Framework for doing Capacity Sizing for zSeries Processors

SHARE - Winter, 2005
EWCP Project
Session: 2517

Kathy Walsh
IBM
Washington Systems Center
Advanced Technical Support

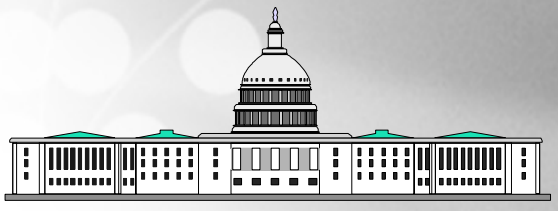

Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries. For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml: AS/400, DBE, e-business logo, ESCON, eServer, FICON, IBM, IBM Logo, iSeries, MVS, OS/390, pSeries, RS/6000, S/390, VM/ESA, VSE/ESA, Websphere, xSeries, z/OS, zSeries, z/VM

The following are trademarks or registered trademarks of other companies

Lotus, Notes, and Domino are trademarks or registered trademarks of Lotus Development Corporation
Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries
LINUX is a registered trademark of Linux Torvalds
UNIX is a registered trademark of The Open Group in the United States and other countries.
Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.
SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC.
Intel is a registered trademark of Intel Corporation
* 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.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

Any proposed use of claims in this presentation outside of the United States must be reviewed by local IBM country counsel prior to such use.

The information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

Agenda

- Is it CPU Sizing or Capacity Planning?
- End to End Process of Capacity Planning
  - Role of Performance Management
  - Describing the Current Environment
  - Consolidation Considerations
  - Specialty Processors
  - Estimation Confidence
  - Post-Install Analysis
- Summary

Is It Capacity Planning or CPU Sizing?

- Terms are often used interchangeably, but they mean different things, and imply different activities

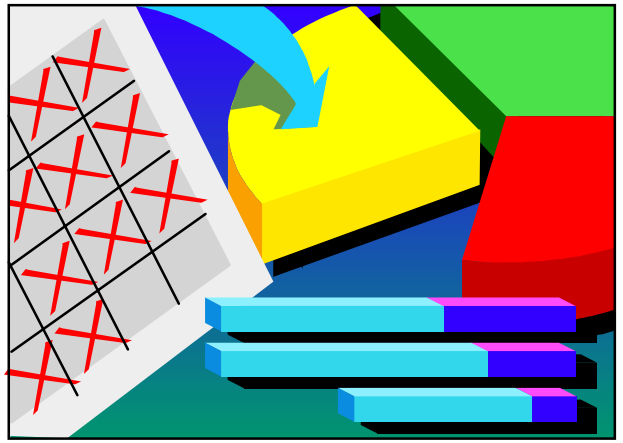

- CPU Sizing
  - Done in preparation for a processor change
  - One time effort
  - Aimed at verifying a proposed change
- Capacity Planning
  - Ongoing, with system utilization checked against a multi-period plan
  - Evaluates new applications
  - Identifies and manages workload growth at a business function level
  - Goal of forecasting capacity upgrades 3-6 months in advance

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

  - Solicit Input
  - Evaluate Current System(s) Performance
  - Create Capacity Relationships of Current Processors
  - Establish "End Game" configuration
  - Establish Capacity Relationships of Future Processors
  - Generate the Plan
  - Set Capacity Expectations
  - 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
  - Not LSPR information because LSPR ratios 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

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 MIP ratings for current processors
    - Planning process will define MIP 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.

Obtain Performance Data

- CPU Sizing process assumes the system is well-tuned
- Generally RMF records 70:78 are used for the analysis
  - SMF 30 records are also 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, software/network problem, holiday, etc.
  - Identify and decide if need to eliminate the data
- Low utilization
  - Processor utilization affects the efficiency of the 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

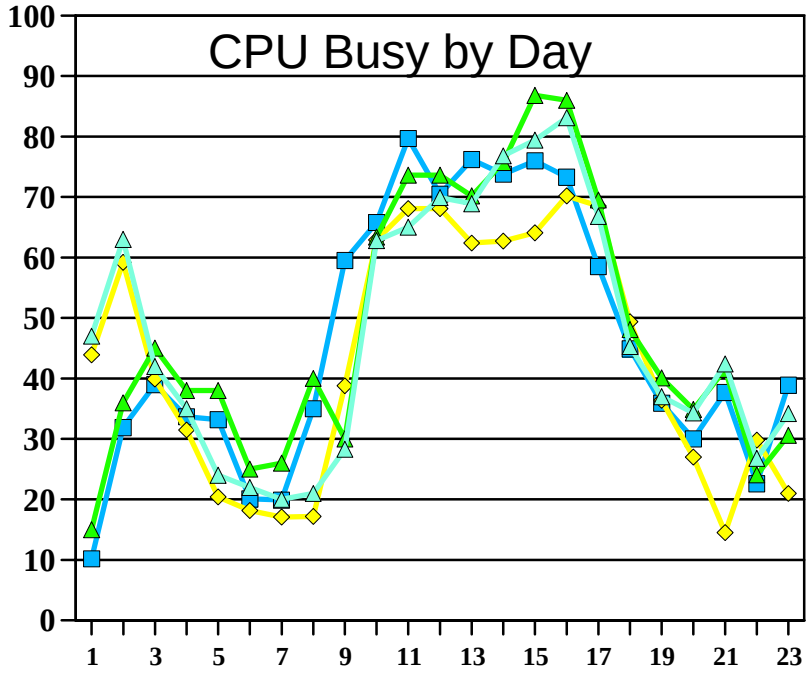

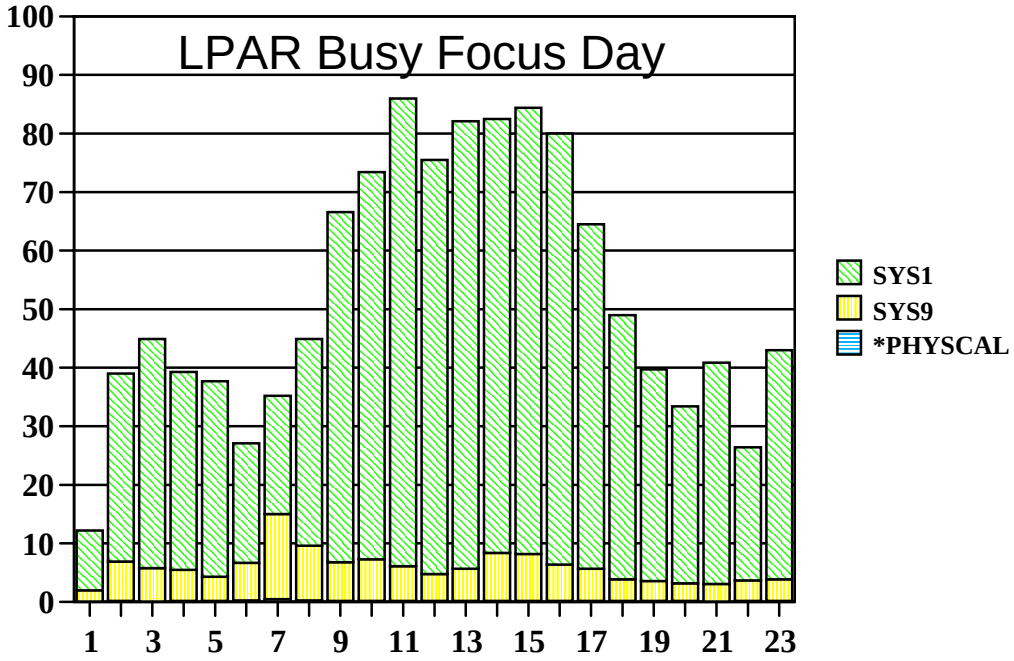

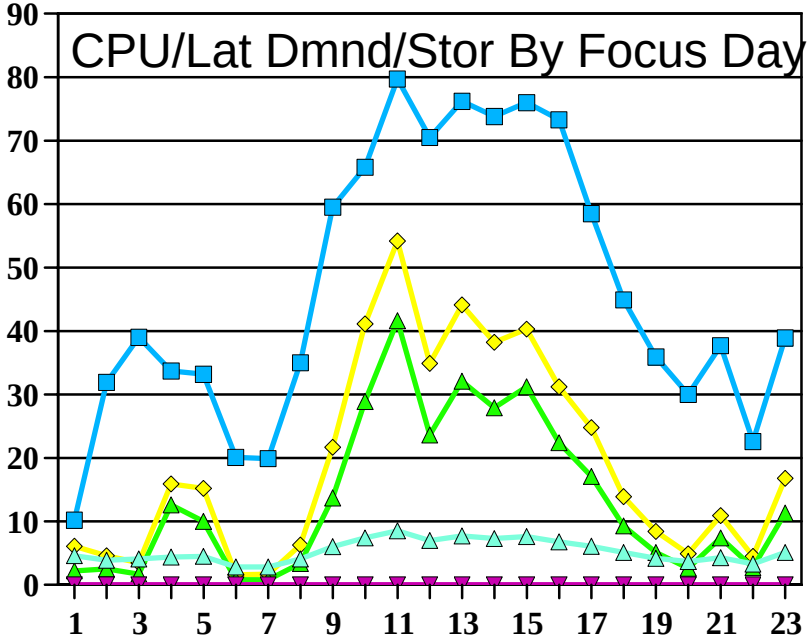

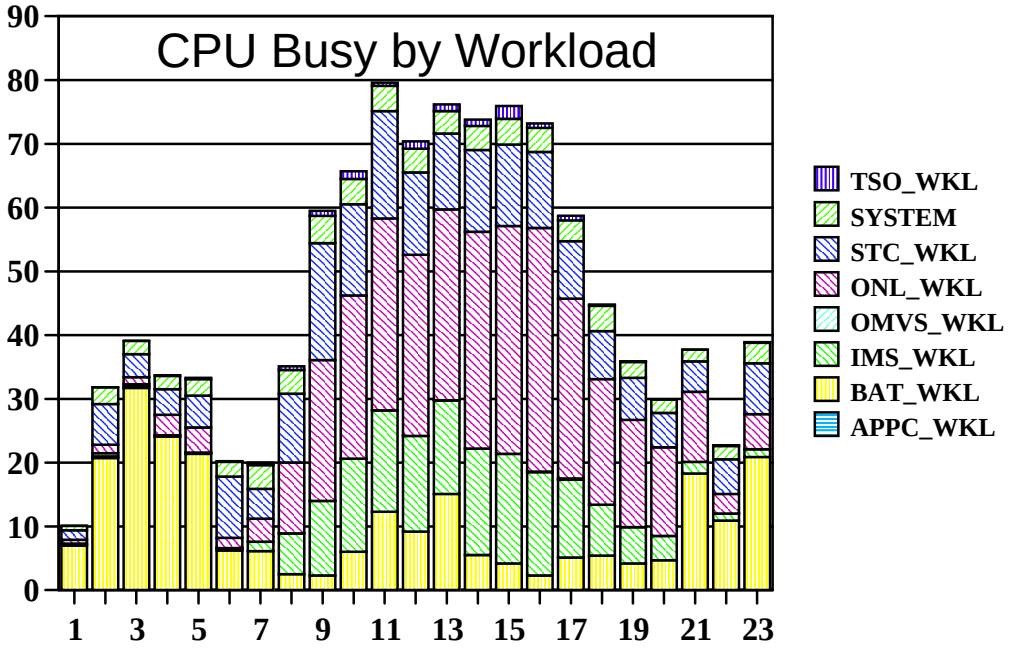

Describe the Current Processor(s)

- Identify current processors involved in the study
- Create a base processor
  - Use only one CEC's input on processor size, even if more CECs are involved
- Calculate custom workload mix for each processor
  - Description of dominant LPAR is often sufficient
  - Verify the custom mix for each identified time period
    - ✓ Prime shift peak hour
    - ✓ Key batch window
    - ✓ Monthly/Quarterly/Yearly close

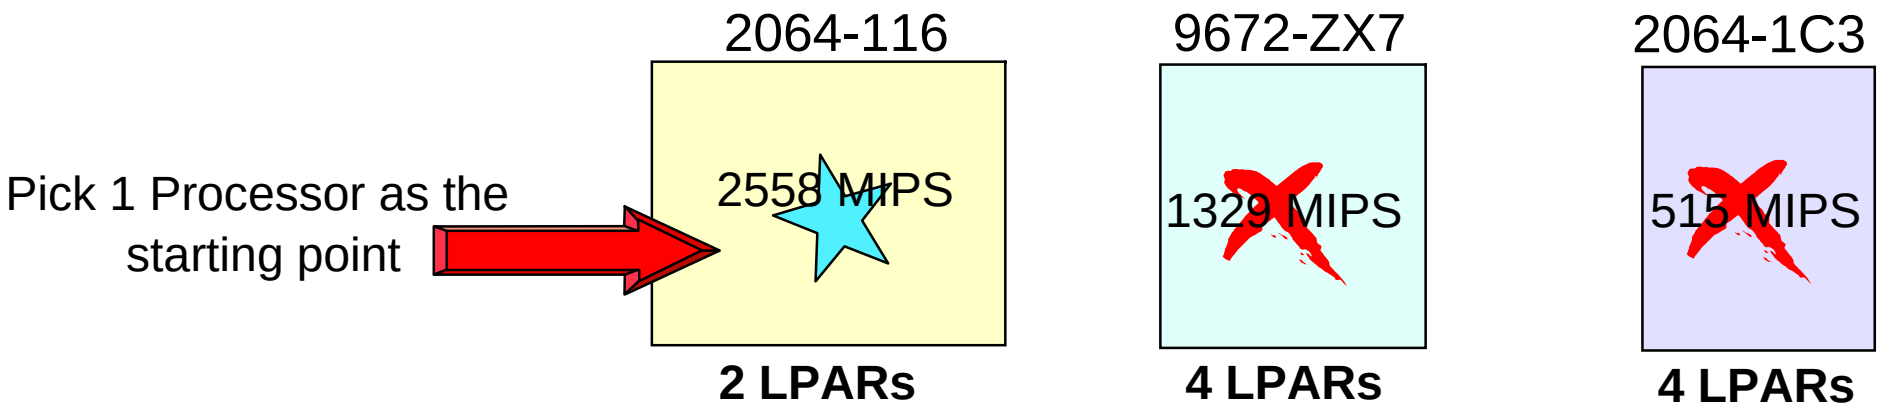

Custom Workload Mixes

- Do not use LSPR primitives to describe capacity relationships
- Build a custom mix for the major system images
- Use the Online to Other methodology
  - http://www-01.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS135

Workload	MIPS	ONLINE	OTHER
SYSTEM	190		190
TSO	30	30	
CICS	400	400	
BATCH	200		200
DDF	150	90	60
WAS	200	200	0
TOTAL	1170	720	450

Online = 62%
Other = 38%

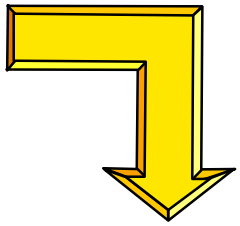

Create a Custom LSPR Mix

Tailor the LSPR Primitives

- Blend the LSPR primitives to match the workload

  Online = 62%
  Other = 38%

Name	CB-L	CB-S	WASDB	OLTP-W	OLTP-T
Production					
Online				31%	31%
Other	19%	19%			

- IBM has defined a set of "common" mixes

Name	CB-L	CB-S	WASDB	OLTP-W	OLTP-T
TI-MIX	30%	10%	0%	30%	30%
TD-MIX	45%	15%	0%	20%	20%
TM-MIX	52.5%	17.5%	0%	15%	15%
CB-MIX	75%	25%	0%	0%	0%
LOIO	60%	0%	20%	20%	0%
WEBMIX	27.3%	18.2%	18.2%	36.4%	0%
LSPR-MIX	20%	20%	20%	20%	20%

Low IO Workloads - Special Case

LPAR	SSCH RT	USED MSU	SSCH/MSU
SYS1	1100	75	15
SYS9	100	10	10
TOTAL	1200	85	14

USED MSU = (Processor MSU rating x utilization of the CEC or LPAR)

For z890s and z990s*
USED MSU = ((Processor MSU rating x utilization of the CEC or LPAR)/.90)

(MSU ratings for these processors have been set 10% below their actual capacities to provide for improved software price/performance)

- SSCH per Used MSU is <= 30
- Systems with a low I/O content should use proportionately higher amounts of WASDB and OLTP-W

Custom Workload Mixes

- If multiple LSPR tables are needed to characterize capacity, the five mixes can be used to assure consistency
  - The underlying LSPR workload primitives available in the various tables are different and cannot be used directly
  - Similar mixes of the same name are available in both the zSeries LSPR data and the legacy LSPR data and can be used
- New Web-Mix is only available in the zSeries LSPR Data
  - Cannot be used if processors prior to zSeries are part of the sizing equation
- LSPR-Mix is not intended to be useful for capacity planning
  - It is a simple average of the z/OS workloads which is used to generate the values used for software pricing purposes
  - The z/OS LSPR-Mix is not directly comparable to the OS/390 LSPR-Mix
    - The underlying LSPR workload primitives are not the same
- LSPR Information can be found at:

  http://www-1.ibm.com/servers/eserver/zseries/lspr/zSerieszOS.html

Generate Capacity Relationships of Current Processors

- Input to zPCR
  - Number of Partitions
  - Custom Mix
  - Number of logical CPs defined
  - Specialty CP (ICF, IFL, ZAAP)

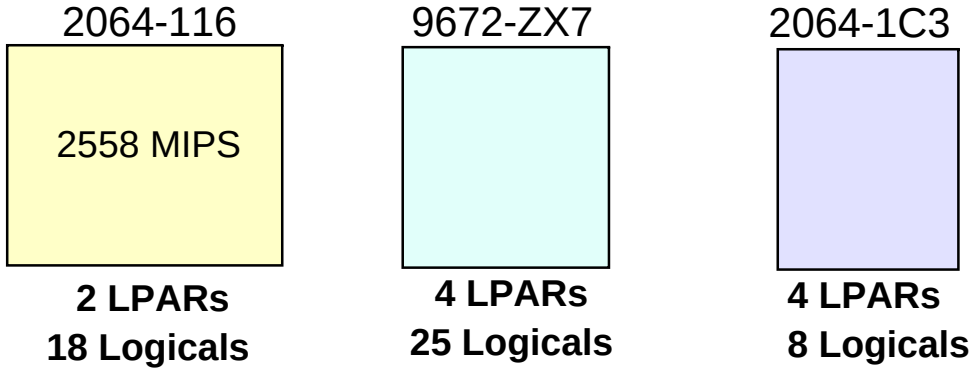

Processor	Relative Capacity**	MIPS	Comments
2064-116	1.0	2558	Base Processor
2064-116	0.9953	2546	2 LPARs, 18 Logicals
9672-ZX7	0.5003	1280	4 LPARs, 25 Logicals
2064-1C3	0.2099	537	4 LPARs, 8 Logicals

** Based on LOIO Mix

MIP Tables*
2558
2558
1433
651

* Based on LSPR Mix

- MIP tables may give misleading results
  - Often based on LSPR-Mix rather than custom mixes
  - Does not take into account the effects of LPAR

Generate Capacity Data

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

2064-116 - 2546 MIPs			Batch Window		Peak Online Hour	
LPAR	Weight	NLP	MIPS	MIPS@ 90%	MIPS	MIPS@ 90%
SYS1	930	16	1023	1137	2302	2558
SYS9	70	2	200	222	170	189
TOTAL	1000	18	1223	1359	2472	2747

Server Consolidation Considerations

- Document the end picture of the upgrade
  - LPAR layout
  - Special purpose CPs, i.e. ICF, IFL, zAAP
  - Number of books

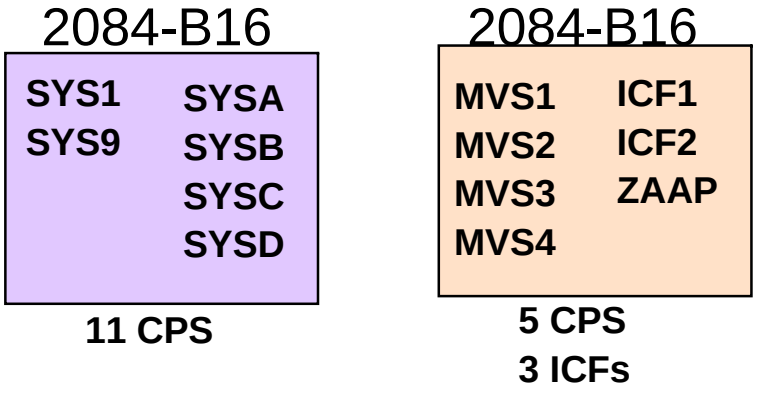

- Lay out the capacity data and determine the new LPAR weights and number of logical CPs
- Determine the number of logical CPs
  - White Paper shows how to get the logical CP requirements for an LPAR
    - http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/TD101238

2084-B16 - 11 CPs

	Batch Window				Peak Online Hour			
LPAR	MIPS	MIPS@ 90%	Weight	NLP	MIPS	MIPS@ 90%	Weight	NLP
SYS1	1023	1137	590	11	2302	2558	760	9
SYS9	200	222	110	1	170	189	60	1
SYSA	418	464	240	2	243	270	80	1
SYSB	30	33	20	1	161	179	50	1
SYSC	30	33	20	1	81	90	30	1
SYSD	30	33	20	1	54	60	20	1
TOTAL	1731	1922	1000	18	3011	3346	1000	14

- Need a processor which delivers 3346 MIPS
- Batch Window needs only 1922 which says the Batch Window is only 57% busy and the LPAR controls don't influence capacity
- Use peak online hour to set LPAR definitions

zPCR Capacity Data for 2084-311

zPCR (2.2a) - Detailed Partition Configuration Capacity Report
Combined: zSeries LSPR Data (10/29/2004) + Legacy LSPR Data (04/07/2004)
LPAR Host: 2084-B16 with 11 Traditional Engines, 0 IFLs, and 0 ICFs
Partitions: Total = 6; 6 Traditional, 0 IFL/ICF
Capacity values are relative to an IBM 2064-116 assumed at 2,558 MIPS (excludes LPAR cost)

	LP Identification				LP Definition					LP Capacity	
Inc	No	Pool	Name	SCP	Workload	Type	LCPs	Weight	Weight%	Capping	Minimum	Maximum
Y	1	Trad	SYS1	z/OS**	LoIO-Mix	SHR	9	760	76.0%		2,521.1	2,714.1
Y	2	Trad	SYS9	z/OS**	LoIO-Mix	SHR	1	60	6.0%		207.1	313.7
Y	3	Trad	SYSA	z/OS**	LoIO-Mix	SHR	1	80	8.0%		276.1	313.7
Y	4	Trad	SYSB	z/OS**	LoIO-Mix	SHR	1	50	5.0%		172.6	313.7
Y	5	Trad	SYSC	z/OS**	LoIO-Mix	SHR	1	30	3.0%		103.5	313.7
Y	6	Trad	SYSD	z/OS**	LoIO-Mix	SHR	1	20	2.0%		69.0	313.7

Totals for all partitions included in configuration	LP capacity
6 Traditional Partitions	3,349.3
0 IFL/ICF Partitions	0.0
6 Partitions Total - Combined Total	3,349.3

2084-B16 - 5 CPs + 3 ICF

	Batch Window				Peak Online Hour			
LPAR	MIPS	MIPS@ 90%	Weight	NLP	MIPS	MIPS@ 90%	Weight	NLP
MVS1	640	711	590	3	512	569	410	3
MVS2	128	142	120	1	320	356	250	2
MVS3	130	144	120	1	300	333	230	2
MVS4	185	206	170	1	128	142	110	1
ICF1				1				1
ICF2				1				1
ZAAP				1				1
ZAAP				1				1
TOTAL	1080	1203	1000	6	1260	1400	1000	8

zPCR Capacity Data for 2084-305, 3 ICFs

zPCR (2.2a) - Detailed Partition Configuration Capacity Report
Combined: zSeries LSPR Data (10/29/2004) + Legacy LSPR Data (04/07/2004)
LPAR Host: 2084-B16 with 5 Traditional Engines, 1 IFL, and 2 ICFs
Partitions: Total = 8; 4 Traditional, 4 IFL/ICF
Capacity values are relative to an IBM 2064-116 assumed at 2,558 MIPS (excludes LPAR cost)

	LP Identification				LP Definition					LP Capacity	
Inc	No	Pool	Name	SCP	Workload	Type	LCPs	Weight	Weight%	Capping	Minimum	Maximum
Y	1	Trad	MVS1-GP	z/OS**	LoIO-Mix	SHR	3	410	41.0%		669.2	979.3
Y	2	Trad	MVS2	z/OS**	LoIO-Mix	SHR	2	250	25.0%		410.3	656.4
Y	3	Trad	MVS3	z/OS**	LoIO-Mix	SHR	2	230	23.0%		377.4	656.4
Y	4	Trad	MVS4	z/OS**	LoIO-Mix	SHR	1	110	11.0%		182.3	331.5
Y	5	ICF	ICF1	CFCC	CFCC	DED	1	n/a			379.3	379.3
Y	6	ICF	ICF2	CFCC	CFCC	DED	1	n/a			379.3	379.3
Y	7	IFL	MVS1	Linux	WASDB/L	SHR	1	410	62.1%		183.9	296.0
Y	8	IFL	MVS2	Linux	WASDB/L	SHR	1	250	37.9%		112.1	296.0

Totals for all partitions included in configuration	LP capacity
4 Traditional Partitions	1,639.2
4 IFL/ICF Partitions	1,054.7
8 Partitions Total - Combined Total	2,693.9

Capacity Relationships for Proposed Processors

Processor	Relative Capacity**	MIPS	Comments	MIP Tables*
2064-116	1.0	2558	Base Processor	2558
2084-311	1.31	3349	6 LPARs, 14 Logicals	3540
2084-311	1.30	3323	6 LPARs, 19 Logicals	3540
2084-305	0.64	1639	8 LPARs, 12 Logicals	1801

** Based on LOIO Mix

* Based on LSPR Mix

- If the 2084-311 is changed to have more of the LPARs run as 2-ways the capacity change is approximately 26 MIPS
- The capacity expectation for the 2084-311 vs LSPR Default Mix is approx. 6%
- The capacity expectation for the 2084-305 vs LSPR Default Mix is approx. 10%

LPAR Impacts on Capacity

- **n-way and MP effects will impact capacity**
- **LPAR 3 is a uni, but the hardware is running as a 6-way shared processor and the capacity is of a 6-way shared processor**
  - **5 GCPs and 1 zAAP**
  - **z/OS 1.6 supports up to 24 CPs per image (32 CPs in 6/2005)**

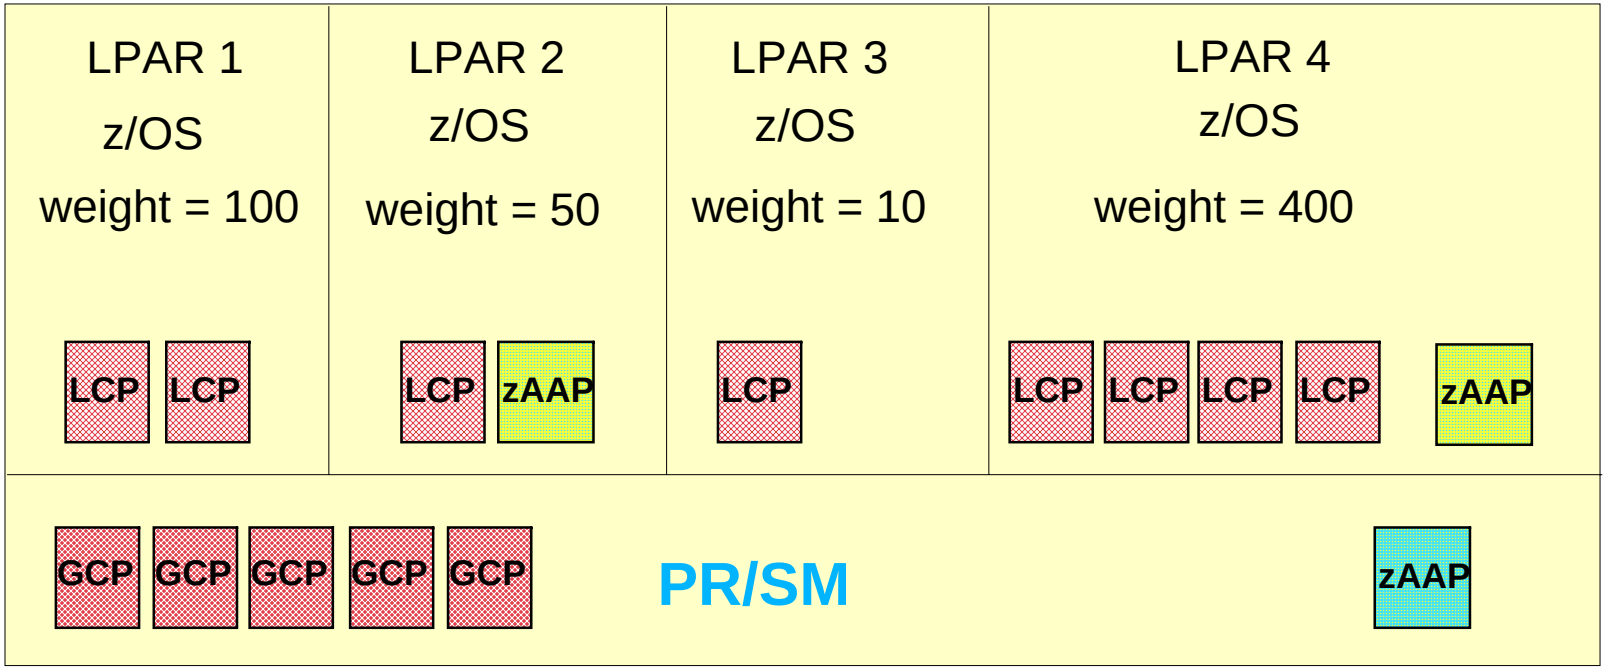

Capacity Planning and LPAR

- Set a 2084-316 as the base processor equated to 1.0
- 5752 MIPS, and defined with a LOIO mixed workload and shared CPs

Case	Mode	# of LPARs	LPARs x LCPs	LCP	ITRR	LCP:PCP
Base	2084-316	1	1 x 16	16	1.00	1:1
1	2084-316	2	2 x 12	24	.9882	1.5:1
2	2084-316	7	2 x 12 2 x 4 3 x 2	38	.9591	2.4:1
3	2084-316	4	2 x 12 2 x 6	36	.9733	2.25:1
4	2084-316	4	4 x 6	24	.9998	1.5:1

Special Purpose CPs

- ICF Impact
  - Within 10% of the performance of a stand-alone CF of the same processor family
  - Impact on SCP Engines
    - Addition of an ICF CP showed little to no degradation to the z/OS CPs
    - Impact of adding an ICF would be <= 1% per ICF
    - Example: Using a 2084-309, rated at 2932 MIPS, the addition of an ICF would cause the new configuration to be rated at 2903 MIPS. This is significantly closer to the non-ICF rating than having a 10-way configuration which is rated at 2875 MIPS
- IFL Impact
  - Depends upon the application activity level, impact may be slight (<10%) up to the full n-way impact of an additional CP
- zAAP Impact
  - Depends upon the amount of Java "eligible" workload, impact may be slight up to the full n-way impact of an additional CP
  - Impact is in terms of an additional partition and an additional CP to the z/OS environment

LPAR Utilization Cautions

- Lightly weighted LPARs may require additional capacity when being migrated to newer processors
- Explore the benefits of doing LPAR consolidation
  - Reduces the need to run z/OS as a uni-processor
    - Availability Issues
  - Virtual storage constraints need to be reviewed
  - Places greater emphasis on doing CICS consolidation to make fewer, larger CICS regions which can use more of a CP's capacity

Estimation Confidence

- Accuracy of the zPCR model 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
- Capacity decisions should be made with knowledge of the confidence factors

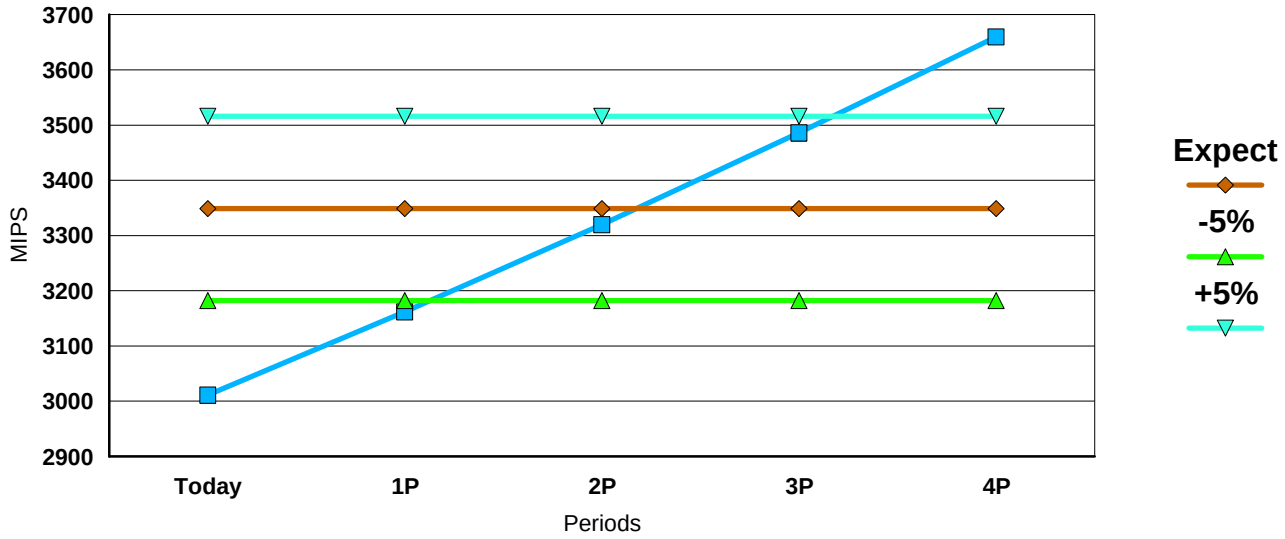

Post Install Analysis

- 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
  - Architecture changes (ESA/390 to z/Architecture)
  - Change in operating system levels
  - Maintenance
  - Change in processor storage (impacts sort-based workloads)
  - Buffer pool changes
  - Use of dynamic SQLs
  - Rebinding of SQL on new processor

Summary

- LPAR environments and their associated complexity have caused staight MIPS charts to become obsolete
- LSPR primitives need to be combined into custom workload mixes
- Understand the current system performance and latent demand indicators of an upgrade
- Use tools like zPCR / CP2000 to get the best view of expected capacity
- Set expectations with knowledge of confidence factors