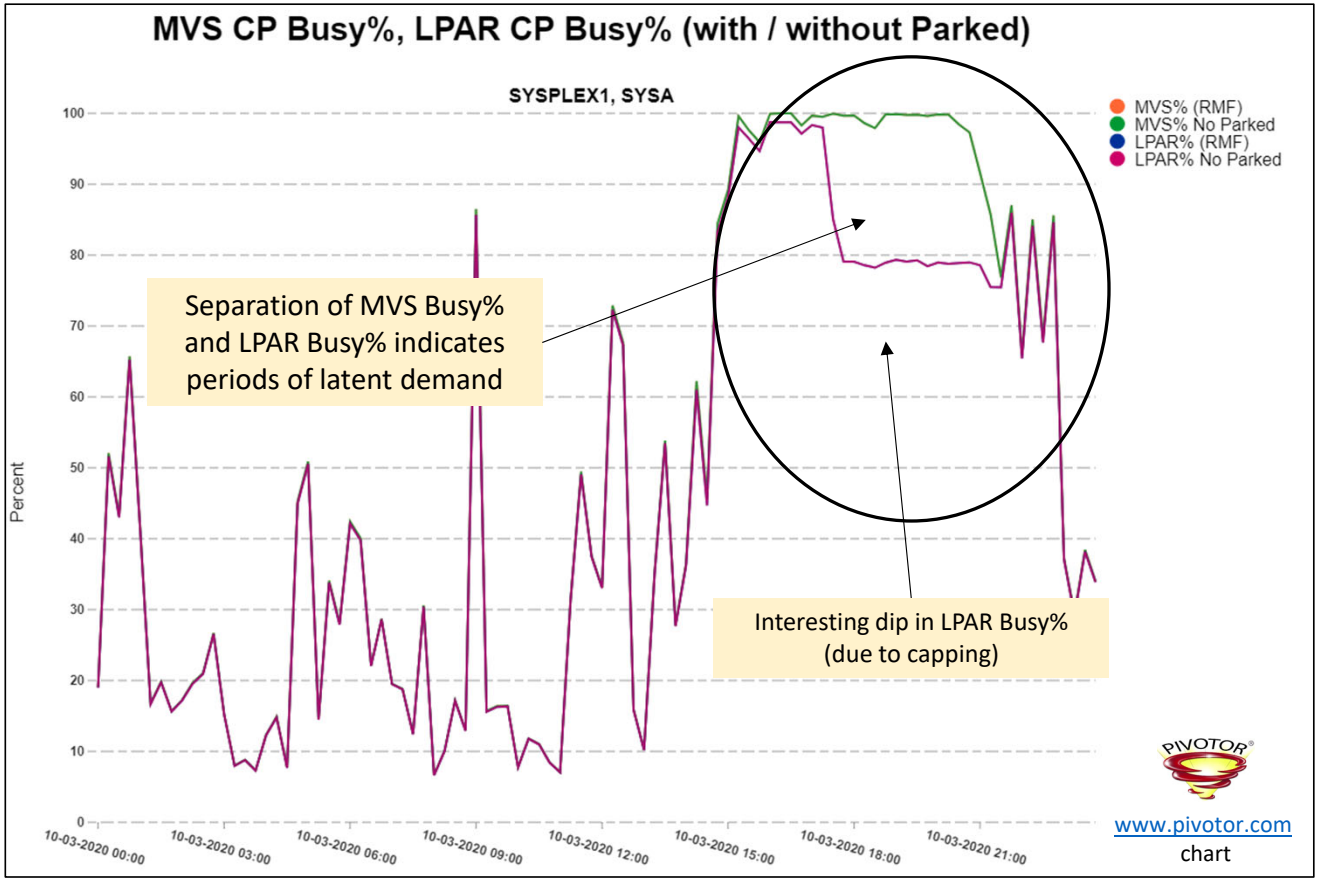

LPAR Busy % with Config CPs and only Unparked CPs

- LPAR Busy % based on configured number of logical processors
  - Reports logical constraint of the LPAR
- LPAR Busy % based on unparked number of logical processors
  - Reports the HiperDispatch constraint

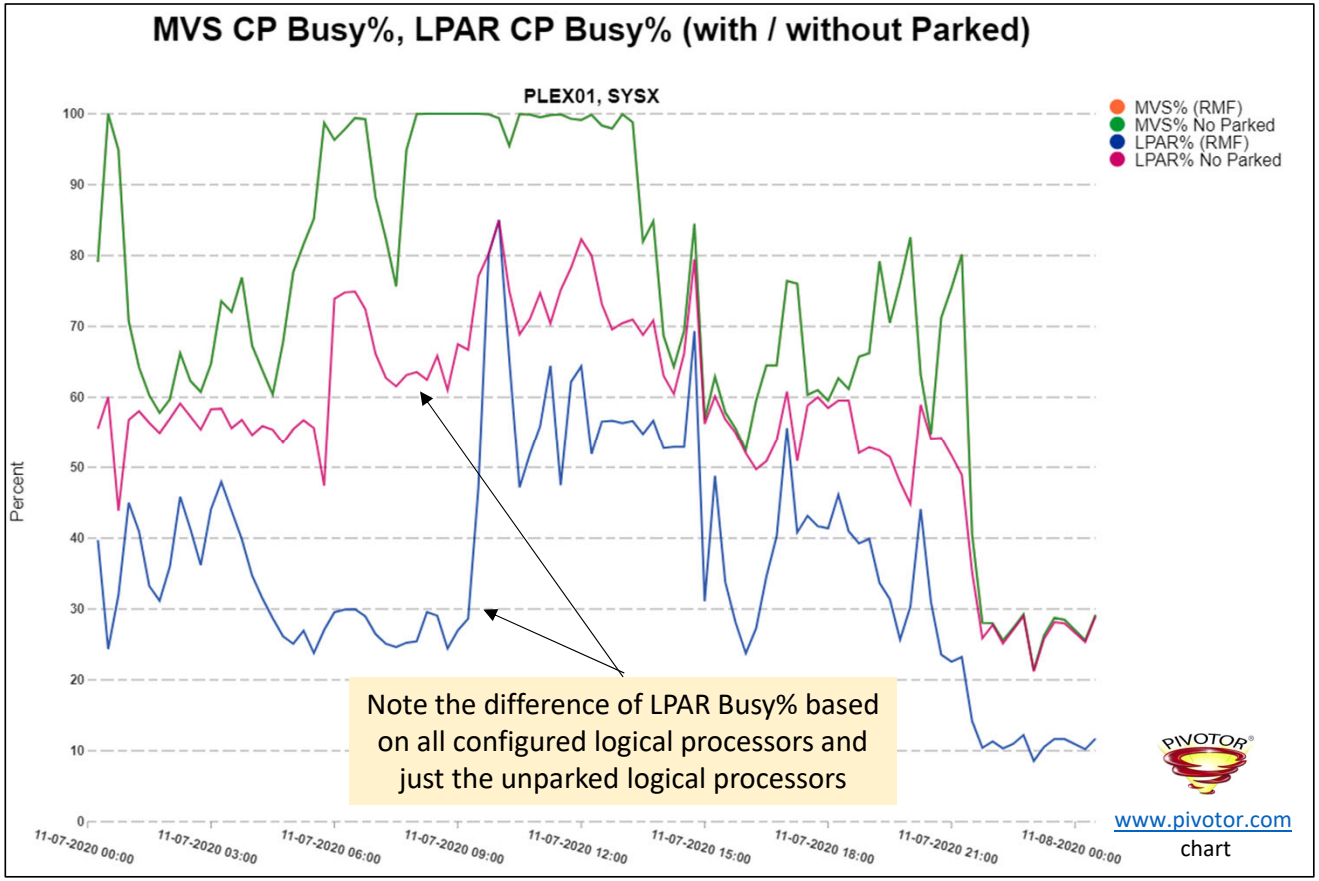

Distribution of work unit queue lengths

- 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:
  - N = number of unparked CP + zIIP engines

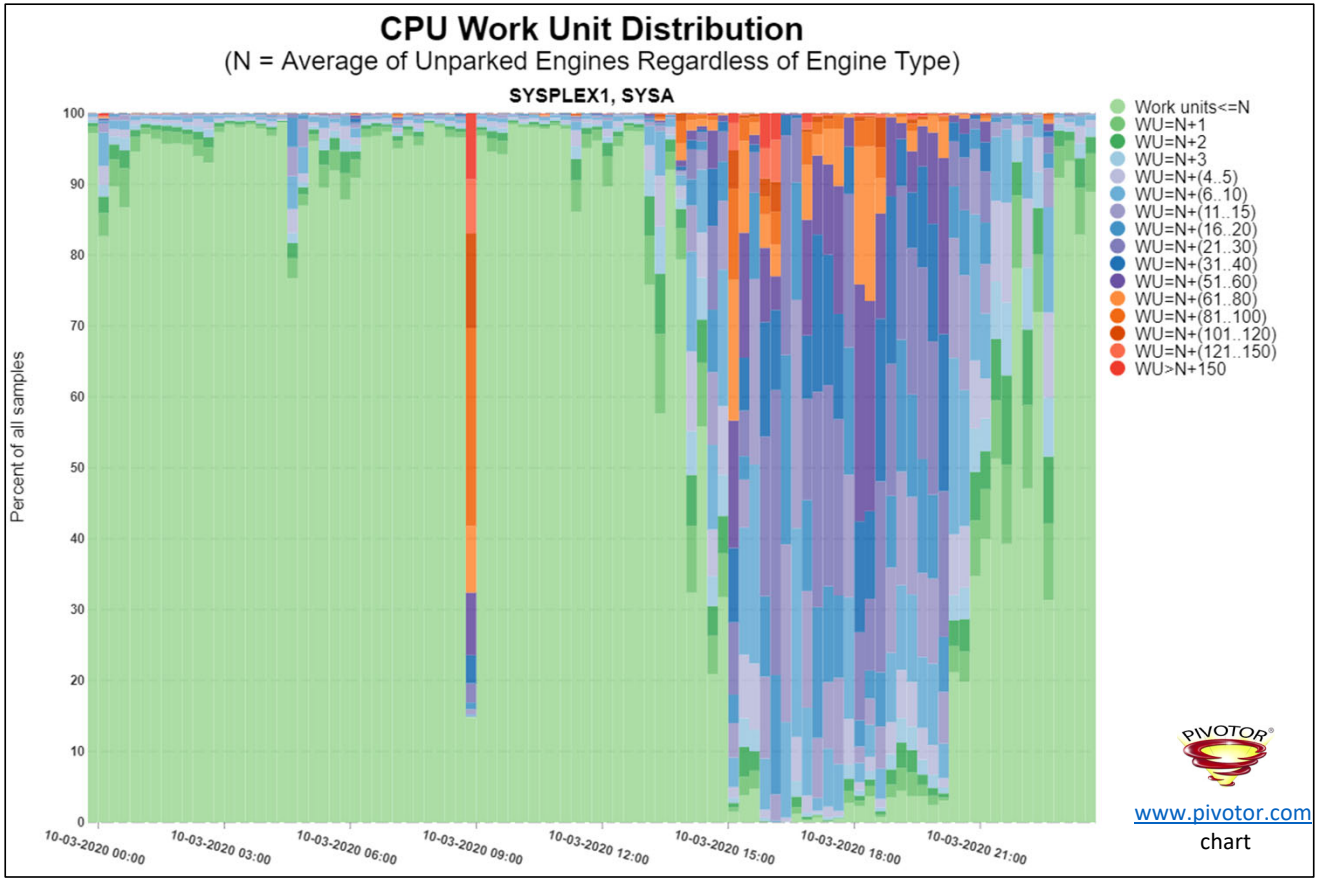

Distribution of work unit queue lengths

- How much latent demand is too much, too unhealthy?
- Assuming a rule of thumb that CP queues lengths of > 3 times the number of CP CPUs is unhealthy latent demand
  - We see here that during the evening hours we have continuous unhealthy latent demand
  - With large percentages of the measurement intervals of more than 100 Work Units queued up

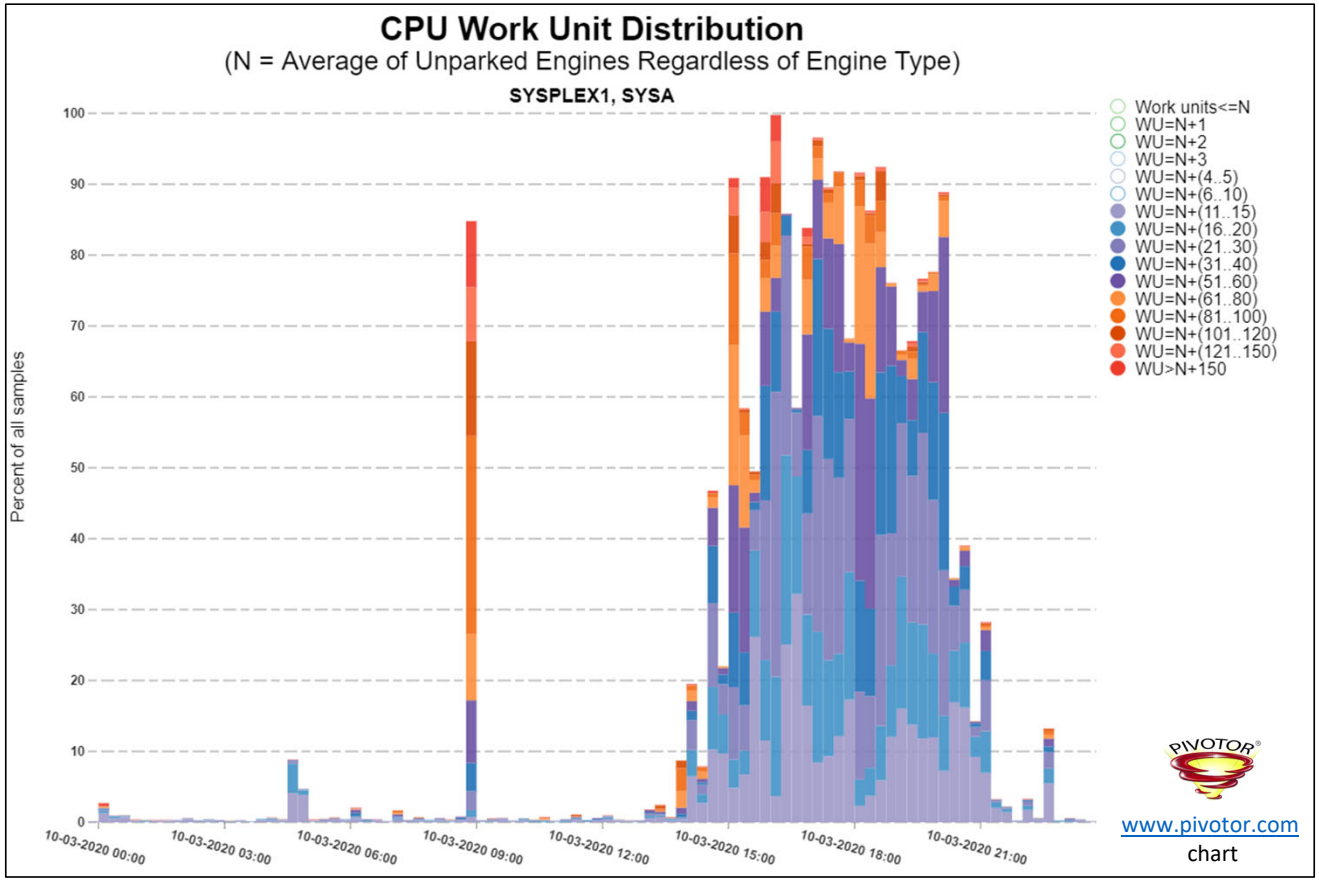

Relationship of LPAR% delta to MVS%, and Work Unit Queuing

- When we overlay the two charts, we see a correlation

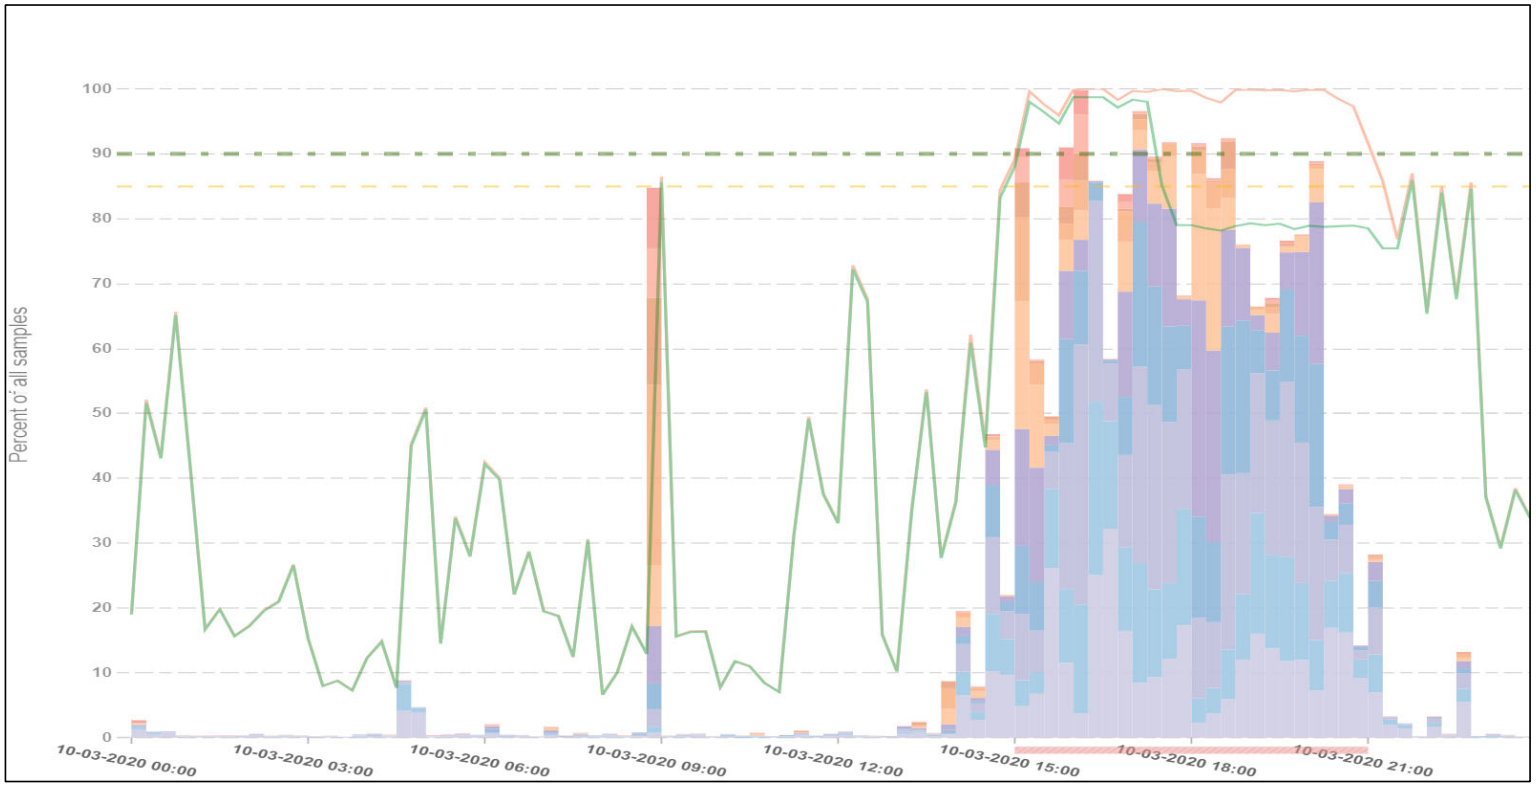

www.pivotor.com
chart

Physical Busy % vs LPAR Busy % vs MVS Busy %

- PR/SM Physical Busy Utilization
  - Helps us gain insights into the physical constraints of the machine / CEC
  - Based on the number of physical processors and dispatch times

- PR/SM LPAR (Logical) Busy Utilization
  - Helps us gain insights into the logical constraints of the LPAR / z/OS system
  - Based on number of logical processors and dispatch times

- MVS Busy Utilization
  - Helps us gain insights into the demand for CPU by the LPAR / z/OS system
  - Based on the number of logical processors and wait times

Questions?