Performance Analysis
and Capacity Planning Report

for

WATECH

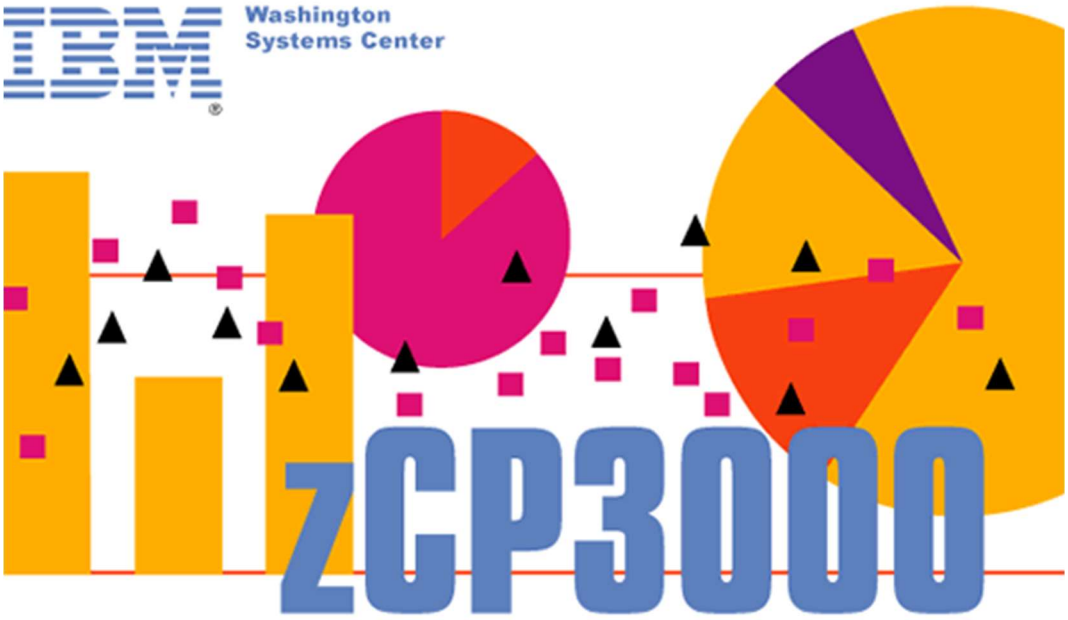

Date Prepared: Aug 20, 2019

About This Report

Please Be Aware

The performance estimates presented are approximations which are believed to be sound. The results are dependent upon the data provided by the user and based on IBM's experience with similar installations. The degree of success which you may achieve in the use of your installed equipment and programs is dependent on a number of factors, many of which are not under IBM's control. Thus, IBM does not warrant or guarantee that you can or will achieve similar results, and it is your responsibility to validate the estimates furnished and determine their relevance to your operation. You should verify all results provided by IBM prior to making a data processing decision.

Processor Power Numbers

Capacity projections are based on LSPR algorithms that project capacity for any LPAR Host processor model with its specific LPAR configuration. For z/OS and z/VM partitions, the workload assignment is based on CPU MF data if available, otherwise "Average" will be used. All capacity results are scaled based on the Reference-CPU setting. For this study, the Reference-CPU is set to a 2094-701 with an assumed capacity of 593.0, representing the productive capacity of a shared single-partition configuration.

Specialty Engines

The information contained in this document provides only general descriptions of the types and portions of workloads that are eligible for execution on Specialty Engines (e.g, zIIPs, zAAPs, and IFLs) ("SEs"). IBM authorizes customers to use IBM SEs only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the "Authorized Use Table for IBM Machines" provided at www.ibm.com/systems/support/machine_warranties/machine_code/aut.html. No other workload processing is authorized for execution on an SE. IBM offers SEs at a lower price than General Processors/Central Processors because customers are authorized to use SEs only to process certain types and/or amounts of workloads as specified by IBM in the AUT.

Trademarks

IBM is a registered trademark of the International Business Machines Corporation in the United States or other countries and may be used throughout this publication. Complete trademark information can be found at http://www.ibm.com/legal/copytrade.shtml . The following terms are trademarks of the International Business Machines Corporation in the United States or other countries, and may be used throughout this publication:

z/OS, z/VM, z/VSE, PR/SM, FICON, ESCON

Document Details

Created by	zCP3000 - ebf4333
Data created by	CP3KEXTR 06/28/19 V3.90
zPCR 9.2f	Based on z/OS-2.2 LSPR Data (04/10/2018) including three different workloads (Low, Average, High) and assuming HiperDispatch is active for z10 and later processors.
GMT offset	-7:00

Table of Contents

Executive Summary for WATECH

- This analysis is based on WATECH's 2828-R05 system referenced in this report as CPC6E5A7
- We were asked to establish a baseline capacity representation of current processing activities
- WATECH provide 1-week of SMF data from 7/21/19 – 7/27/19. This data is assumed reflective of normal processing activities
- SMF data were in 30-min RMF intervals from PROD/TEST partitions V1RA, V4YM, VLLL, VV51, V2DL, NET1, V3HC & C2JR. We selected peak interval from prime shift on 7/25/19 at 0900hr as the study interval for capacity comparison purposes
- Detected mostly Average (V1RA, V4YM) to High (remaining LPARs) Relative Nest Intensity (RNI) workload categories from the provided SMF113 records
- Observed acceptable usage of the **1x configured zIIP engine** with additional room for growth
- **Unable to detect actual utilization of configured 2x IFL engines** due to IFL resources defined as DEDICATED instead of SHARED to respective partitions
- Proposed z14 ZR1 models are potential workable processor options with varying capacity levels to replace existing 2828-R05
- Lastly, please **be aware** that it is reasonable to expect a **±5% margin of error in our estimates**

2828-R05's GCP Utilization by Partition

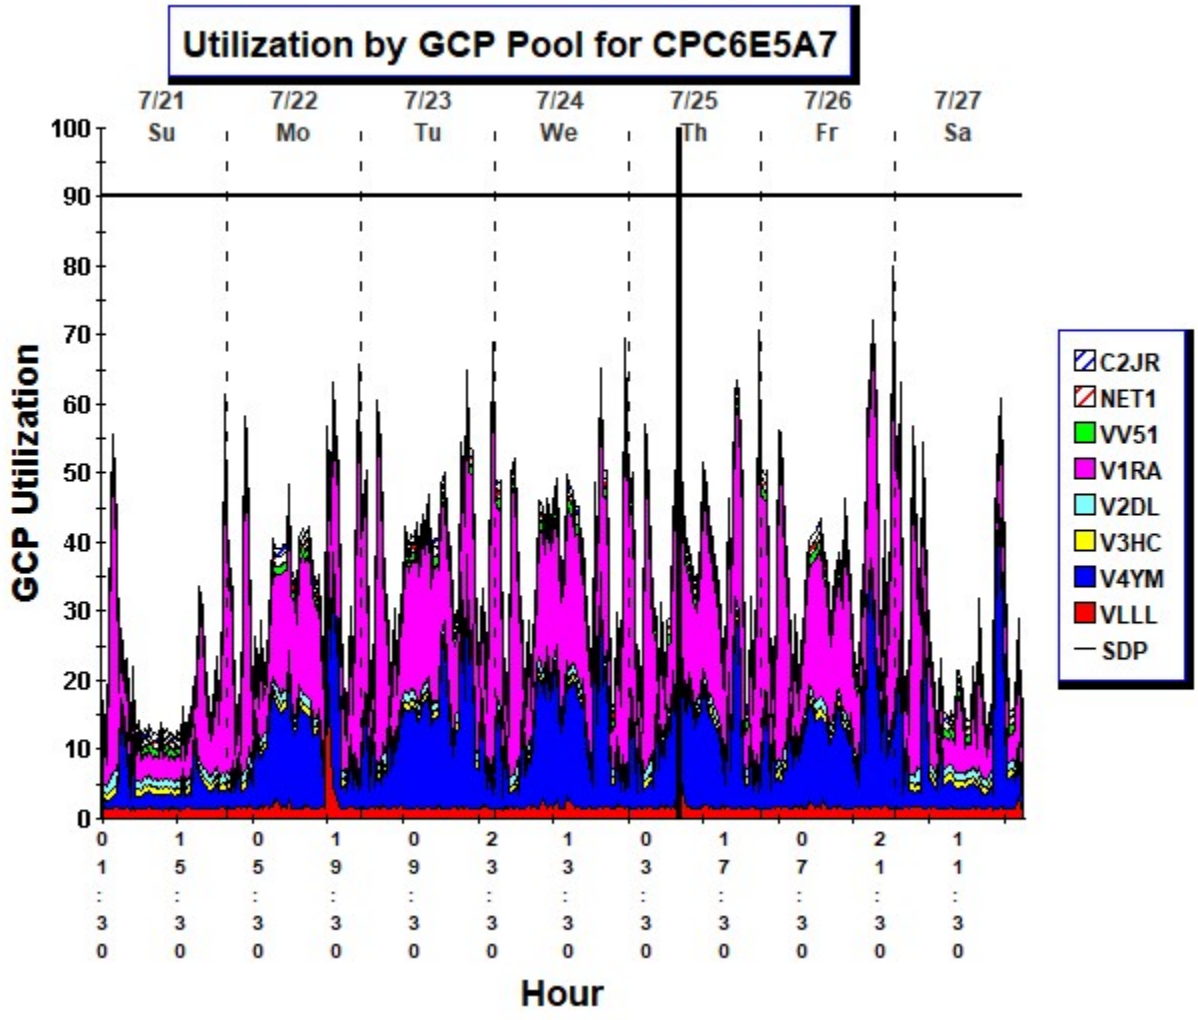

This graph represents the accumulated CPU utilizations for all the partitions running with GCPs in this processor. The utilization is the percent of the 5 GCPs configured on the 2828-R05.

The GCP utilization in the study interval is 62.1%. The largest component of the utilization in the study interval is contributed by V1RA. The individual utilizations by SYSID are listed below.

The saturation design point (SDP) is set at 90.0% for this study. The table below contains the study interval and maximum physical CPU% in the samples. There is also the partitions Fair Share which is the percentage of the shared engine pool to which this partition is entitled. The GCPs by Weight represents the number of GCPs required to satisfy the given weight of this partition. The field Image Capacity refers to the capacity of this partition in MSUs. If a Defined Capacity has been set, then this will be the value. Otherwise, it will be the share of the entire CPCs capacity that this partitions logical engines can possibly use.

	Study	Max.	Avg.	LPAR	Fair	Logical	GCPs By	Hiper	Image		Workload
Partition	CPU%	CPU%	CPU%	Weight	Share	GCPs	Weight	Dispatch	Capacity	Cap	Category
C2JR	1.3%	2.8%	1.2%	40	2%	2.0	0.1	Y	7	S	High
NET1	1.4%	1.6%	1.2%	100	6%	3.0	0.3	Y	7	S	High
VV51	1.6%	2.5%	1.4%	100	6%	2.0	0.3	Y	7	S	High
V1RA	**32.1%**	**61.0%**	**16.2%**	**706**	**44%**	**4.0**	**2.2**	**Y**	**53**	**S**	**Average**
V2DL	1.5%	2.6%	1.4%	160	10%	3.0	0.5	Y	7	S	High
V3HC	1.3%	1.5%	1.1%	50	3%	2.0	0.2	Y	7	S	Average
V4YM	**16.1%**	**41.5%**	**9.6%**	**350**	**22%**	**4.0**	**1.1**	**Y**	**30**	**S**	**Average**
VLLL	6.7%	22.7%	1.7%	100	6%	2.0	0.3	Y	7	S	Average
Total	62.1%	80.0%	33.8%	1,606		22.0			213		

In the table above, the codes for the "Cap" column mean: A is absolute capping, H is hard capping, S is soft capping and G is group capping.

The Study Interval (the peak interval from the Prime shift) is shown on the graph as a solid line drawn on 7/25/19 at 09:00.

The following partitions were not active for all the samples.

- C3PO
- RAD1

zIIP Utilization by Partition (1x zIIP CP)

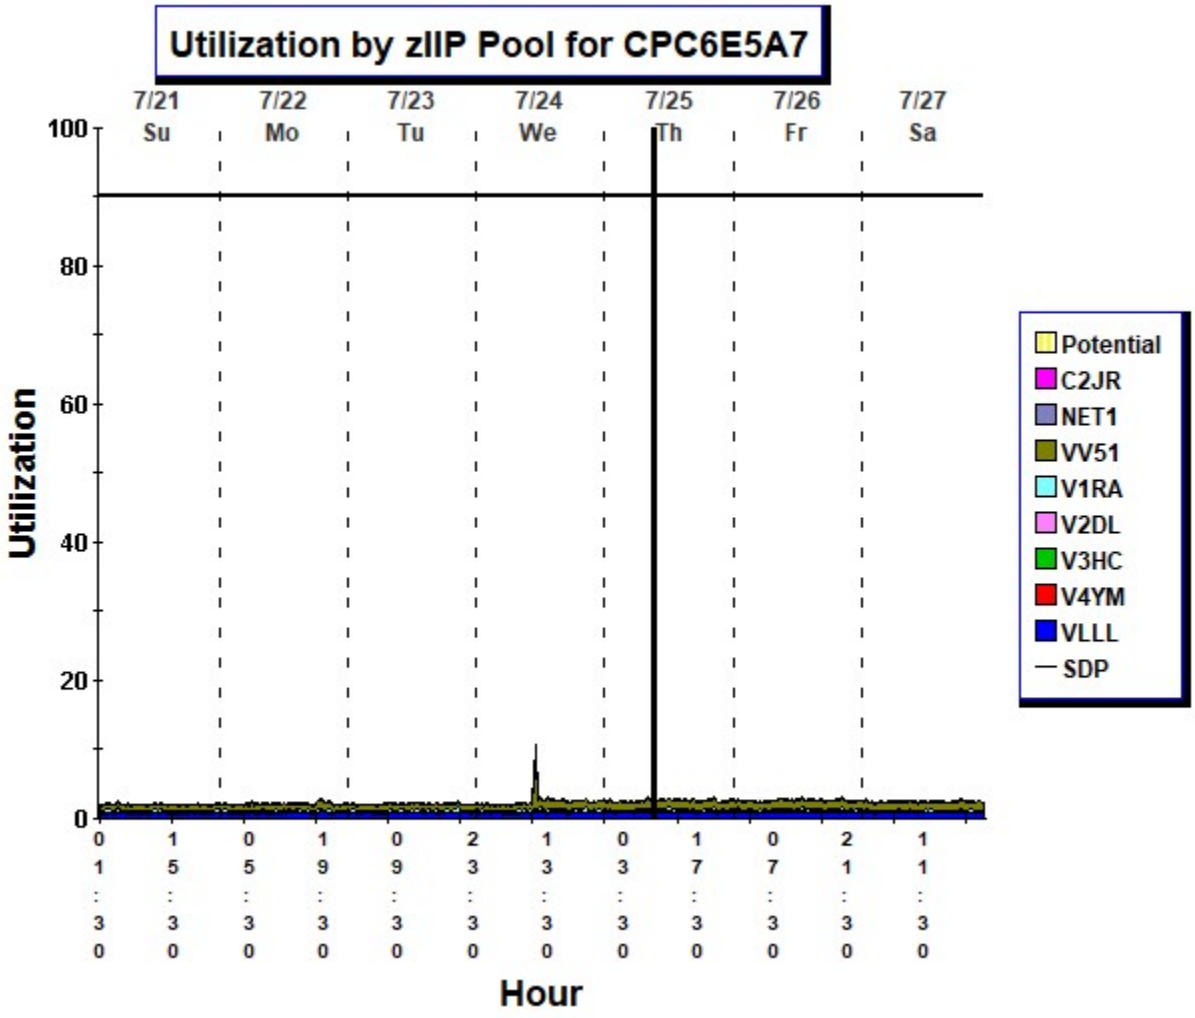

This graph represents the accumulated physical zIIP utilizations for all the partitions running with zIIPs in this processor. The utilization is the percent of 1.0 engines configured on the 2828-R05.

The zIIP utilization in the study interval is 3.7%. The largest component of the utilization in the study interval is contributed by VV51. The individual utilizations by SYSID are listed below.

The saturation design point (SDP) for zIIP is set at 90.0% for this study.

The study interval is shown on the graph as a line drawn on 7/25/19 at 09:00.

The 2828-R05 processor runs zIIP engines about 2.5 times faster than GCP engines.

	Study	Max.	Avg.	LPAR				Honor
Partition	**zIIP%**	**zIIP%**	**zIIP%**	**Weight**	**#zIIPs**	**SMT**	**Capacity**	**Priority**
C2JR	0.1%	0.1%	0.1%	3%	1.0	N	S	Y
NET1	0.1%	0.1%	0.1%	7%	1.0	N	S	Y
VV51	1.6%	7.1%	1.0%	3%	1.0	N	S	Y
V1RA	1.0%	1.0%	0.3%	47%	1.0	N	S	Y
V2DL	0.1%	0.1%	0.1%	11%	1.0	N	S	Y
V3HC	0.1%	0.1%	0.1%	3%	1.0	N	S	Y
V4YM	0.1%	0.1%	0.1%	23%	1.0	N	S	Y
VLLL	0.6%	1.0%	0.8%	3%	1.0	N	S	Y
Total	3.7%	8.6%	2.4%	100%	1.0			
Potential	0.2%	2.1%	0.0%		1.0			

Note: the total line is based on the sum of the utilizations and is not simply the sum of the columns.

In the table above, the codes for the "Cap" column mean: A is absolute capping, H is hard capping, S is soft capping, and G is group capping. Soft and group capping are based on GCP-MSUs, but a GCP-capped partition may have an influence on zIIP utilization.

IFL Utilization by Partition (2x IFL Engines)

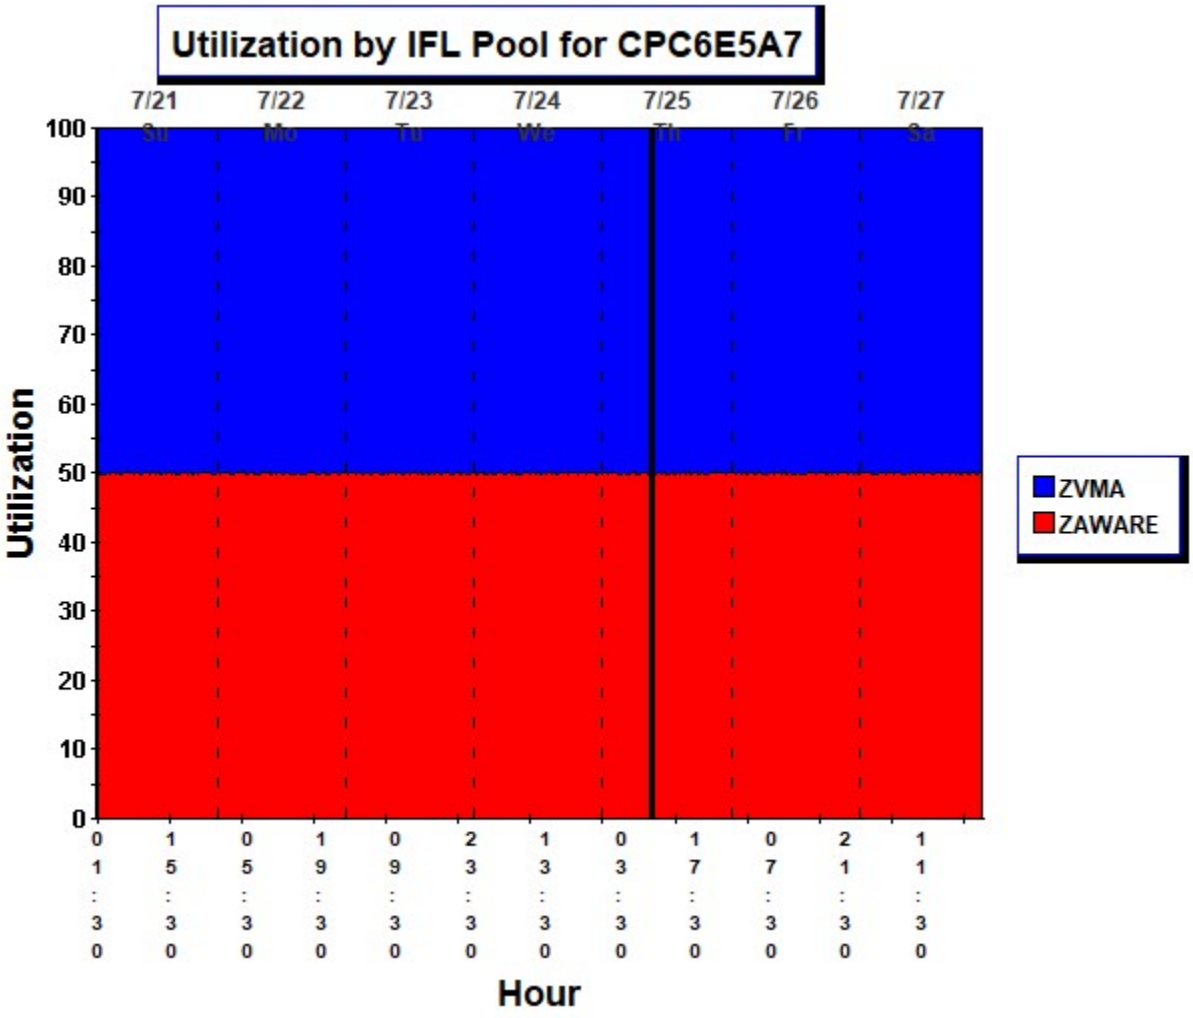

This graph represents the accumulated CPU utilizations for all the partitions running with IFLs under LPAR in this processor. The utilization is the percent of 2.0 engines configured on the 2828-R05.

The IFL utilization in the study interval is 100.0%. The largest component of the utilization in the study interval is contributed by ZVMA. The individual utilizations by SYSID are listed below. The study interval is shown on the graph as a solid line drawn on 7/25/19 at 09:00.

Partition	Study CPU%	Max. CPU%	Avg. CPU%	LPAR Weight	#IFLs	SMT	Capacity	
							Min CPU%	Max CPU%
ZVMA	50.0%	50.0%	50.0%	Ded	1.0	?	50%	50%
ZAWARE	50.0%	50.0%	50.0%	Ded	1.0	?	50%	50%
Total	100.0%	100.0%	100.0%	0	200%			

Note: the total line is based on the sum of the utilizations and is not simply the sum of the columns.

The input data did not contain detailed performance data for one or more partitions using IFL resource. **For partitions using dedicated IFL resource, there is no information available about how much of the dedicated resource is being used, so it appears as 100% used in the graph.** The IFL partitions in which the input EDF data is missing are [ZVMA, ZAWARE].

No partition was found to have either absolute or hard capping of IFL resource.

ICF Utilization by Partition (2x ICF Engines)

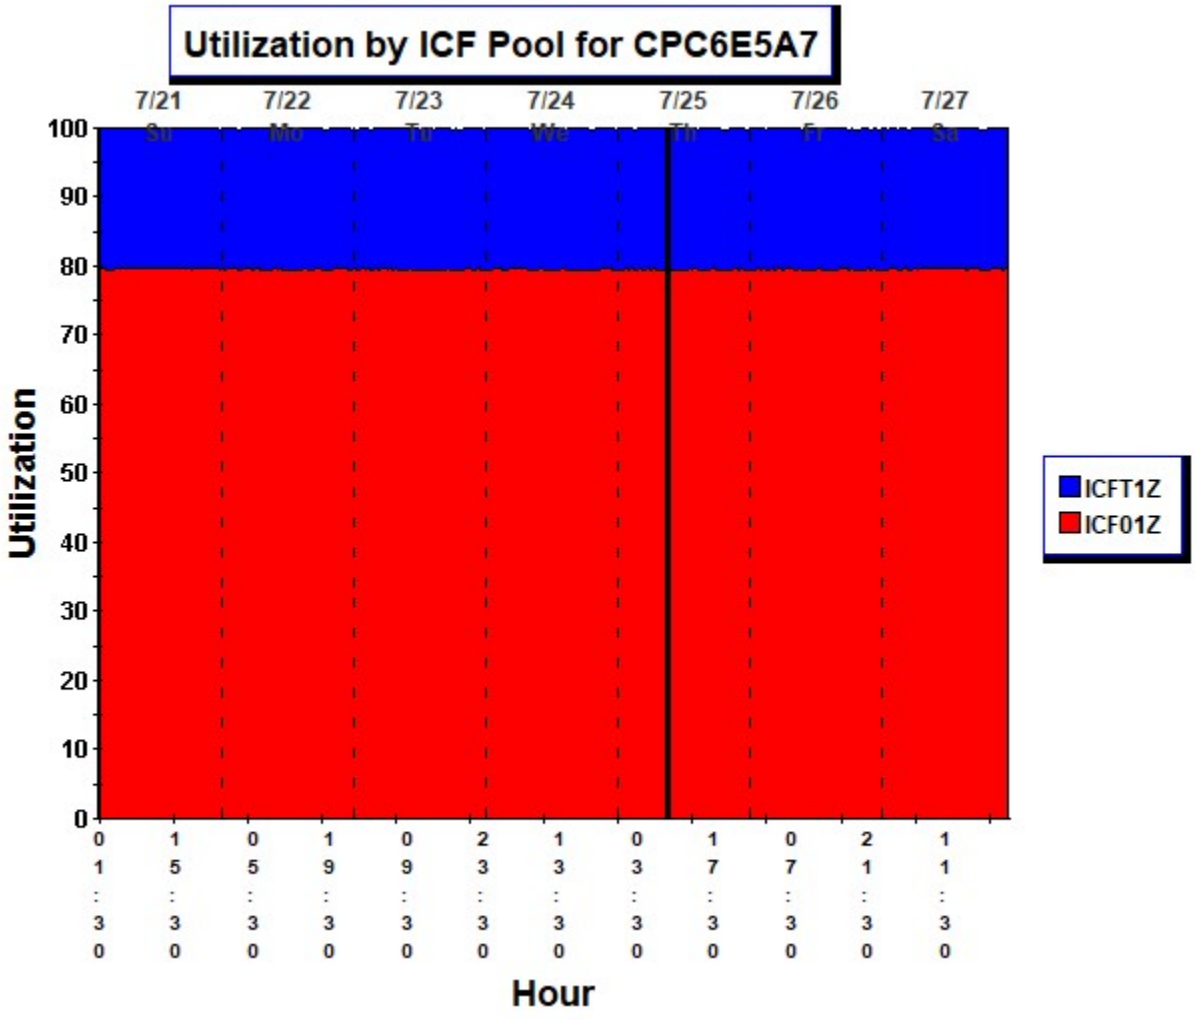

This graph represents the accumulated CPU utilizations for all the partitions running with ICFs under LPAR in this processor. The utilization is the percent of 2.0 engines configured on the 2828-R05.

The ICF utilization in the study interval is 99.8%. The largest component of the utilization in the study interval is contributed by ICF01Z. The individual utilizations by SYSID are listed below. The study interval is shown on the graph as a solid line drawn on 7/25/19 at 09:00.

Partition	Study CPU%	Max. CPU%	Avg. CPU%	LPAR Weight	#ICFs	Effective Engines	Capacity	
							Min CPU%	Max CPU%
ICFT1Z	20.3%	20.3%	20.3%	20%	2.0	0.4	20%	100%
ICF01Z	79.5%	79.6%	79.6%	80%	2.0	1.6	80%	100%
Total	99.8%	99.9%	99.9%	1,000	400%			

Note: the total line is based on the sum of the utilizations and is not simply the sum of the columns.

IBM recommends that dedicated engines should always be used for a CF whose response times are critical. The use of shared engines instead of dedicated engines for the CF always impacts CF response times, and in turn, increases the overhead on the z/OS sysplex members using the coupling facility. The number of effective engines should be kept at or above 1.0 for coupling facilities whose response times are critical. Effective engines of less than 0.5 should be considered non-responsive.

No partition was found to have either absolute or hard capping of ICF resource.

Partition Rolling 4 Hour Average by System

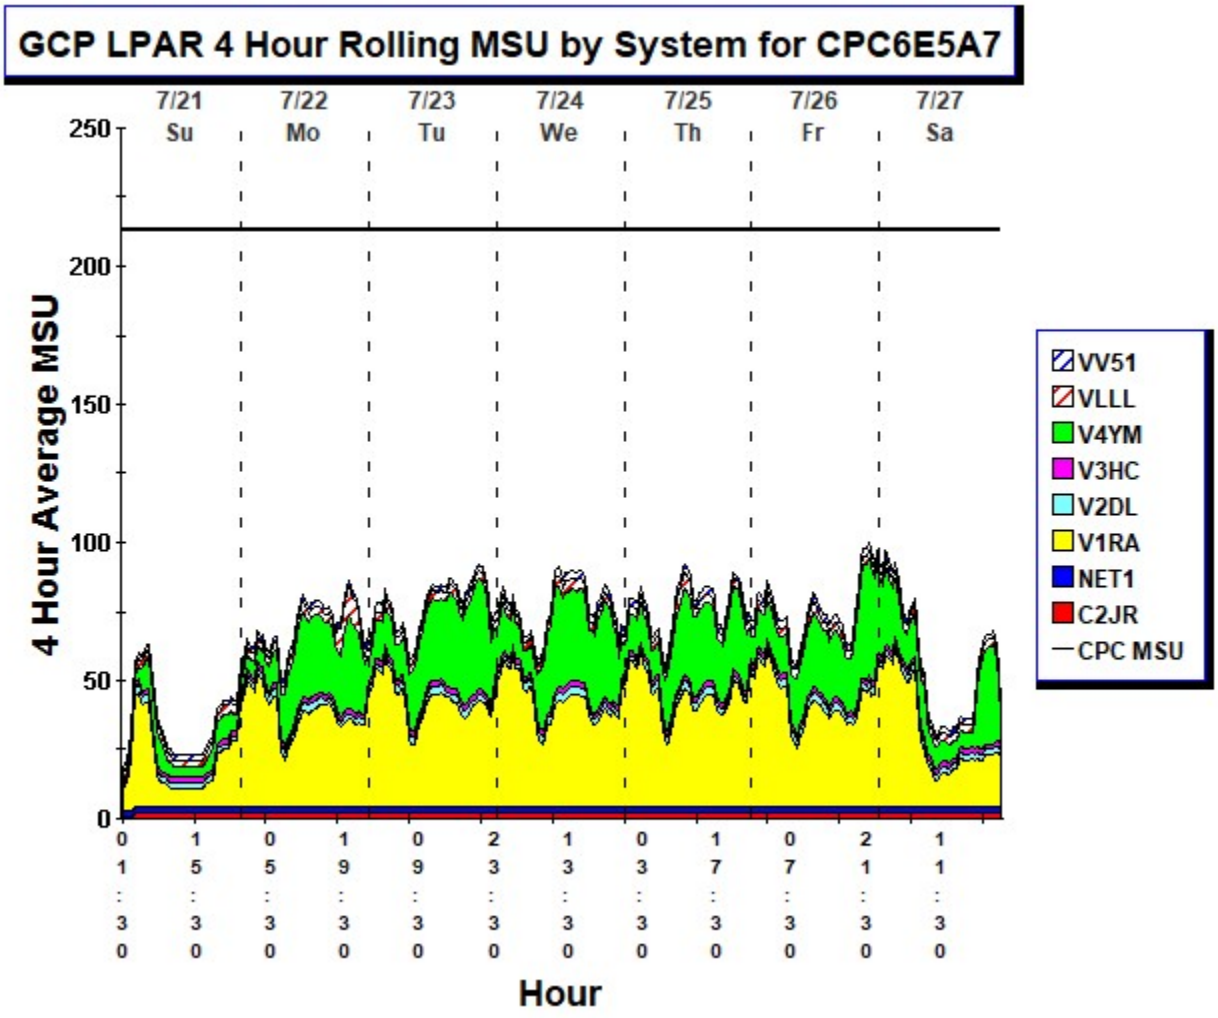

This graph represents the rolling 4 hour average for the processor GCP utilization of CPC6E5A7 and for the 8 GCP partitions running under LPAR on this 2828-R05 (MSU Rating of 213). The values in the graph are the 4 hour average at that time. **The maximum value of 100 at 22:00 on 7/26/19, would be for the previous 4 hours.**

Note: The rolling 4 hours average is very sensitive to the number of samples and intervals selected. And, it is imperative that the samples selected reflect the periods of business interest such as period closing, end-of-month, prime shift, etc.

		Study Interval	Max Interval		
LPAR	WLM Data	MSU	Date	Time	MSU
C2JR	Y	2	7/21/19	04:00	2
NET1	Y	2	7/21/19	01:30	2

V1RA	Y	31	7/27/19	03:00	58
V2DL	Y	3	7/21/19	04:00	3
V3HC	Y	2	7/21/19	04:00	2
V4YM	Y	24	7/26/19	21:30	42
VLLL	Y	4	7/22/19	18:30	10
VV51	Y	3	7/22/19	08:00	3
Total	Y	71	7/26/19	22:00	100

The study interval in the table is 7/25/19 at 09:00 and is the peak interval from the Prime shift.

It is important to note that this analysis presents values either that were measured by WLM or calculated by zCP3000. In either case the Defined Capacity was not taken into account. A Defined Capacity may be lower than either the value that WLM measures or than is calculated by zCP3000. In both cases the customer will never be charged for a value greater than the Defined Capacity.

Memory Allocation by LPAR (112GB Configured)

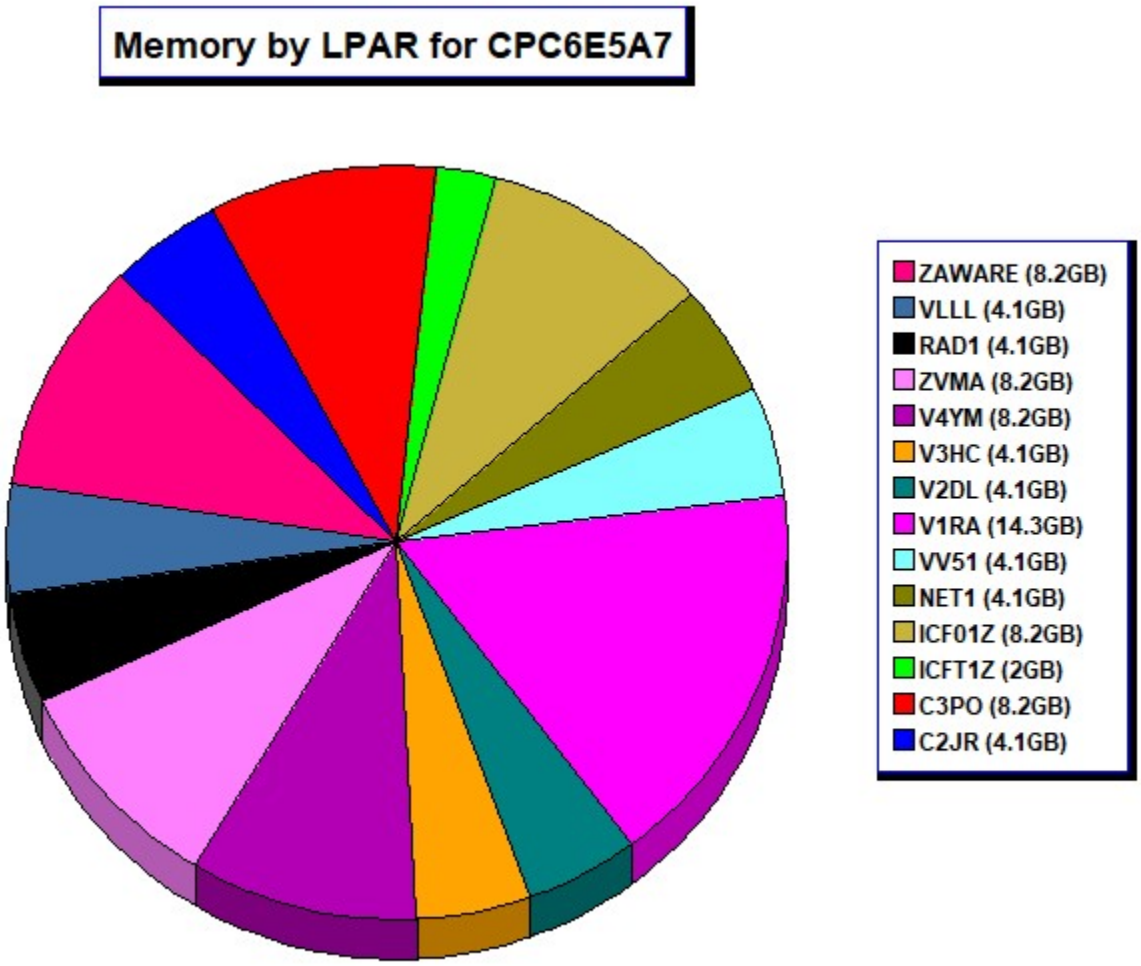

This graph represents the memory defined to each of the partitions on CPC CPC6E5A7. Inactive partitions with no memory defined are omitted.

Partition	Memory MB
C2JR (4.1GB)	4,096
C3PO (8.2GB)	8,192
ICFT1Z (2GB)	2,048
ICF01Z (8.2GB)	8,192
NET1 (4.1GB)	4,096
VV51 (4.1GB)	4,096
V1RA (14.3GB)	14,336
V2DL (4.1GB)	4,096
V3HC (4.1GB)	4,096

V4YM (8.2GB)	8,192
ZVMA (8.2GB)	8,192
RAD1 (4.1GB)	4,096
VLLL (4.1GB)	4,096
ZAWARE (8.2GB)	8,192
Total	86,016

This graph represents the memory configuration in the study interval which is 7/25/19 at 09:00 and is the peak interval from the Prime shift. All other measured intervals have the same memory configuration.

Workload GCP Physical Utilization for V1RA

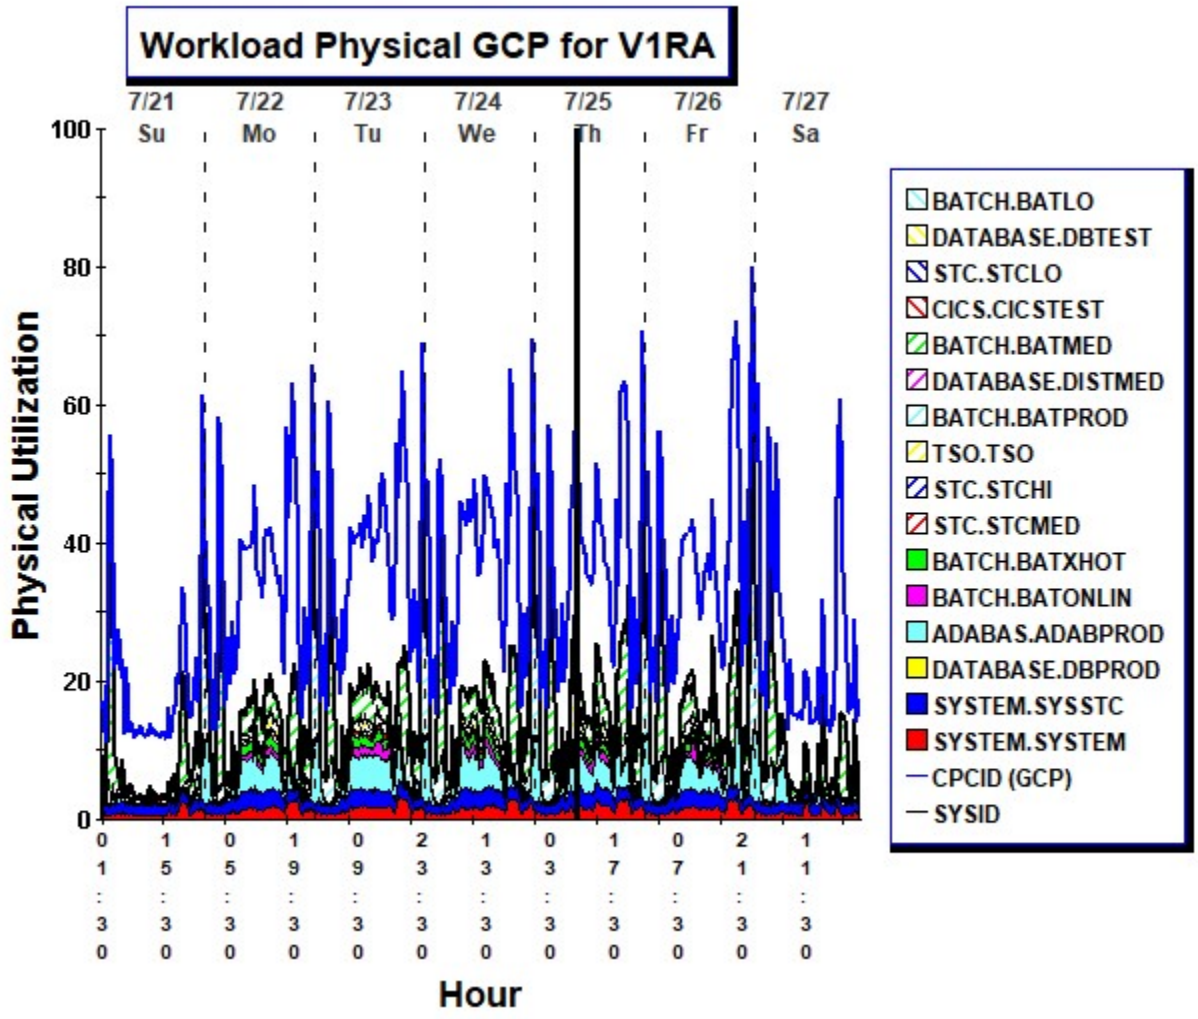

This graph plots the workload CPU samples with the z/OS capture ratio applied. The workload data is accumulated by priority. The partition utilization is also shown.

The CPCID data is the utilization of the CPC. The difference between the partition line and the CPCID line would be the capacity used by the other partitions.

The work on V1RA is broken down into 16 components. With the capture ratio applied, the total for the workloads should be close to the partition value.

The workload which consumed the most CPU over the entire interval was BATCH.BATMED with a utilization of 3.6% in the selected interval. The utilization varied from 0.0% to 42.5% over the entire set of intervals.

The highest priority workload was SYSTEM.SYSTEM. Its utilization in the selected interval was 2.1%. This contribution varied from 0.6% to 3.6%.

The following workloads did not have CPU utilization greater than zero for all sample intervals -

- TSO.TSO
- DATABASE.DISTMED
- CICS.CICSTEST
- BATCH.BATLO

The 90th percentile and peak values in the following table are the workload values at the partition CPU% 90th percentile and peak.

Description	Average	@90th %ile Interval	@Peak Interval	Workload Peak	Peak Date & Time	Study Interval
SYSTEM.SYSTEM	1.3%	2.6%	1.7%	3.6%	7/26/19 20:00	2.1%
SYSTEM.SYSSTC	1.8%	1.6%	1.4%	3.0%	7/22/19 11:00	2.7%
DATABASE.DBPROD	0.0%	0.0%	0.0%	0.3%	7/23/19 20:30	0.0%
ADABAS.ADABPROD	**2.4%**	**3.9%**	**8.8%**	**13.9%**	**7/26/19 20:30**	**4.3%**
BATCH.BATONLIN	0.4%	0.0%	0.0%	3.5%	7/26/19 10:30	1.0%
BATCH.BATXHOT	0.4%	0.0%	0.0%	1.7%	7/22/19 14:30	1.3%
STC.STCMED	0.5%	0.4%	0.4%	1.0%	7/21/19 01:30	1.0%
STC.STCHI	0.2%	0.4%	0.3%	2.5%	7/24/19 13:00	0.3%
TSO.TSO	**0.4%**	**0.0%**	**0.0%**	**14.3%**	**7/25/19 09:00**	**14.3%**
BATCH.BATPROD	**2.8%**	**2.1%**	**36.4%**	**36.4%**	**7/26/19 23:30**	**0.7%**
DATABASE.DISTMED	0.0%	0.0%	0.0%	0.0%	7/24/19 10:00	0.0%
BATCH.BATMED	**5.3%**	**17.0%**	**11.6%**	**42.5%**	**7/21/19 03:30**	**3.6%**
CICS.CICSTEST	0.1%	0.1%	0.1%	1.0%	7/26/19 10:00	0.2%
STC.STCLO	0.2%	0.1%	0.1%	1.6%	7/27/19 12:00	0.3%
DATABASE.DBTEST	0.2%	0.3%	0.2%	2.2%	7/26/19 10:00	0.3%
BATCH.BATLO	0.0%	0.0%	0.0%	0.1%	7/22/19 04:30	0.0%
Total	16.2%	28.5%	61.0%	127.6%		32.1%

Note: The Total line of the Workload Peak column represents the sum of all the peak values. If the workloads were to peak in the same interval, this value would be the utilization.

Workload Central Storage Usage for V1RA

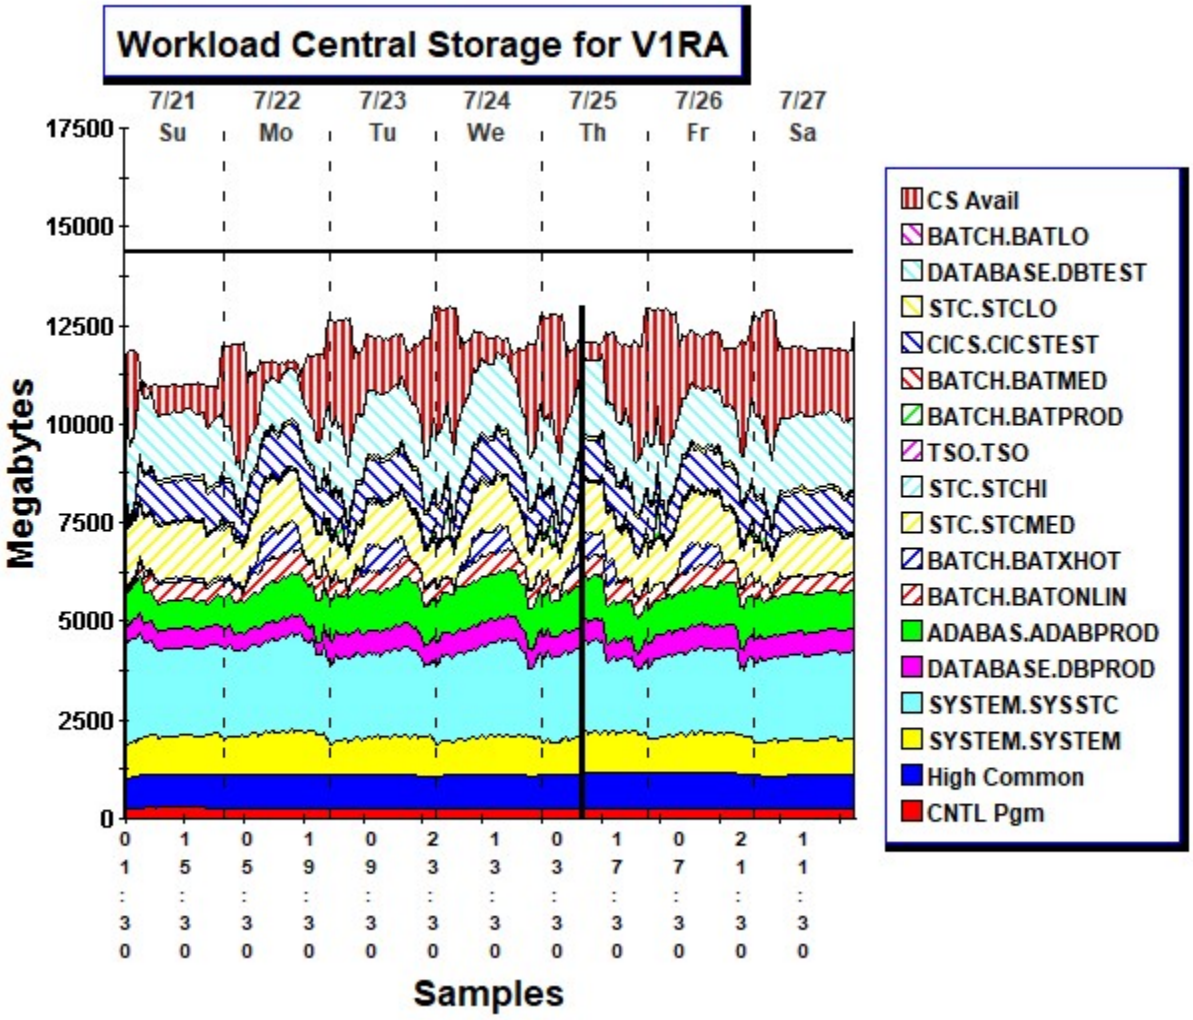

This graph plots the workload Central Storage. The workload data is accumulated. Also plotted is the average amount of Central Storage installed (Central Storage INSTALL) and the average amount of Central Storage used. This is defined as Central Storage_installed - Central Storage_available (Central Storage USED).

The work on V1RA is broken down into 18 workload components. The total for the workloads should be close to the partition value.

The largest amount of Central Storage usage was contributed by SYSTEM.SYSSTC. This was 2,269.9 MB for the study interval. This contribution varied from 1,627.3 to 2,745.3 MB.

The highest priority workload was SYSTEM.SYSTEM. The study interval central memory usage was 947.7 MB. This contribution varied from 770.2 to 1,136.2 MB.

The following workloads did not have Central Storage usage greater than zero for all sample intervals -

- CICS.CICSINTL
- CICS.CICSPROD
- TSO.TSO
- DATABASE.DISTMED
- BATCH.BATMED
- CICS.CICSTEST
- BATCH.BATLO

The space between Central Storage available and the amount of Central Storage online could not be accounted for. In OS/390 this is usually Logical Swap.

The Study Interval in the table below is 7/25/19 at 09:00 and is the peak interval from the Prime shift.

	Central Storage		
Workload	**S. I.**	**Min**	**Max**
CNTL Pgm	276	263	282
High Common	905	727	911
SYSTEM.SYSTEM	**948**	**770**	**1,136**
SYSTEM.SYSSTC	**2,270**	**1,627**	**2,745**
DATABASE.DBPROD	545	323	605
ADABAS.ADABPROD	**1,055**	**572**	**1,169**
CICS.CICSINTL	0	0	0
CICS.CICSPROD	0	0	0
BATCH.BATONLIN	506	25	594
BATCH.BATXHOT	615	6	794
STC.STCMED	**1,228**	**886**	**1,744**
STC.STCHI	70	30	106
TSO.TSO	12	0	12
BATCH.BATPROD	**5**	**0**	**1,043**
DATABASE.DISTMED	0	0	0
BATCH.BATMED	11	0	219
CICS.CICSTEST	**956**	**0**	**1,090**
STC.STCLO	107	85	151
DATABASE.DBTEST	**1,888**	**973**	**3,002**
BATCH.BATLO	0	0	1
CS Avail	632	125	4,002

Workload GCP Physical Utilization for V4YM

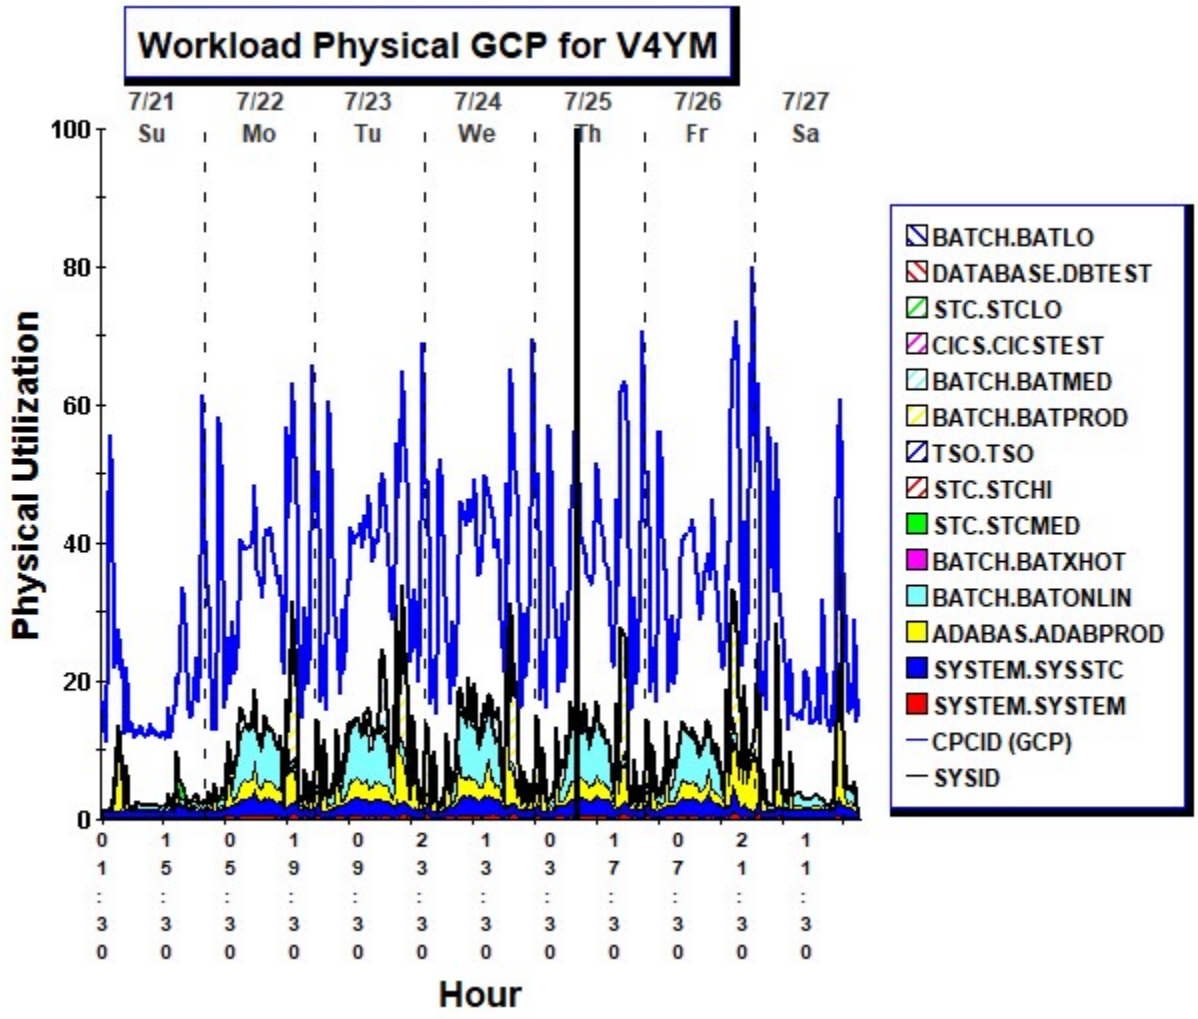

This graph plots the workload CPU samples with the z/OS capture ratio applied. The workload data is accumulated by priority. The partition utilization is also shown.

The CPCID data is the utilization of the CPC. The difference between the partition line and the CPCID line would be the capacity used by the other partitions.

The work on V4YM is broken down into 14 components. With the capture ratio applied, the total for the workloads should be close to the partition value.

The workload which consumed the most CPU over the entire interval was BATCH.BATONLIN with a utilization of 9.3% in the selected interval. The utilization varied from 0.0% to 9.3% over the entire set of intervals.

The highest priority workload was SYSTEM.SYSTEM. Its utilization in the selected interval was 0.5%. This contribution varied from 0.3% to 1.3%.

The following workloads did not have CPU utilization greater than zero for all sample intervals -

- BATCH.BATONLIN
- STC.STCHI
- TSO.TSO
- BATCH.BATPROD
- BATCH.BATMED
- CICS.CICSTEST
- BATCH.BATLO

The 90th percentile and peak values in the following table are the workload values at the partition CPU% 90th percentile and peak.

Description	Average	@90th %ile Interval	@Peak Interval	Workload Peak	Peak Date & Time	Study Interval
SYSTEM.SYSTEM	0.4%	0.5%	0.5%	1.3%	7/26/19 19:30	0.5%
SYSTEM.SYSSTC	1.6%	2.2%	2.6%	3.7%	7/22/19 11:00	2.8%
ADABAS.ADABPROD	**2.1%**	**2.2%**	**16.8%**	**18.6%**	**7/27/19 19:00**	**2.9%**
BATCH.BATONLIN	**2.8%**	**5.8%**	**2.4%**	**9.3%**	**7/25/19 09:00**	**9.3%**
BATCH.BATXHOT	0.0%	0.0%	0.0%	0.2%	7/22/19 20:00	0.0%
STC.STCMED	**0.0%**	**0.0%**	**0.0%**	**2.9%**	**7/21/19 18:30**	**0.0%**
STC.STCHI	0.0%	0.1%	0.4%	1.5%	7/23/19 15:00	0.0%
TSO.TSO	0.1%	1.6%	0.0%	3.2%	7/25/19 11:00	0.2%
BATCH.BATPROD	**2.0%**	**0.2%**	**18.8%**	**22.7%**	**7/24/19 19:00**	**0.1%**
BATCH.BATMED	**0.3%**	**4.9%**	**0.0%**	**10.8%**	**7/23/19 15:00**	**0.1%**
CICS.CICSTEST	0.0%	0.0%	0.0%	0.2%	7/26/19 10:00	0.1%
STC.STCLO	0.0%	0.0%	0.0%	0.0%	7/26/19 19:30	0.0%
DATABASE.DBTEST	0.0%	0.2%	0.0%	0.5%	7/26/19 13:30	0.0%
BATCH.BATLO	0.0%	0.0%	0.0%	0.5%	7/25/19 11:30	0.0%
Total	9.6%	17.8%	41.5%	75.4%		16.1%

Note: The Total line of the Workload Peak column represents the sum of all the peak values. If the workloads were to peak in the same interval, this value would be the utilization.

Workload Central Storage Usage for V4YM

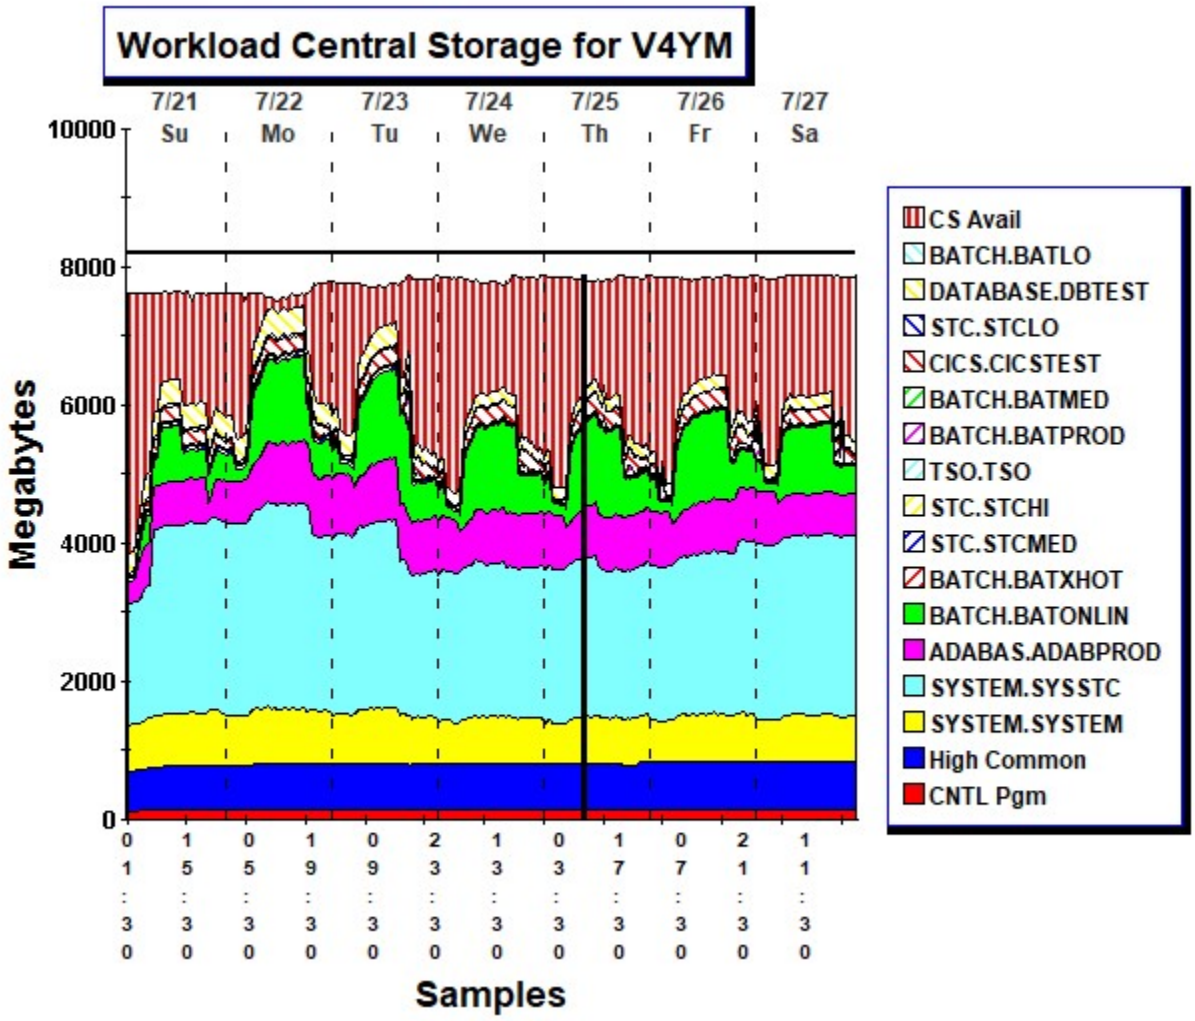

This graph plots the workload Central Storage. The workload data is accumulated. Also plotted is the average amount of Central Storage installed (Central Storage INSTALL) and the average amount of Central Storage used. This is defined as Central Storage_installed - Central Storage_available (Central Storage USED).

The work on V4YM is broken down into 16 workload components. The total for the workloads should be close to the partition value.

The largest amount of Central Storage usage was contributed by SYSTEM.SYSSTC. This was 2,291.6 MB for the study interval. This contribution varied from 1,738.4 to 2,973.2 MB.

The highest priority workload was SYSTEM.SYSTEM. The study interval central memory usage was 661.1 MB. This contribution varied from 576.5 to 845.9 MB.

The following workloads did not have Central Storage usage greater than zero for all sample intervals -

- CICS.CICSINTL
- CICS.CICSPROD
- BATCH.BATONLIN
- TSO.TSO
- BATCH.BATPROD
- BATCH.BATMED
- CICS.CICSTEST
- BATCH.BATLO

The space between Central Storage available and the amount of Central Storage online could not be accounted for. In OS/390 this is usually Logical Swap.

The Study Interval in the table below is 7/25/19 at 09:00 and is the peak interval from the Prime shift.

	Central Storage		
Workload	**S. I.**	**Min**	**Max**
CNTL Pgm	153	128	159
High Common	654	568	696
SYSTEM.SYSTEM	661	577	846
SYSTEM.SYSSTC	**2,292**	**1,738**	**2,973**
ADABAS.ADABPROD	**722**	**250**	**914**
CICS.CICSINTL	0	0	0
CICS.CICSPROD	0	0	0
BATCH.BATONLIN	**1,229**	**0**	**1,303**
BATCH.BATXHOT	2	1	13
STC.STCMED	28	24	82
STC.STCHI	19	7	30
TSO.TSO	0	0	7
BATCH.BATPROD	**0**	**0**	**618**
BATCH.BATMED	1	0	287
CICS.CICSTEST	229	0	259
STC.STCLO	17	16	44
DATABASE.DBTEST	169	167	380
BATCH.BATLO	0	0	2
CS Avail	1,621	142	3,794

Projected Peak Interval Utilization from Prime Shift onto z14 ZR1 Models for Capacity Comparison

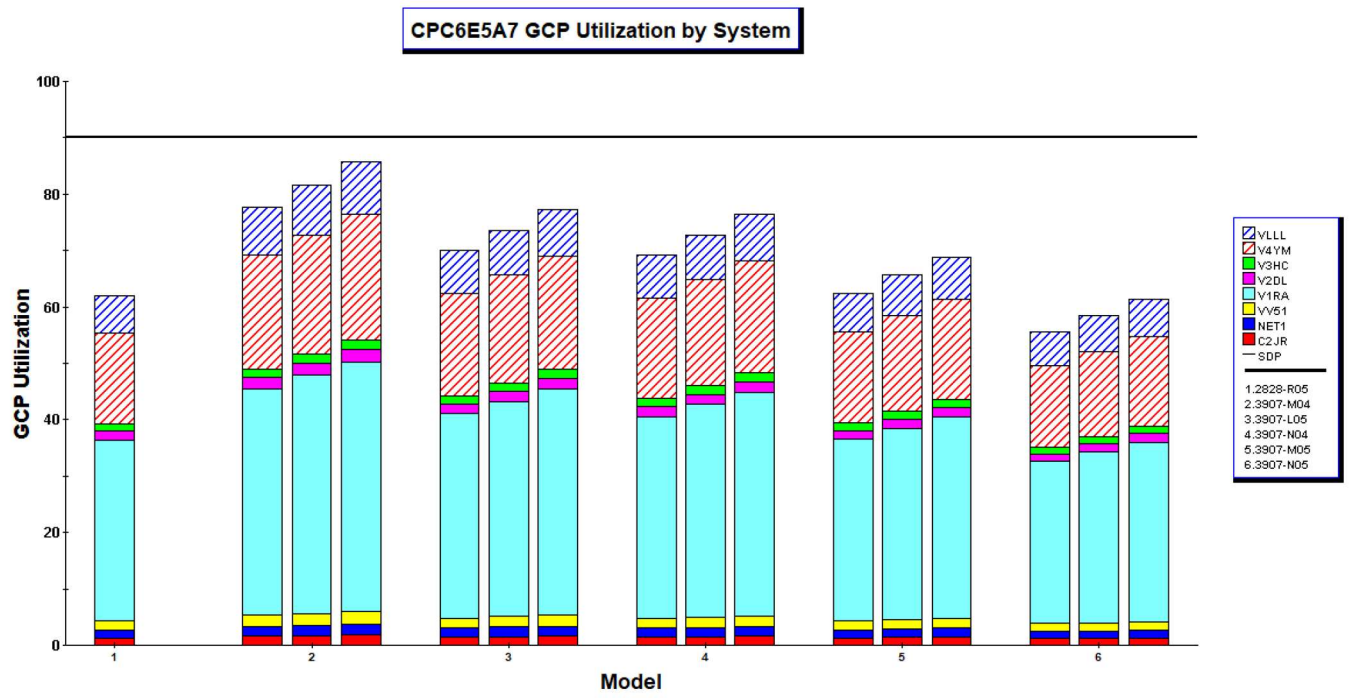

This graph plots the GCP busy value for the partitions running PU type of GCP (8 partitions on CPC6E5A7). The base utilization is computed using 5.0 GCPs. The power value of the GCPs is computed using all the configured PUs on the machine 5 CPs, 2 ICFs, 2 IFLs, 0 zAAPs and 1 zIIPs.

Partition	#PUs	Type	Shared?	z/OS Level	Workload
C2JR	2	GCP	Y	z/OS-2.3*	High
NET1	2	GCP	Y	z/OS-2.3*	High
VV51	2	GCP	Y	z/OS-2.3*	High
V1RA	3	GCP	Y	z/OS-2.3*	Average
V2DL	2	GCP	Y	z/OS-2.3*	High
V3HC	2	GCP	Y	z/OS-2.3*	Average
V4YM	2	GCP	Y	z/OS-2.3*	Average
VLLL	2	GCP	Y	z/OS-2.3*	Average

Three values are given for each alternate. For example, the 3907-M04 has 3 MIPS values [1,165, 1,226, 1,287].

- The center bar for the 3907-M04 is the projected utilization for change from the currently used MIPS value of 1,001 to the MIPS value of 1,226 on the alternate Model for the equivalent workload category.

- The left bar and right bar represent a reasonable expected variation of ± 5% in the attained **GCP utilization percent**. Notice that if the lower CPU% was attained, it would imply a higher relative power ratio to the base.

The following table summarizes the capacities of the processors based on their configurations.

Model Name	MSU	# Of GCPs	GCP Speed	Total GCP	GCP Ratio	# of zIIPs	zIIP Speed	Total zIIP	zIIP Ratio	# of IFLs	IFL Speed	Total IFL	IFL Ratio	Total MIPS
2828-R05	213	5	323	1,613	1.00	1	850	850	1.00	2	1,045	2,089	1.00	4,552
3907-M04	168	4	306	1,226	0.76	1	1,311	1,311	1.54	2	1,553	3,105	1.49	5,642
3907-L05	184	5	272	1,360	0.84	1	1,311	1,311	1.54	2	1,548	3,097	1.48	5,768
3907-N04	189	4	344	1,376	0.85	1	1,311	1,311	1.54	2	1,551	3,103	1.49	5,790
3907-M05	206	5	305	1,526	0.95	1	1,311	1,311	1.54	2	1,547	3,094	1.48	5,932
3907-N05	231	5	343	1,713	1.06	1	1,311	1,311	1.54	2	1,546	3,091	1.48	6,115

Using these configurations we can project the following capacities. Be aware that the utilizations are approximations which do not take into account potential latent demand. Latent Demand is additional work that may get service when capacity is available.

	GCP Utilization			Model					
Model	-5%	Expected	+5%	GCP MIPS	MSU Rating	#GCPs	GCP Config.	zAAP Config.	zIIP Config.
2828-R05		62.1%		1,613	213	5	2 2 2 3 2 2 2 2	0 0 0 0 0 0 0 0	1 1 1 1 1 1 1 1
3907-M04	78%(0.80)	82%(0.76)	86%(0.72)	1,226	168	4	2 2 2 3 2 2 2 2	0 0 0 0 0 0 0 0	1 1 1 1 1 1 1 1
3907-L05	70%(0.88)	74%(0.84)	77%(0.80)	1,360	184	5	2 2 2 3 2 2 2 2	0 0 0 0 0 0 0 0	1 1 1 1 1 1 1 1
3907-N04	69%(0.90)	73%(0.85)	76%(0.81)	1,376	189	4	2 2 2 3 2 2 2 2	0 0 0 0 0 0 0 0	1 1 1 1 1 1 1 1
3907-M05	62%(0.99)	66%(0.95)	69%(0.90)	1,526	206	5	2 2 2 3 2 2 2 2	0 0 0 0 0 0 0 0	1 1 1 1 1 1 1 1
3907-N05	56%(1.11)	58%(1.06)	61%(1.01)	1,713	231	5	2 2 2 3 2 2 2 2	0 0 0 0 0 0 0 0	1 1 1 1 1 1 1 1

In the table above the partitions' configurations are listed in the following order in whole numbers: C2JR, NET1, VV51, V1RA, V2DL, V3HC, V4YM and VLLL.

The MIPS value for a processor running LPAR is a function of the MIPS of the same CPU model in single image mode and the number of partitions, whether the GCPs are shared or dedicated, and the number of GCPs in each partition. This analysis focuses on GCP processors. The MIPS values presented were calculated based on the customer's current configuration including partition configuration and specialty engines defined. The configuration specified above includes all defined processors, but in calculating the MIPS value the parked time, if any, was subtracted from the logical number of engines.

In addition, the MIPS value for any processor is a function of the LSPR workload category. The category for each partition is shown in the first table. This means that the MIPS value is ***not*** a constant.

When an alternate processor is chosen, the number in parentheses in the table is the ratio of the estimated GCP utilization percent. The LPAR overhead will also vary with the processor family. Newer technologies generally have better LPAR implementations.

The reference processor is set to a 2094-701 with an assumed capacity of 593.0 MIPS.

Projected GCP Utilization onto z14 ZR1 Models

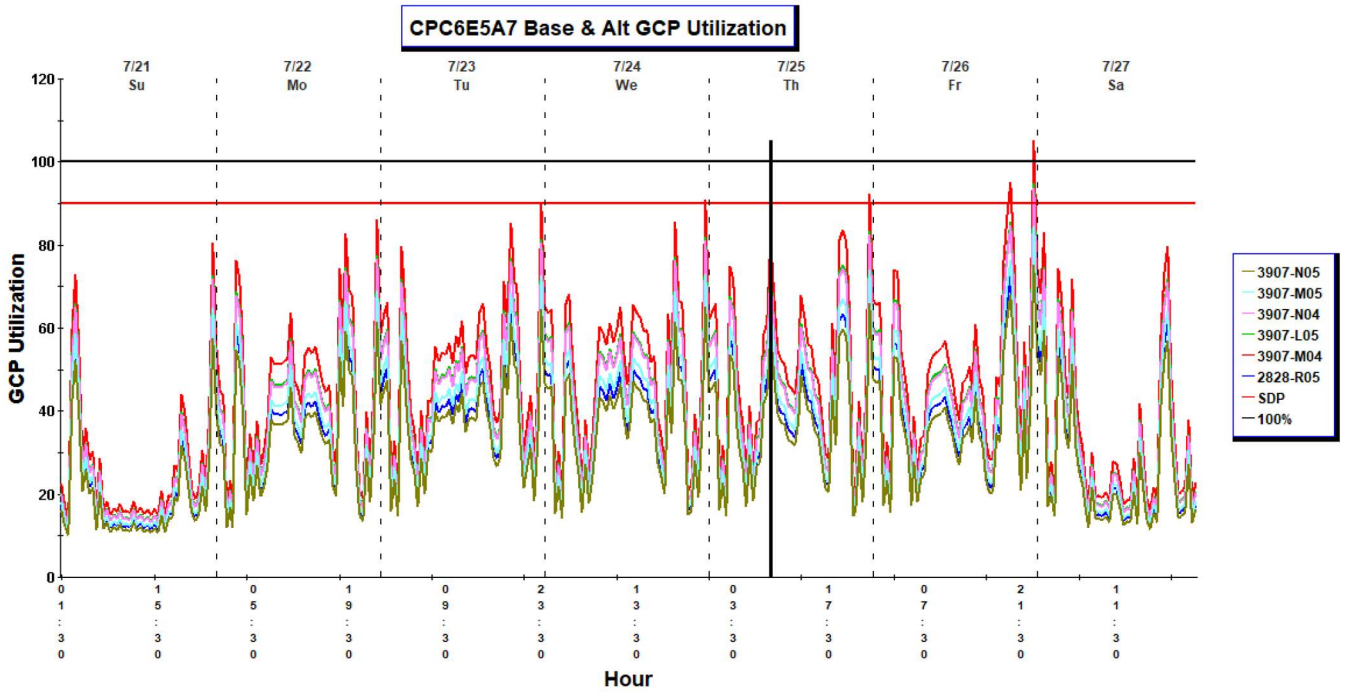

This graph plots the projected GCP utilization for all intervals and for the base processor (2828-R05), and each alternate processor in the study. There is no growth added to the workload. The projected values for the alternate processors are based strictly on the LSPR ratios between the base processor and each alternate.

The Study Interval displayed here is 7/25/19 at 09:00 and is the peak interval from the Prime shift.

Processor	Selected Interval	Average	Maximum
2828-R05	62.1%	33.8%	80.0%
3907-M04	81.7%	44.3%	105.2%
3907-L05	73.6%	39.9%	94.7%
3907-N04	72.8%	39.5%	93.8%
3907-M05	65.6%	35.6%	84.4%
3907-N05	58.5%	31.7%	75.2%

Warning: There are instances of alternate processors exceeding 100%. This workload will not fit on one or more of the defined alternate processors.