Framework for Doing Capacity Sizing on System z Processors

Boston Share
Bradley Snyder
Email Address: bradley.snyder@us.ibm.com
Phone: 972-561-6998

Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

AlphaBlox*
APPN*
CICS*
CICS/VSE*
Cool Blue
DB2*
DFSMS
DFSMShsm
DFSMSrmm
DirMaint
DRDA*
DS6000
DS8000
ECKD
ESCON*
FICON*
FlashCopy*
GDPS*
HiperSockets
HyperSwap
IBM*
IBM eServer
IBM logo*
IMS
Language Environment*
Lotus*
Large System Performance Reference™ (LSPR™)
Multiprise*
MVS
OMEGAMON*
Parallel Sysplex*
Performance Toolkit for VM
PowerPC*
PR/SM
Processor Resource/Systems Manager
RACF*
Redbooks*
Resource Link
RETAIN*
REXX
RMF
S/390*
Scalable Architecture for Financial Reporting
Sysplex Timer*
Systems Director Active Energy Manager
System/370
System p*
System Storage
System x*
System z
System z9*
System z10
Tivoli*
Tivoli Storage Manager
TotalStorage*
VSE/ESA
VTAM*
WebSphere*
xSeries*
z9*
z10
z10 BC
z10 EC
z/Architecture*
Zenterprise 196*
z/OS*
z/VM*
z/VSE
zSeries*

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.
Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.
Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.
IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

* All other products may be trademarks or registered trademarks of their respective companies.

Notes:
Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.
IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.
All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.
This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.
All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.
Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.
Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

Agenda

- **Capacity Planning or CPU Sizing?**

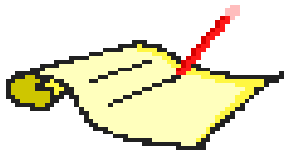

- **End to End Process of CPU Sizing**
- **Summary**
- **Suggested Review Sessions:**
  - zPCR Capacity Sizing/Advanced Lab – Part 1 - Introduction and Overview
  - zPCR Capacity Sizing/Advanced Lab – Part 2 - Hands on Lab

Is it Capacity Planning or CPU Sizing?

- **Terms are often used interchangeably, but they mean different things, and imply different activities**
- **Capacity Planning**
  - Ongoing, with system utilization checked against a multi-period plan
  - Evaluates new applications
  - Identifies and manages workload growth at a business function level
  - Goal of forecasting capacity upgrades 3-6 months in advance
- **CPU Sizing**
  - Done in preparation for a processor change
  - One time effort
  - Aimed at verifying a proposed change

End to End CPU Sizing Process

- **Describe the steps and considerations in the process**
  - Identify points where expectations should be clearly set
- **Identify areas which cause increased complexity and may raise the risk associated with the plan**
- **Identify practical approaches to handling unknowns**

1. Solicit input
2. Evaluate current system(s) performance
3. Create Capacity Relationships of Current Processors
4. Establish “End Game” configuration
5. Establish Capacity Relationships of Future Processors
6. Generate the Plan
7. Identify Post-Install Requirements

Acceptable Use of MIPS

- **It is acceptable to use a MIPS designation for a processor in the planning process as long as the capacity ratios between relative processors agrees with the output of a zPCR study!**
  - Do Not use primitive LSPR data or single number MIPS tables because they do not include LPAR effects of specific processor configurations
  - zPCR is based on LSPR information but factors additional information into the relative capacity relationships it creates
    - zCP3000 uses zPCR for detailed capacity planning

Solicit Input and Document Assumptions

- **Understand rational for the processor change**
- **Identify key parameters involved in the study**
  - Data requirements
  - Specific time of day to evaluate capacity
  - Client defined MIPS ratings for current processors
    - Planning process will define MIPS ratings for proposed processors
  - Available information on growth rates or new workloads
- **Identify key capacity guidelines, i.e.,**
  - New processor can't be more than 90% busy
  - Certain LPARs can't be on the same footprint
  - Batch window can't elongate
  - Etc.

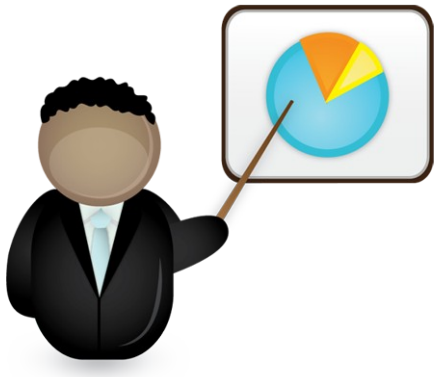

Evaluate Current System(s) Performance Data

- **CPU Sizing ASSUMES the system is well tuned**
- **Generally SMF Records 70:78 are used for Analysis**
  - SMF 30 records sometimes used
- **A good planning process will still make some rudimentary checks to evaluate the performance of the system**
  - Latent demand in an LPAR
  - Latent demand in a CP (single TCB architectures)
  - Latent demand in Job queues
  - Consistently high utilization
  - Well-running I/O subsystem
  - No processor storage contention
  - Good z/OS capture ratio
- **Evaluate the WLM setup to ensure the workloads have enough granularity to get a reasonable view of the system**
  - Need to look at the report class granularity

Performance Data – Red Flags

- **Uneven Utilization patterns**
  - Could have been an outage, problem, holiday, etc.
  - Identify and decide if need to eliminate data
- **Low utilization**
  - Processor utilization affects the efficiency hardware and software
- **High amounts of Latent Demand**
  - Needs to be identified in the plan
- **Poorly performing I/O subsystem**
- **Processor storage contention**

A Few Charts can tell a lot....

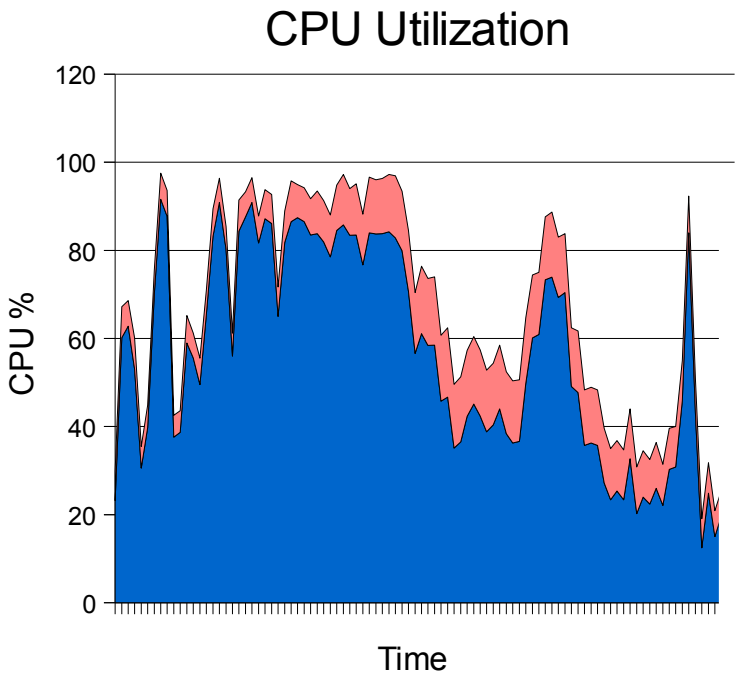

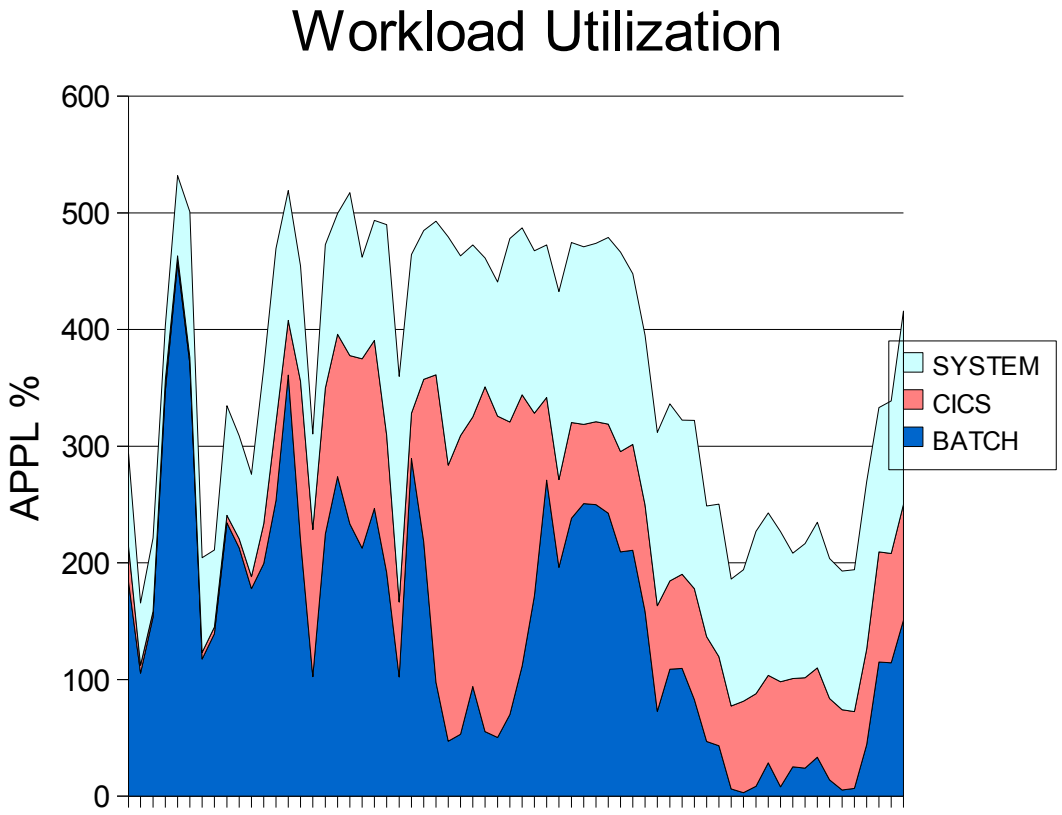

Describe the Current Environment

- **Identify current processors involved in the study**
- **Use pre-defined default processor in zPCR**
  - All Processors MUST use same base, or ratios are invalid
  - Most likely all MIPS values will need to be adjusted
  - Ratios between processors will now be valid
- **Pick pre-defined workload mix for each LPAR**
  - Description of dominant LPAR is often sufficient
  - Verify the mix for each identified time period
    - Prime shift peak hour
    - Key batch window
    - Monthly/Quarterly/Yearly close

Will need to rerun zPCR to establish proper capacity ratios

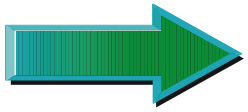

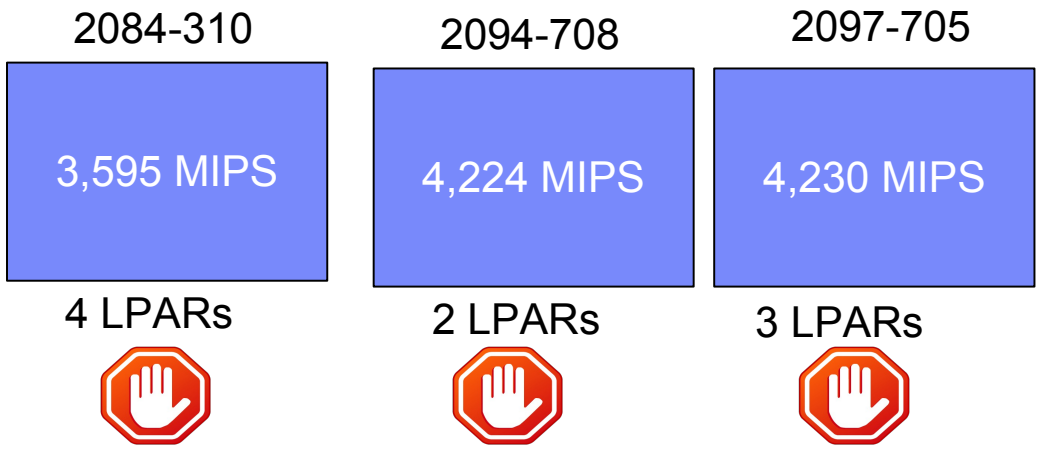

zPCR Workload Mixes

- **Do NOT use LSPR primitives to describe capacity relationships**
  - z/OS V1R9 primitves:
    - ODE-B, CB-L, WASDB, OLTP-T, and OLTP-W
- **IBM recommends using pre-built mixes**
  - Most customer workloads will fit closely with one of several pre-built mixes in zPCR
  - z/OS V1R9 mixes are listed below, used in zPCR V6.3c
    - LoIO-Mix
    - CB-Mix
    - TM-Mix
    - TD-MIX
    - TI-Mix
    - Web-Mix
    - LSPR-MIX
  - New Mixes for zPCR V7.1c with z/OS V1R11
    - Low
    - Average
    - High

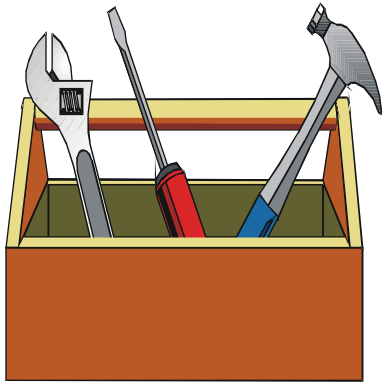

Generate Capacity Relationships of Current Processors

- **Input into zPCR**
  - Number of Partitions
  - Number of processors
    - Include all Specialty CPs
  - Workload Mix
- **Example**
  - Relative capacity vs. 2094-701 set to 1.00
  - zPCR MIPS rated with 2094-701 set to 593 MIPS

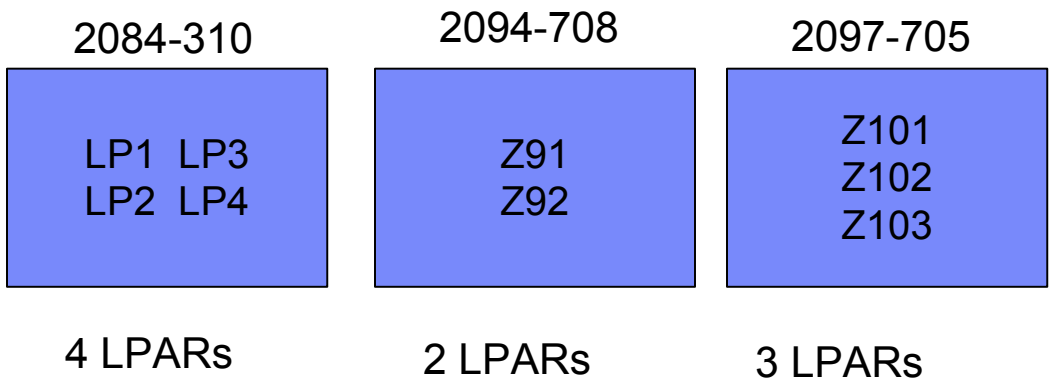

Processor	Relative Capacity**	New zPCR MIPS**	Previous 'MIPS' Value
2084-310	6.3667	3,778	3,595
2094-708	7.1929	4,265	4,224
2097-705	6.8331	4,052	4,230
Total MIPS		12,095	Total MIPS

** based on all LPARS rated with AVERAGE mix

Generate Capacity Data

- **Using performance data, generate capacity requirements**
- **Identify any capacity needed for latent demand**
  - Evaluate all appropriate time periods
- **Create a table for each current processor**

2084-B16 310 – 3,778 MIPS Batch Window Online Window

LPAR	Weight	# LPs	MIPS	MIPS @ 90%	MIPS	MIPS @ 90%
LP1	600	7	957	1,034	1,300	1,990
LP2	300	5	345	433	955	1,115
LP3	50	3	155	201	201	355
LP4	50	2	45	56	145	185
Total	1000	24	1,502	1,724	2,591	3,645

Repeat table for each processor to be upgraded

Document “End Game” Configuration

- **LPAR Layout**
  - Number and types of LPARs
  - Specialty CPs
  - Number of Books

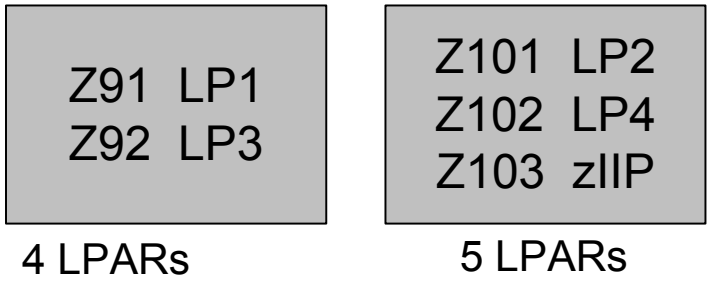

- **Layout the capacity data and determine the new LPAR weights and number of logical CPs**
- **Determine number of logical CPs for each partition**

zEnterprise 196 Sizing

- **Need Processor that delivers 6,600 MIPS**
- **Use Peak Hour to set LPAR definitions**
  - If using IRD, be sure MAX/MIN ranges are appropriate

	Batch Window		Online Window	
LPAR	MIPS	MIPS @ 90%	MIPS	MIPS @ 90%
Z91	1,450	1,625	2,325	2,550
Z92	345	545	1,200	1,715
LP1	957	1,034	1,300	1,990
LP3	155	201	201	355
Total	2,907	3,405	5,026	6,600

zPCR Capacity Data for 2817-706

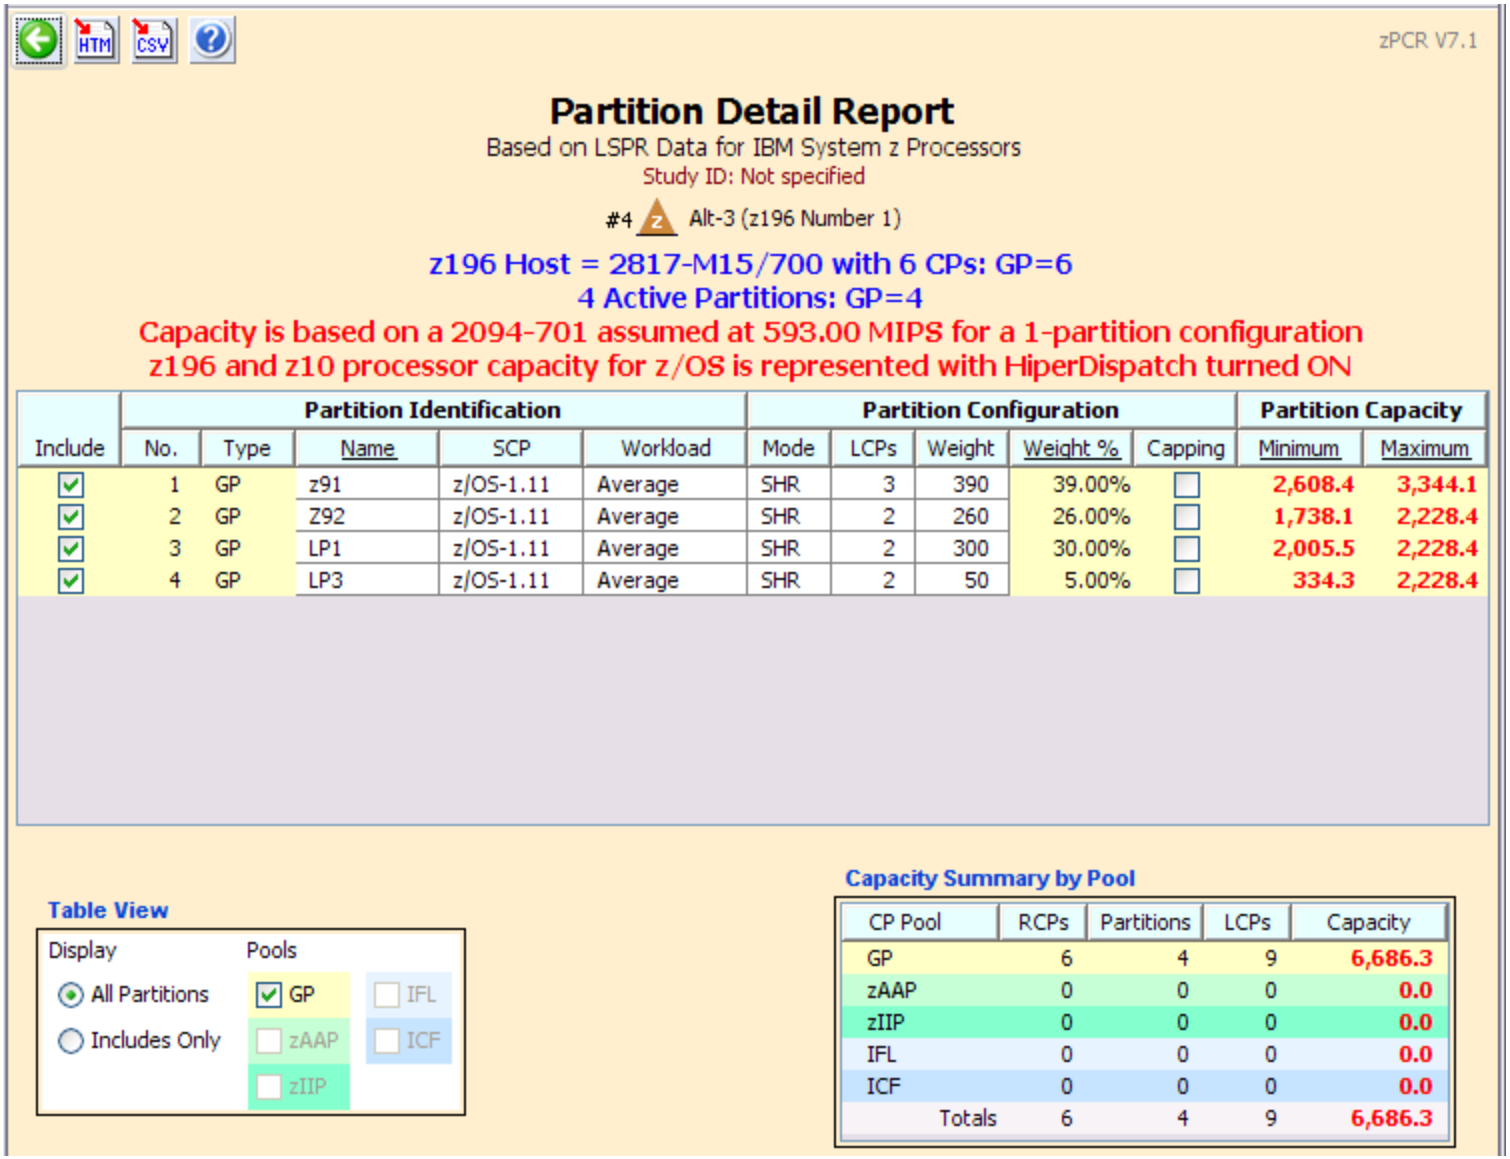

zPCR V7.1

Partition Detail Report

Based on LSPR Data for IBM System z Processors

Study ID: Not specified

#4 Alt-3 (z196 Number 1)

z196 Host = 2817-M15/700 with 6 CPs: GP=6

4 Active Partitions: GP=4

Capacity is based on a 2094-701 assumed at 593.00 MIPS for a 1-partition configuration

z196 and z10 processor capacity for z/OS is represented with HiperDispatch turned ON

	Partition Identification					Partition Configuration					Partition Capacity	
Include	No.	Type	Name	SCP	Workload	Mode	LCPs	Weight	Weight %	Capping	Minimum	Maximum
☑	1	GP	z91	z/OS-1.11	Average	SHR	3	390	39.00%	☐	2,608.4	3,344.1
☑	2	GP	Z92	z/OS-1.11	Average	SHR	2	260	26.00%	☐	1,738.1	2,228.4
☑	3	GP	LP1	z/OS-1.11	Average	SHR	2	300	30.00%	☐	2,005.5	2,228.4
☑	4	GP	LP3	z/OS-1.11	Average	SHR	2	50	5.00%	☐	334.3	2,228.4

Table View

Display: ◉ All Partitions ○ Includes Only

Pools: ☑ GP ☐ zAAP ☐ zIIP ☐ IFL ☐ ICF

Capacity Summary by Pool

CP Pool	RCPs	Partitions	LCPs	Capacity
GP	6	4	9	6,686.3
zAAP	0	0	0	0.0
zIIP	0	0	0	0.0
IFL	0	0	0	0.0
ICF	0	0	0	0.0
Totals	6	4	9	6,686.3

Capacity Relationships for Proposed Processors

- **In this case, the three existing processors plan to be consolidated to two z196s**
  - Existing z10 and z9 upgraded to new z196
  - Total capacity for two CECs needed to be enough for existing workload, plus workload growth and include effects of new zIIP processor
- **Prior to upgrade, zPCR MIPS for three processors was combined 12,095 MIPS**
- **After upgrade, zPCR MIPS for two processors is combined 14,260 MIPS**

Processor	Relative Capacity	New zPCR MIPS
2094-701	1.000	593
2817-706	11.2754	6,686
2817-707	12.772	7,574
Total MIPS		14,260

MIPS Value in Table does not include extra Capacity for zIIP

LPAR Impacts on Capacity

- **n-way and MP effects will impact capacity**
- **LPAR 3 is a uni, but the hardware is running as an 8-way shared processor and the capacity is of an 8-way shared processor**
  - 5 GCPs, 2 zIIPs, 1 zAAPs
- **Number and how busy they are will affect capacity**
- **Only zPCR can help determine what true capacity delivered is**

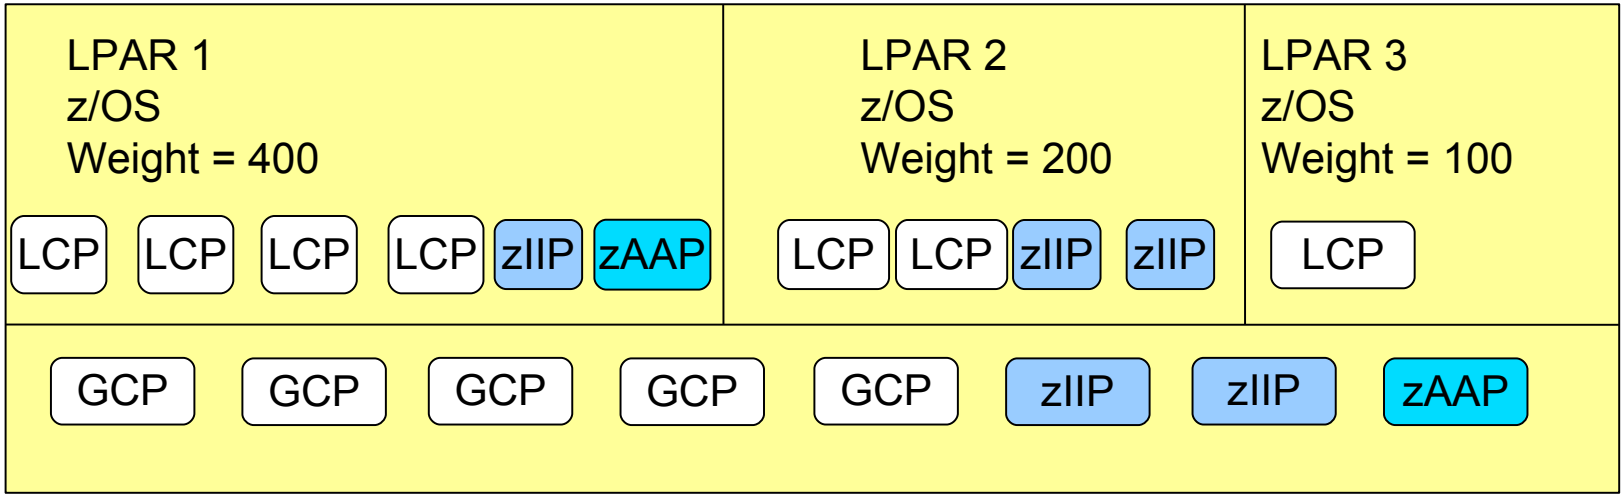

Capacity Planning and LPAR

- Examples of single z196 CEC with multiple LPAR configurations
  - On z196 or z10 with HIPERDISPATCH=YES, or z9 with IRD Vary CPU Management, logical engine configuration will closely match what is guaranteed by LPAR weight
- ITRRs shown are relative to z9-701 set at 1.00

Case	Mode	# of LPs	LP x LCP	LCP	ITRR	% Change	LCP:PCP
Base	2817-720	1	1 x 20	20	31.680	Base	1:1
1	2817-720	2	2 x 10	20	33.279	5%	1:1
Base	2817-710	6	3 x 10 2 x 3 1 x 2	38	16.570	Base	3.8:1
1	2817-710	6	1 x 3 4 x 2 1 x 1	12	17.676	6.7%	1.2:1
2	2817-710	10	10 x 1	10	17.292	4.4%	1:1

Impact of Specialty CPs

- **Impact of Specialty Engines on GP CPUs**
  - Impact will vary based on utilization of specialty CP's
    - Can be slight (less than 10%) to the impact of a full n-way impact of another GP CPU
- **Capacity is characterized as independent partitions with their own LCPs that compete for resources within their assigned CP pool**
- **zPCR is the best source for Specialty CP Impact**
- **Estimation given in zPCR assumes specialty processors are 90% busy**
  - Example: impact of 6 zAAPs running at 50% busy

Specialty CP Example

Partition Type	2094-712 with no zAAPs			2094-712 with 6 zAAPs			% Change
	CPs	LCPs	Capacity	CPs	LCPs	Capacity	
GP	12	12	5,946	12	12	5,349	-10%
zAAP	0	0	0	6	7	2,712	--

- **zAAPs are running at 50% busy**
  - Physically 50% busy, combined usage of both LPARs
- **Capacity of GP CPs with zAAPs is 10% less than without zAAPS**
  - Reduction of 597 MIPS
- **If zAAPs are only 50% busy, than GP CP capacity is expected to be 5,648 instead of 5,349 (1/2 of 597 MIPS)**
- **zAAP capacity adds 2,712 MIPS to environment with no MSU change**

LPAR Utilization Cautions

- **Lightly weighted LPARs might need more capacity when moving to newer, faster processors**
  - A capacity planning concern when weight of LPAR will be less than 50% of an engine on new processor
  - Explore potential LPAR consolidation
    - Reduces need to run z/OS on uniprocessor
    - Virtual storage constraints need to be reviewed
- **zPCR will not model or be indicator of this issue**

Estimation Confidence

- **Major Configuration Changes**
  - Accuracy of zPCR model for an upgrade is +/- 5% of the estimate
  - Variability comes from multiple sources
    - Workload mix used is an estimate, actual workload can vary throughout time
    - Interactions of LPAR peaks and valleys
    - Efficiency of buffering techniques which impact I/O, and hence quantity of interrupts, which drives rate of preemption
    - Hardware changes made after LSPR benchmarks
- **Minor Configuration Changes**
  - Adding 1 LPAR, 1 engine, or changing number of LCP
  - Much higher confidence
  - Newer versions of zPCR will include information on scope of change
- **Capacity decisions should be made with knowledge of the confidence factors**

Post Install Analysis

- **Success Factors:**
  - Evaluation is done as close to the install of the new processor as possible
    - Rebuild the capacity expectations to match the installed configuration
  - Critical applications are isolated into WLM definitions which allow a clear view of capacity
- **Performance data is retained and available for analysis**
- **Changes not included in capacity estimation but should be factored**
  - Change in operating system or middleware levels
  - Maintenance
  - Change in processor storage (impacts sort-based workloads)
  - Buffer pool changes
  - Use of dynamic SQLs
  - Rebinding of SQL on new processor

Summary

- **Long ago, LPAR environments and associated complexity have caused straight MIPS charts to become obsolete**
- **Use pre-built mixes in zPCR and zCP3000**
  - Use CPUMF Data where possible when upgrading from z10 to z196
- **Understand the current system performance and latent demand indicators of an upgrade**
- **Use tools like zPCR / zCP3000 to get the best view of expected capacity**
- **Set expectations with knowledge of confidence factors**
  - Confidence factor of +/- 5% on all upgrades

Additional Resources

- **Understanding the impacts of LPAR on a uni-processor**
  - Managing CPU-Intensive Work on Uniprocessor LPARs - white paper WP100925
  - http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP100925
- **Running IBM System z at High Utilization**
  - Running and how to manage processors at high utilizations – white paper WP101208
  - http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101208