#SHAREorg

SHARE

Understanding the RMF CPU Activity Report

Kathy Walsh
IBM

March 11, 2014
Session: 15285

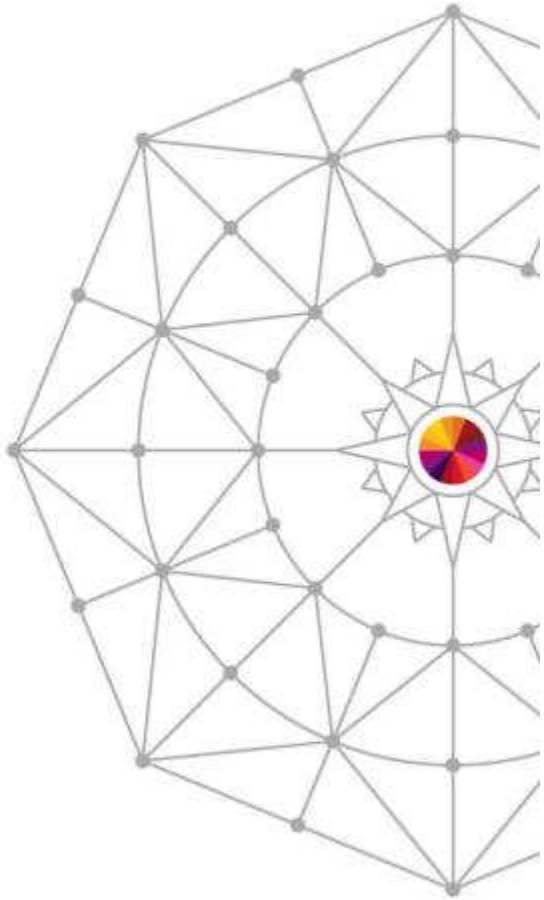

SHARE in Anaheim

Copyright (c) 2014 by SHARE Inc. Except where otherwise noted, this work is licensed under http://creativecommons.org/licenses/by-nc-sa/3.0/

Agenda

- Short RMF Overview
- How to Generate RMF Post-Processor Report
- Review Different CPU Activity Sections
  - Approach Used When Reviewing
  - Understanding Terms
  - How to tie the different metrics together

The RMF Examples used throughout this presentation are:

- From multiple systems and / or time frames
- Modified to make the information readable
- From RMF Reports before APARs were taken

© 2014 IBM Corporation

RMF Overview

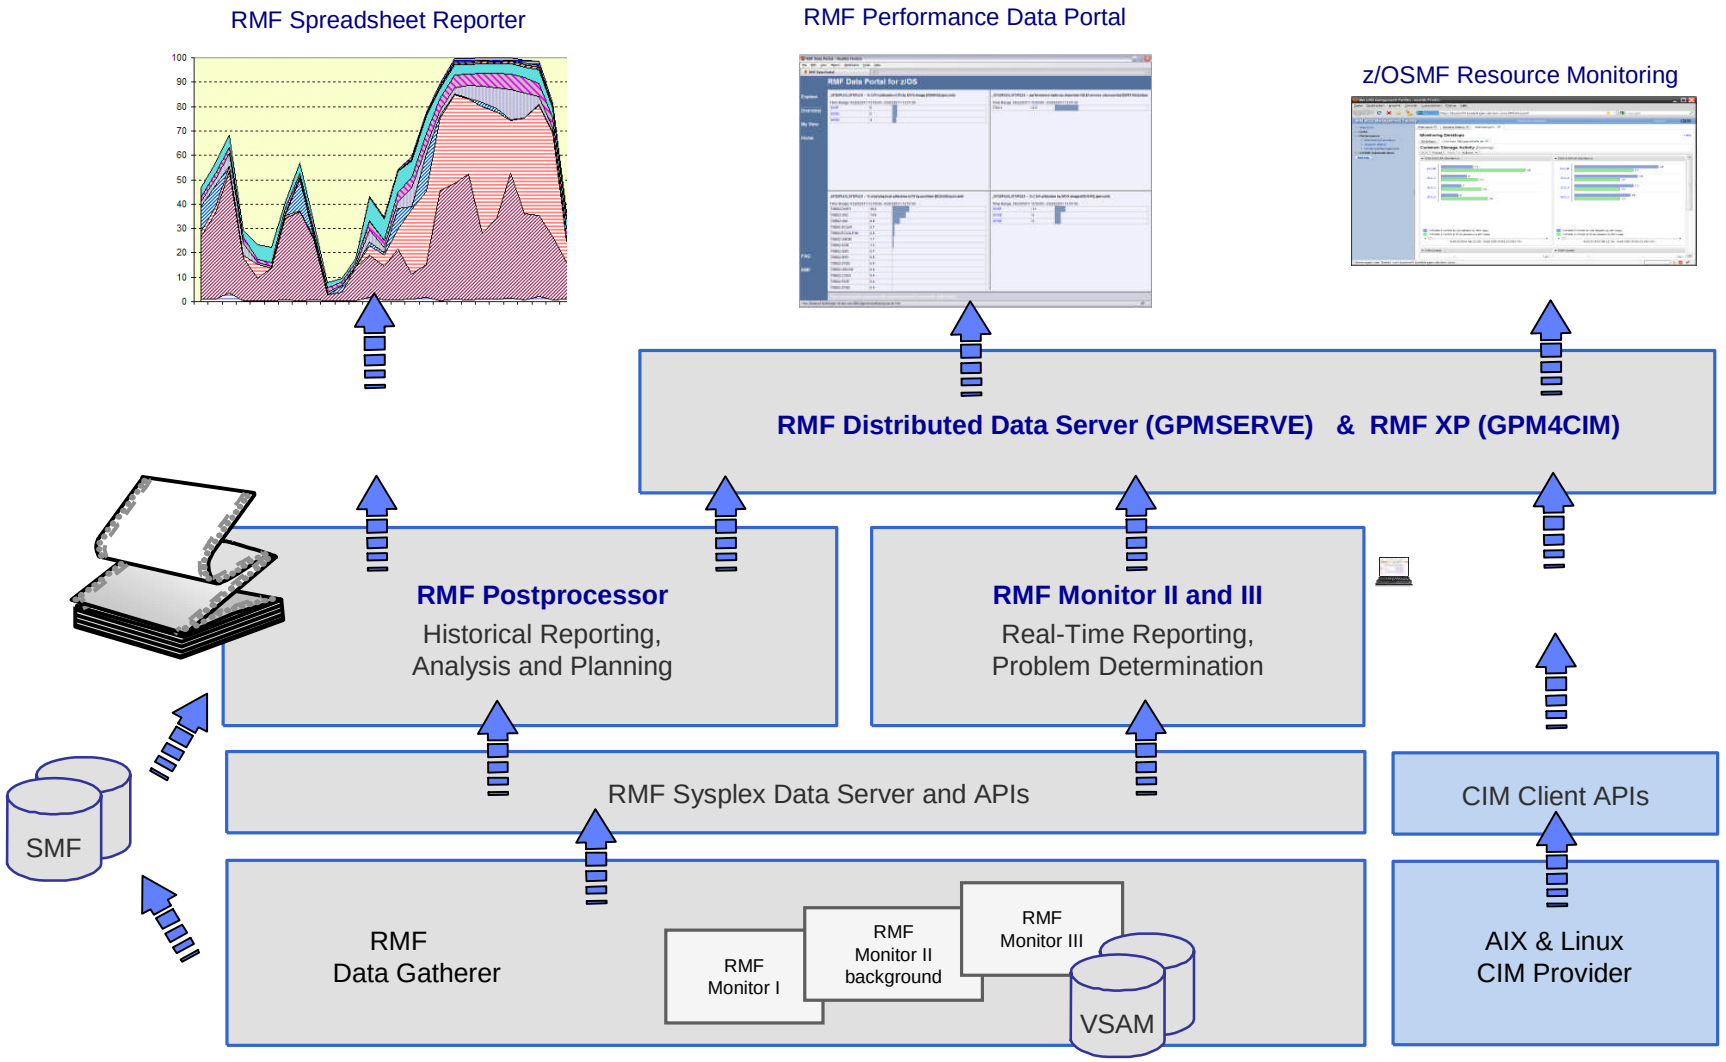

© 2014 IBM Corporation

RMF Data and SMF Records

- SMF Records are generated by RMF Monitor I, Monitor II, Monitor III
- Records can now be written out to either the MANxx data sets or, beginning with z/OS V1R9, use the MVS System Logger
- RMF Postprocessor reads the SMF type 70 records to generate the RMF CPU Activity Report

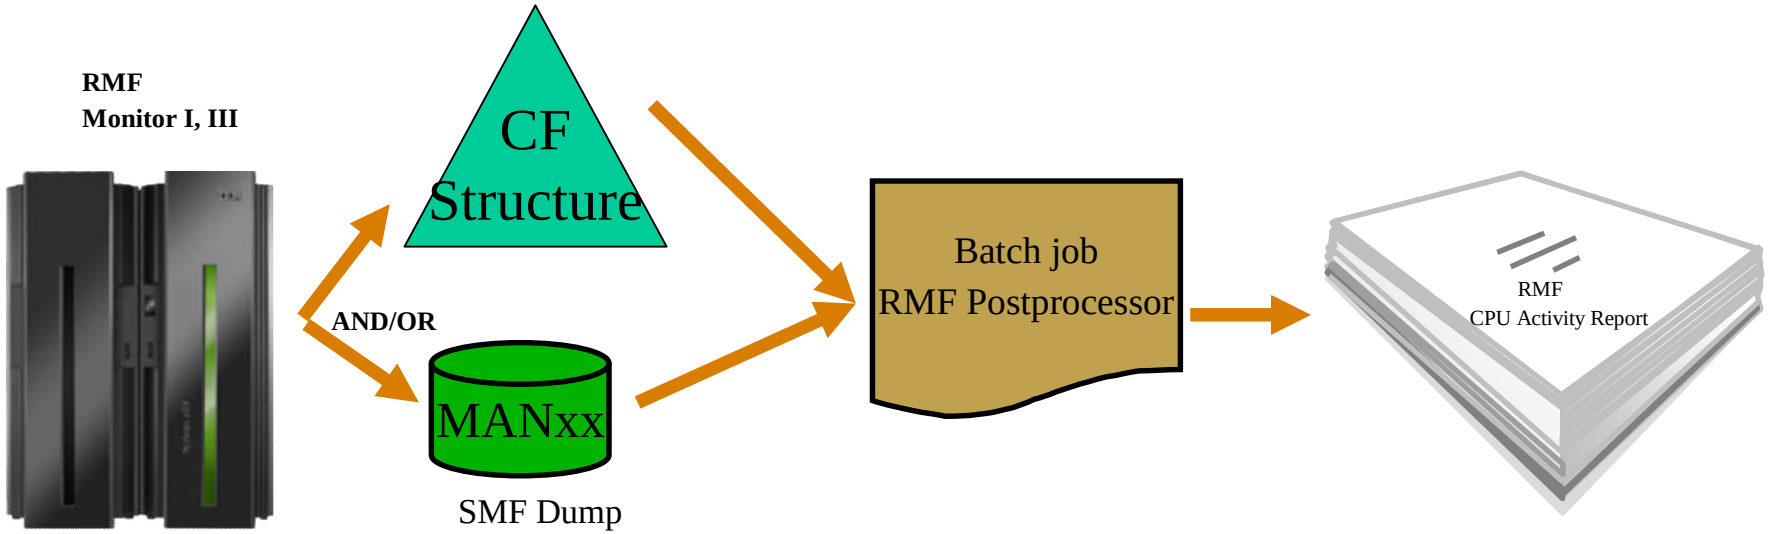

© 2014 IBM Corporation

Postprocessor XML Formatted Reports

Rationale:

- RMF Postprocessor reports are limited to a page width of 132 characters
- No state-of-the-art display capability of Postprocessor reports
- No easy access to RMF Postprocessor data for application programs
  - Cumbersome to parse the text output
  - Each report has its own layout

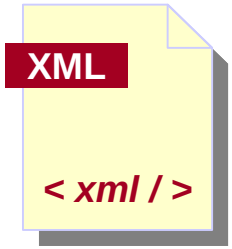

© 2014 IBM Corporation

Postprocessor XML Formatted Reports...

z/OS V1R11 RMF	z/OS V1R12 RMF	z/OS V1R13 RMF	z/OS V2R1 RMF
CPU Activity CRYPTO Activity FICON Director Activity ESS Disk Systems Activity OMVS Kernel Activity report OVERVIEW Report	DEVICE Activity WORKLOAD Activity	PAGING Activity SDELAY (XML only)	CACHE Subsystem Activity CF Activity CHANNEL Path Activity ENQUEUE Avtivity HFS Statistics IOQ Activity PAGESP Activity PCIE Activity (XML Only) SDEVICE Activity report VSTOR Activity XCF Activity

- Summary and Exception reports as well as interval reports based on data collected by a Monitor II background session are not available in XML format
- The XML format is the preferred RMF Postprocessor Report format for the future
- The XML Format supersedes the Text format. New Reports might not be implemented in Text format

© 2014 IBM Corporation

How to Generate a CPU Activity Report

Sample JCL

```
//RMFPP  EXEC PGM=ERBRMFPP
//MFPINPUT DD DISP=SHR,DSN=SMF.DATASET.SORTED
//SYSIN DD *
  SYSOUT(0)
  RTOD(0900,1700)                /*Report time of day*/
  DINTV(0060)                    /*Duration Interval*/
  DATE(02022009,02062009)  /*Days*/
  REPORT(CPU)
```

© 2014 IBM Corporation

CPU Report - Layout - 4 Parts

- CPU Activity
  - Detailed Individual CP data
  - Address Space Analysis
  - Blocked Workload Analysis
- Partition Data Report
  - Detailed LPAR information
- LPAR Cluster Report
  - Automatically created if clusters defined
  - Not available in a z/VM guest environment
- Group Capacity Report
  - Automatically created if group capacities are defined
  - Not available in a z/VM guest environment

© 2014 IBM Corporation

Part 1
CPU Activity

© 2014 IBM Corporation

What does this report tell me?

- What questions can I answer?
  - How many Logical CPs are defined?
  - Are Logical CPs being changed in the interval?
  - How busy is z/OS on each processor?
  - Is Hiperdispatch enabled?
    - How many Vertical Highs, Vertical Mediums and Vertical Lows are there
  - The I/O Interrupt rate and CPENABLE info
  - How much of each processor is the LPAR getting?

© 2014 IBM Corporation

Hiperdispatch Mode

- PR/SM
  - Supplies topology information/updates to the z/OS guest
  - Ties high priority logicals to physicals (gives 100% share)
  - Distributes remaining share to medium priority logicals
  - Distributes any additional service to unparked low priority logicals
- z/OS
  - Associates tasks with a small subsets of logical processors
  - Dispatches work to associated subset of logicals when possible
  - Dispatches work to some other CPU when necessary
  - Parks low priority processors if not needed or will not get service
- The combination provides the processor affinity that maximized the efficiency of the hardware caches

© 2014 IBM Corporation

CPU Activity Report Part 1 – Detailed CP Data

INTERVAL 14.59.302

CPU 2827 CPC CAPACITY 2755 SEQUENCE CODE 00000000000BBBBB
MODEL 722 CHANGE REASON=NONE HIPERDISPATCH=YES
H/W MODEL H89

---CPU---		TIME %			LOG PROC		--I/O INTERRUPTS--	
NUM	TYPE	ONLINE	LPAR BUSY	MVS BUSY	PARKED	SHARE %		RATE	% VIA TPI
0	CP	99.99	73.72	73.72	0.00	100.0	HIGH	332.0	49.79
1	CP	99.99	98.00	97.97	0.00	100.0	HIGH	405.5	45.92
2	CP	99.99	72.54	72.55	0.00	100.0	HIGH	5005	11.91
3	CP	99.99	77.80	79.50	0.00	94.5	MED	449.8	51.79
4	CP	99.99	21.94	41.65	43.99	0.0	LOW	0.00	0.00
5	CP	99.99	19.56	40.20	47.86	0.0	LOW	0.00	0.00
6	CP	99.99	15.86	40.68	57.22	0.0	LOW	0.00	0.00
7	CP	99.99	11.02	42.23	70.74	0.0	LOW	0.00	0.00
8	CP	99.99	25.29	44.79	40.40	0.0	LOW	0.00	0.00
TOTAL/AVERAGE			46.19	66.35		394.5		6192	19.06
9	IIP	99.99	70.42	70.36	0.00	100.0	HIGH		
A	IIP	99.99	62.36	62.32	0.00	83.3	MED		
B	IIP	99.99	24.28	47.48	48.90	0.0	LOW		
C	IIP	99.99	18.29	44.77	59.17	0.0	LOW		
D	IIP	99.99	7.42	43.20	82.83	0.0	LOW		
E	IIP	99.99	2.80	37.89	92.60	0.0	LOW		
F	IIP	99.99	1.16	37.09	96.88	0.0	LOW		
10	IIP	99.99	0.00	-----	100.00	0.0	LOW		
11	IIP	99.99	0.00	-----	100.00	0.0	LOW		
12	IIP	99.99	0.00	-----	100.00	0.0	LOW		
TOTAL/AVERAGE			18.67	58.38		183.3			

© 2014 IBM Corporation

I/O Total Interrupt Rate

- CPENABLE = (x,y)
  - A trade off of I/O responsiveness vs. throughput (ITR)
  - IEAOPTxx parameter to control the number of processors enabled for interrupts
  - High order CP is typically enabled, usually Vertical High if available
  - Any Logical CP in a wait is eligible to handle interrupts

- New support in z/OS 1.12 to show Interrupt Delay Time on RMF Device Activity Report

- Flash W9634A MVS CPENABLE Setting
  - Generally recommend CPENABLE=(10,30) for most environments
  - Use the flash for recommended settings for your environment
  - www.ibm.com/support/techdocs

© 2014 IBM Corporation

CPU Terminology - Busy Times

Dedicated Partition:

$$\text{MVS BUSY \%} = \frac{\text{INTERVAL TIME - WAIT TIME}}{\text{INTERVAL TIME}} * 100$$

Non-Dedicated Partition and WAIT Complete = NO (default)

$$\text{LPAR BUSY \%} = \frac{\text{Partition Dispatch Time}}{\text{INTERVAL TIME}} * 100$$

Non-Dedicated Partition and WAIT Complete = YES

$$\text{LPAR BUSY \%} = \frac{\text{Partition Dispatch Time - Wait Time}}{\text{INTERVAL TIME}} * 100$$

© 2014 IBM Corporation

MVS Busy MVS View of Processor Utilization

HD=NO

$$\text{MVS Busy Time} = \frac{\text{Online Time} - \text{Wait Time}}{\text{Online Time}} * 100$$

HD=YES

$$\text{MVS Busy Time} = \frac{\text{Online Time} - (\text{Wait Time} + \text{Parked Time})}{(\text{Online Time} - \text{Parked Time})} * 100$$

© 2014 IBM Corporation

CPU to Dispatch Ratio

Interval / Ready Work	CP 0	CP 1	CP 2	CP 3
1 CICS,STC,Batch,Batch	CICS L=P	BATCH L=P	STC L=P	BATCH L=P
2 CICS,STC,Batch	CICS L	BATCH L=P	STC L=P	0
3 CICS,Batch,Batch,Batch	CICS L=P	BATCH L=P	BATCH L=P	BATCH L=P
4 CICS	CICS L	0	0	0

```
CICS Active          4:4  = 100%
CICS Dispatched      2:4  =  50%

LPAR BUSY           10:16 =  63%
MVS BUSY            12:16 =  75%
```

© 2014 IBM Corporation

Evaluate LPAR Busy vs MVS Busy

- LPAR Busy ~= MVS Busy
  - MVS is voluntarily giving back the physical CP to the LPAR hipervisor
    - MVS completed its work
    - MVS in a wait state
- LPAR Busy < MVS Busy
  - LPAR hipervisor is taking the physical CP away from the logical CP before work is completed
  - z/OS with the Warning Track support on a zEC12 / zBC12 is now aware when the physical CP is taken away from the logical CP
    - On any prior server z/OS is not aware the Physical CP is taken away
  - Indicator of Latent Demand
  - A function of weights and number of logical CP's assigned
- LPAR Busy > MVS Busy
  - LPAR is using the physical CP to support CF Link operations
  - With z196 GA2 this activity is reported in *PHYSCAL partition

© 2014 IBM Corporation

Parked Engines and Logical Processor Share %

- Parked time is percentage of time parked
  - Time not dispatched to z/OS, and does not attempt to run work. This field is not applicable without Hiperdispatch enabled
- Logical Proc Share %
  - Percentage of the physical processor to which the logical processor is entitled
  - HiperDispatch=YES
    - Can be up to 100%, or a low of 0
  - HiperDispatch=NO
    - Processing weight is divided equally between the online processors

© 2014 IBM Corporation

Latent Demand: LPAR Busy vs MVS Busy

CPU 2097 CPC CAPACITY 1451
MODEL 719 CHANGE REASON=N/A **HIPERDISPATCH=YES**

---CPU---		TIME %			LOG PROC	
NUM	TYPE	ONLINE	LPAR BUSY	MVS BUSY	**PARKED**	SHARE %	
0	CP	100.00	96.77	96.80	0.00	100.0	HIGH
1	CP	100.00	94.91	94.95	0.00	100.0	HIGH
2	CP	100.00	96.72	96.74	0.00	100.0	HIGH
3	CP	100.00	95.07	95.10	0.00	100.0	HIGH
4	CP	100.00	**50.18**	**93.55**	0.00	66.0	MED
5	CP	100.00	50.15	93.56	0.00	66.0	MED
6	CP	100.00	**20.30**	**89.09**	**56.00**	0.0	LOW
7	CP	100.00	11.40	90.19	72.00	0.0	LOW
8	CP	100.00	22.12	88.49	50.79	0.0	LOW
9	CP	100.00	**46.12**	**87.87**	**0.00**	0.0	LOW
A	CP	100.00	45.37	86.74	0.00	0.0	LOW
B	CP	100.00	38.46	86.76	11.21	0.0	LOW
C	CP	100.00	35.08	86.96	19.43	0.0	LOW
D	CP	100.00	19.29	84.13	57.66	0.0	LOW
E	CP	100.00	0.00	-----	100.00	0.0	LOW
F	CP	100.00	0.00	-----	100.00	0.0	LOW
10	CP	100.00	0.00	-----	100.00	0.0	LOW
TOTAL/AVERAGE			**42.47**	**91.45**		**532.0**	

CEC Busy = 98.85

.0115 * 19 CP = .22 CPs available

Weight: 5.32 CPs

Using: 42.47/100 * 17 LCP = 7.22 CPs

© 2014 IBM Corporation

Understanding the Numbers

CEC is 98.85% busy

			LCP		LOG PROC	POLARITY	UNPARKED	LPAR	UNPARKED
	ONLINE	LPAR BUSY	MVS BUSY	PARKED	SHARE %		CPs	MVS BUSY	EFF
1	100	96.77	96.80	0	100	HIGH	100	96.80	96.77
2	100	94.91	94.95	0	100	HIGH	100	94.95	94.91
3	100	96.72	96.74	0	100	HIGH	100	96.74	96.72
4	100	95.07	95.10	0	100	HIGH	100	95.10	95.07
5	100	50.18	93.55	0	66	MED	100	93.55	50.18
6	100	50.15	93.56	0	66	MED	100	93.56	50.15
7	100	20.30	89.09	56.00	0	LOW	44.00	39.20	46.14
8	100	11.40	90.19	72.00	0	LOW	28.00	25.25	40.71
9	100	22.12	88.49	50.79	0	LOW	49.21	43.55	44.95
10	100	46.12	87.87	0	0	LOW	100	87.87	46.12
11	100	45.37	86.74	0	0	LOW	100	86.74	45.37
12	100	38.46	86.76	11.21	0	LOW	88.79	77.03	43.32
13	100	35.08	86.96	19.43	0	LOW	80.57	70.06	43.54
14	100	19.29	84.13	57.66	0	LOW	42.34	35.62	45.56
15	100	0.00	0.00	100	0	LOW	0		
16	100	0.00	0.00	100	0	LOW	0		
17	100	0.00	0.00	100	0	LOW	0		
	17	42.47		5.67	5.32		11.33	1036.03	
		7.22	74.76					91.45	

Unparked Effective – Percent of time dispatched by LPAR when not parked

$$\frac{\text{(LPAR Busy / 100)}}{\text{100 – Parked Time / 100}}$$

© 2014 IBM Corporation

Understanding the Numbers – Next Interval

CEC is 97.85% busy

			LCP		LOG PROC	POLARITY	UNPARKED	LPAR	UNPARKED
	ONLINE	LPAR BUSY	MVS BUSY	PARKED	SHARE %		CPs	MVS BUSY	EFF
1	100	96.15	96.15	0	100	HIGH	100	96.15	96.15
2	100	93.72	93.75	0	100	HIGH	100	93.75	93.72
3	100	96.03	96.02	0	100	HIGH	100	96.02	96.03
4	100	94.06	94.06	0	100	HIGH	100	94.06	94.06
5	100	60.87	93.23	0	66	MED	100	93.23	60.87
6	100	60.85	93.19	0	66	MED	100	93.19	60.85
7	100	32.59	88.61	40.88	0	LOW	59.12	52.39	55.13
8	100	5.16	84.41	90.54	0	LOW	9.46	7.99	54.55
9	100	0	0	100	0	LOW	0	0	
10	100	36.29	88.05	32	0	LOW	68	59.96	53.29
11	100	40.58	86.46	23	0	LOW	77	66.19	53.01
12	100	54.02	84.87	0	0	LOW	100	84.87	54.02
13	100	53.13	83.78	0	0	LOW	100	83.78	53.13
14	100	53.83	84.59	0	0	LOW	100	84.59	53.83
15	100	0	0	100	0	LOW	0		
16	100	0	0	100	0	LOW	0		
17	100	0	0	100	0	LOW	0		
	17	45.72		5.87	5.32		11.13	1006.16	
		7.77						90.38	

MVS Busy: $\frac{\text{Online Time - (Wait Time + Parked Time)}}{\text{Online Time - Parked Time}}$

LPAR MVS Busy: $\frac{\text{Unparked Time * (LCP MVS BUSY)}}{\Sigma\ \text{(Unparked CPs)}}$

© 2014 IBM Corporation

Hiperdispatch Summary

- Important to ensure LPAR weights are close to actual LPAR usage
  - Drives better allocation of Vertical Highs
- Still be realistic in number of logical CPs assigned to an LPAR
  - e.g. if using 7.7 LCPs at max specify 9-11 not 17
  - LPAR MVS Busy is key metric driving unparking
- Calculate Unparked Effectiveness and evaluate workload delays
  - Impacts are very workload specific
  - Check CPU to Dispatch ratios
- Latent Demand indicators now need to include knowledge of:
  - Parked CPs over time
  - Unparked Effectiveness
- Watch LPAR weights for small LPARs with low utilization
  - Weight = 1.98 CPs then 1 VH, 1 VM (2 LCPs)
  - Change Weight to: 2.01 then 1 VH, 2 VM (3 LCPs)

© 2014 IBM Corporation

Part 2
Address Space Analysis

© 2014 IBM Corporation

System Address Space Analysis

- **IN**
  - In central storage, includes in ready count
- **IN READY**
  - Ready to execute or currently in execution
- **OUT READY**
  - Physically swapped out of memory, ready to execute
  - Tuning Issues
  - Delay in work processing
  - Want this number to approach zero
- **OUT WAIT**
  - Physically swapped out of memory, not ready to execute
- **LOGICAL Out RDY**
  - Logically swapped out of memory, ready to execute
  - Tuning issue - possible memory problem
- **LOGICAL OUT WAIT**
  - Logically swapped out, not ready to execute

© 2014 IBM Corporation

System Address Space Analysis

SYSTEM ADDRESS SPACE AND WORK UNIT ANALYSIS

---------NUMBER OF ADDRESS SPACES---------

QUEUE TYPES	MIN	MAX	AVG
IN	716	787	732.3
IN READY	1	224	15.9
OUT READY	0	7	0.1
OUT WAIT	0	0	0.0
LOGICAL OUT RDY	0	10	0.1
LOGICAL OUT WAIT	305	375	360.0

ADDRESS SPACE TYPES

ADDRESS SPACE TYPES	MIN	MAX	AVG
BATCH	8	18	11.0
STC	810	901	889.6
TSO	75	82	78.7
ASCH	0	0	0.0
OMVS	103	193	113.2

---------NUMBER OF WORK UNITS-------------

CPU TYPES	MIN	MAX	AVG
CP	0	2,399	19.5
IIP	0	435	5.1

---------------------------DISTRIBUTION OF IN-READY WORK U

NUMBER OF WORK UNITS	(%)	0 10 20 30 40 50 |....|....|....|....|....|....
<= N	59.4	>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>
= N + 1	4.0	>>>
= N + 2	3.7	>>
= N + 3	2.8	>>
<= N + 5	4.7	>>>
<= N + 10	6.6	>>>>
<= N + 15	4.3	>>>
<= N + 20	2.8	>>
<= N + 30	3.7	>>
<= N + 40	1.8	>
<= N + 60	1.8	>
<= N + 80	0.8	>
<= N + 100	0.5	>
<= N + 120	0.3	>
<= N + 150	0.4	>
> N + 150	1.5	>

Doesn't differentiate by CP type

N = NUMBER OF PROCESSORS ONLINE UNPARKED (9.6 ON AVG)

© 2014 IBM Corporation

Part 3
Blocked Workload Analysis

© 2014 IBM Corporation

Blocked Workload Support

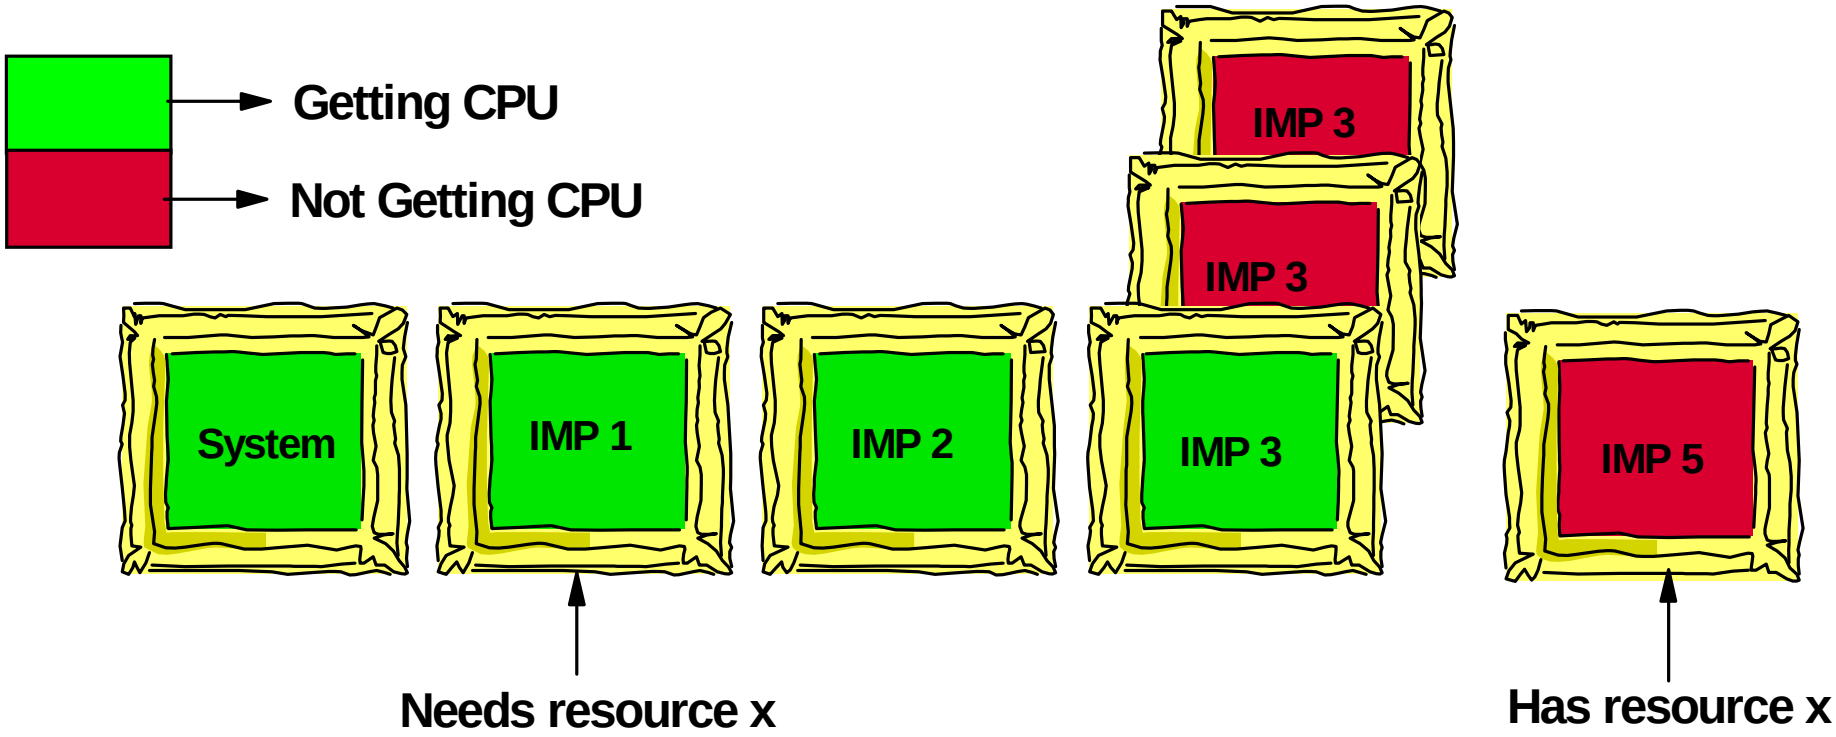

High Priority work is now blocked by lower priority work

© 2014 IBM Corporation

Blocked Workload Support

```
BLOCKED WORKLOAD ANALYSIS
  OPT PARAMETERS: BLWLTRPCT (%)    0.5     PROMOTE RATE:  DEFINED       85     WAITERS FOR PROMOTE:  AVG      0.518
                  BLWLINTHD         20                    USED (%)      0                           PEAK      23
```

Keyword / Field	Description	Value	Default
BLWLTRPCT	•Specifies how much CPU capacity may be used to promote blocked work •Influences how many address spaces can be promoted at the same time	0 - 200	5% or 0.5%
BLWLINTHD	Specifies the threshold time interval, in seconds, a swapped-in address space or enclave must wait before being considered blocked	5 – 65535 (18+ hours)	20 seconds
Promote Rate Defined	Value derived from the BLWLTRPCT parameter. It is the number of dispatchable work units which may get promoted per second.		
Promote Rate Used	The utilization of the defined promote rate during the interval		
Waiters for Promote	Shows the Average and Peak number found blocked during the reporting interval		

© 2014 IBM Corporation

Part 4
Partition Data Report

© 2014 IBM Corporation

What do these reports tell me

- What questions can I answer?
  - Processor configuration by CP type
  - Actual physical utilization of the processor
  - Number of and busy of the partitions
    - Identify dominant partitions - who is using the processor
  - Logical CP's definitions
    - Are the weights aligned to the number of logical CP's
  - Is the LPAR capped
  - Is IRD being used
  - Are Capacity Groups being used
  - Impacts of PR/SM on capacity

© 2014 IBM Corporation

Partition Data Report

```
MVS PARTITION NAME                    WSC1       NUMBER OF PHYSICAL PROCESSORS        44              GROUP NAME       N/A
IMAGE CAPACITY                        1753                    CP                     22              LIMIT            N/A
NUMBER OF CONFIGURED PARTITIONS          7                    IIP                    22              AVAILABLE        N/A
WAIT COMPLETION                         NO
DISPATCH INTERVAL                  DYNAMIC
```

--------- PARTITION DATA ----------							-- LOGICAL PARTITION PROCESSOR DATA --				-- AVERAGE PROCESSOR UTILIZATION PERCENTAGES --				
			----MSU----		-CAPPING--		PROCESSOR-		----DISPATCH TIME DATA----		LOGICAL PROCESSORS		--- PHYSICAL PROCESSORS ---		
NAME	S	WGT	DEF	ACT	DEF	WLM%	NUM	TYPE	EFFECTIVE	TOTAL	EFFECTIVE	TOTAL	LPAR MGMT	EFFECTIVE	TOTAI
WSC1	A	327	0	521	NO	0.0	9.0	CP	01.01.46.138	01.02.18.334	45.79	46.19	0.16	18.73	18.90
WSC2	A	158	0	162	NO	0.0	5.0	CP	00.19.06.106	00.19.20.306	25.49	25.81	0.07	5.79	5.86
WSC3	A	474	0	922	NO	0.0	15.0	CP	01.49.30.249	01.50.21.341	48.71	49.09	0.26	33.21	33.47
WSC4	A	84	0	11	NO	0.0	3.0	CP	00.01.14.697	00.01.17.152	2.77	2.86	0.01	0.38	0.39
WSC5	A	173	0	99	NO	0.0	3.0	CP	00.11.39.491	00.11.47.926	25.93	26.24	0.04	3.54	3.58
WSC6	A	97	0	12	NO	0.0	3.0	CP	00.01.24.656	00.01.27.361	3.14	3.24	0.01	0.43	0.44
WSC7	A	510	0	143	NO	0.0	6.0	CP	00.16.52.600	00.17.10.442	18.77	19.10	0.09	5.12	5.21
PHYSICAL		1823								00.01.34.475			0.48		0.48
									---	---			---	---	---
TOTAL									03.41.33.940	03.45.17.340			1.13	67.19	68.32
WSC1	A	100			NO		10	IIP	00.27.55.074	00.27.59.165	18.63	18.67	0.02	8.47	8.49
WSC2	A	200			NO		2	IIP	00.01.14.519	00.01.15.979	4.14	4.22	0.01	0.38	0.38
WSC3	A	300			NO		15	IIP	01.00.20.522	01.00.32.002	26.84	26.93	0.06	18.30	18.36
WSC5	A	200			NO		1	IIP	00.00.04.683	00.00.04.937	0.52	0.55	0.00	0.02	0.02
WSC6	A	200			NO		1	IIP	00.00.03.072	00.00.03.315	0.34	0.37	0.00	0.02	0.02
WSC7	A	200			NO		1	IIP	00.00.17.676	00.00.18.319	1.97	2.04	0.00	0.09	0.09
PHYSICAL		1200								00.00.18.310			0.09		0.09
									---	---			---	---	---
TOTAL									01.29.55.548	01.30.32.029			0.18	27.27	27.46

© 2014 IBM Corporation

Partition Data Report

```
MVS PARTITION NAME                    WSC1      NUMBER OF PHYSICAL PROCESSORS    44
IMAGE CAPACITY                        1753                CP                     22
NUMBER OF CONFIGURED PARTITIONS          7                IIP                    22
WAIT COMPLETION                         NO
DISPATCH INTERVAL                    DYNAMIC
```

- Image capacity – MSUs – Calculated as the minimum of the following:
  - Capacity based on the partition's logical CP configuration
    - Sum of online and offline Logical CPs
    - Issue D M=CPU command to see maximum number of logical CPs
  - Defined capacity limit of the partition, if available (image softcap)
  - Capacity limit of the related Group Capacity if the partition belongs to a capacity group

Model	MSU	MSU PER CP	WSC1	LCPs	WSC1
2827-722	2755	2755/22=125.2	9*125.2 = 1127	9 online 5 offline	14*125.2=1753

MSU is always based on the MSU of the overall CEC, i.e. a 722 not a 709

© 2014 IBM Corporation

LPAR Busy Calculation

INTERVAL 14.59.302
CYCLE 2.000 SECONDS

MVS PARTITION NAME	WSC1	NUMBER OF PHYSICAL PROCESSORS		44
IMAGE CAPACITY	1753		CP	22
NUMBER OF CONFIGURED PARTITIONS	10		IIP	22
WAIT COMPLETION	NO			
DISPATCH INTERVAL	DYNAMIC			

	-- LOGICAL PARTITION PROCESSOR DATA --				-- AVERAGE PROCESSOR UTILIZATION PERCENTAGES --				
	PROCESSOR-		----DISPATCH TIME DATA----		LOGICAL PROCESSORS		--- PHYSICAL PROCESSORS ---		
NAME	NUM	TYPE	EFFECTIVE	TOTAL	EFFECTIVE	TOTAL	LPAR MGMT	EFFECTIVE	TOTAL
WSC1	9.0	CP	01.01.46.138	01.02.18.334	45.79	46.19	0.16	18.73	18.90
WSC2	5.0	CP	00.19.06.106	00.19.20.306	25.49	25.81	0.07	5.79	5.86
WSC3	5.0	CP	01.49.30.249	01.50.21.341	48.71	49.09	0.26	33.21	33.47
WSC4	3.0	CP	00.01.14.697	00.01.17.152	2.77	2.86	0.01	0.38	0.39
WSC5	3.0	CP	00.11.39.491	00.11.47.926	25.93	26.24	0.04	3.54	3.58
WSC6	3.0	CP	00.01.24.656	00.01.27.361	3.14	3.24	0.01	0.43	0.44
WSC7	6.0	CP	00.16.52.600	00.17.10.442	18.77	19.10	0.09	5.12	5.21
				00.01.34.475			0.48		0.48
			03.41.33.940	03.45.17.340			1.13	67.19	68.32

LPAR Busy = Partition dispatch time/Online time *100
Use WSC1 General CPs as example:

Partition Dispatch time = 3738 seconds (1.02.18.334)
Online time = 8091 seconds (14.59.302 = 899 seconds * 9 LCPs)
LPAR Busy = (3738/8091) * 100 = 46.19%
Physical Bus = (3738/19778) * 100 = 18.90%
22 CPs * 899 second interval = 19778 total seconds

© 2014 IBM Corporation

Calculate processor(s) guaranteed to each LPAR

1. $$\textit{LPAR Fair Share} = \frac{\textit{LPAR Weight}}{\textit{Sum of Weights (509)}}$$

2. ***# of General Purpose Physical CPs * LPAR share % = # Processors guaranteed to the partition***

- WSC1 capacity = 22 * .18 = 3.96 CPs
- WSC2 capacity = 22 * .09 = 1.98 CPs
- WSC3 capacity = 22 * .26 = 5.72 CPs
-WSC4 capacity = 22 * .05 = 1.1 CPs
-WSC5 capacity = 22 * .10 = 2.2 CPs
-WSC6 capacity = 22 * .05 = 1.1 CPs
-WSC7 capacity = 22 * .28 = 6.16 CPs

© 2014 IBM Corporation

Let’s Put Some of it Together

© 2014 IBM Corporation

System Address Space Analysis

SYSTEM ADDRESS SPACE AND WORK UNIT ANALYSIS

---------NUMBER OF ADDRESS SPACES---------			
QUEUE TYPES	MIN	MAX	AVG
IN	716	787	732.3
IN READY	1	224	15.9
OUT READY	0	7	0.1
OUT WAIT	0	0	0.0
LOGICAL OUT RDY	0	10	0.1
LOGICAL OUT WAIT	305	375	360.0
ADDRESS SPACE TYPES			
BATCH	8	18	11.0
STC	810	901	889.6
TSO	75	82	78.7
ASCH	0	0	0.0
OMVS	103	193	113.2

---------NUMBER OF WORK UNITS----------			
CPU TYPES	MIN	MAX	AVG
CP	0	2,399	19.5
IIP	0	435	5.1

----------	----------	---------DISTRIBUTION OF IN-READY WORK U
NUMBER OF WORK UNITS	(%)	0 10 20 30 40 50 |....|....|....|....|....|....
<= N	59.4	>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>
= N + 1	4.0	>>>
= N + 2	3.7	>>
= N + 3	2.8	>>
<= N + 5	4.7	>>>
<= N + 10	6.6	>>>>
<= N + 15	4.3	>>>
<= N + 20	2.8	>>
<= N + 30	3.7	>>
<= N + 40	1.8	>
<= N + 60	1.8	>
<= N + 80	0.8	>
<= N + 100	0.5	>
<= N + 120	0.3	>
<= N + 150	0.4	>
> N + 150	1.5	>

Doesn't differentiate by CP type

N = NUMBER OF PROCESSORS ONLINE UNPARKED (9.6 ON AVG)

© 2014 IBM Corporation

Back to the Logical CP View

---CPU---		---------- TIME % ----------			LOG PROC		--I/O INTERRUPTS--	
NUM	TYPE	ONLINE	LPAR BUSY	MVS BUSY	PARKED	SHARE %		RATE	% VIA TPI
0	CP	99.99	73.72	73.72	0.00	100.0	HIGH	332.0	49.79
1	CP	99.99	98.00	97.97	0.00	100.0	HIGH	405.5	45.92
2	CP	99.99	72.54	72.55	0.00	100.0	HIGH	5005	11.91
3	CP	99.99	77.80	79.50	0.00	94.5	MED	449.8	51.79
4	CP	99.99	21.94	41.65	43.99	0.0	LOW	0.00	0.00
5	CP	99.99	19.56	40.20	47.86	0.0	LOW	0.00	0.00
6	CP	99.99	15.86	40.68	57.22	0.0	LOW	0.00	0.00
7	CP	99.99	11.02	42.23	70.74	0.0	LOW	0.00	0.00
8	CP	99.99	25.29	44.79	40.40	0.0	LOW	0.00	0.00
TOTAL/AVERAGE			46.19	66.35		394.5		6192	19.06
9	IIP	99.99	70.42	70.36	0.00	100.0	HIGH		
A	IIP	99.99	62.36	62.32	0.00	83.3	MED		
B	IIP	99.99	24.28	47.48	48.90	0.0	LOW		
C	IIP	99.99	18.29	44.77	59.17	0.0	LOW		
D	IIP	99.99	7.42	43.20	82.83	0.0	LOW		
E	IIP	99.99	2.80	37.89	92.60	0.0	LOW		
F	IIP	99.99	1.16	37.09	96.88	0.0	LOW		
10	IIP	99.99	0.00	-----	100.00	0.0	LOW		
11	IIP	99.99	0.00	-----	100.00	0.0	LOW		
12	IIP	99.99	0.00	-----	100.00	0.0	LOW		
TOTAL/AVERAGE			18.67	58.38		183.3			

© 2014 IBM Corporation

LPAR Weights are Important

```
MVS PARTITION NAME                    WSC1        NUMBER OF PHYSICAL PROCESSORS        44
IMAGE CAPACITY                        1753                CP                           22
NUMBER OF CONFIGURED PARTITIONS          7                IIP                          22
WAIT COMPLETION                         NO
DISPATCH INTERVAL                    DYNAMIC

--------- PARTITION DATA ----------------  -- LOGICAL PARTITION PROCESSOR DATA --   -- AVERAGE PROCESSOR
                    ----MSU----  -CAPPING--  PROCESSOR-  ----DISPATCH TIME DATA----   LOGICAL PROCESSORS
NAME        S   WGT  DEF    ACT  DEF   WLM%  NUM   TYPE   EFFECTIVE      TOTAL       EFFECTIVE    TOTAL
WSC1        A   327    0    521  NO     0.0  9.0   CP     01.01.46.138  01.02.18.334      45.79     46.19
WSC2        A   158    0    162  NO     0.0  5.0   CP     00.19.06.106  00.19.20.306      25.49     25.81
WSC3        A   474    0    922  NO     0.0 15.0   CP     01.49.30.249  01.50.21.341      48.71     49.09
WSC4        A    84    0     11  NO     0.0  3.0   CP     00.01.14.697  00.01.17.152       2.77      2.86
WSC5        A   173    0     99  NO     0.0  3.0   CP     00.11.39.491  00.11.47.926      25.93     26.24
WSC6        A    97    0     12  NO     0.0  3.0   CP     00.01.24.656  00.01.27.361       3.14      3.24
WSC7        A   510    0    143  NO     0.0  6.0   CP     00.16.52.600  00.17.10.442      18.77     19.10
*PHYSICAL*                                               00.01.34.475
                                                         ------------  ------------
  TOTAL                                                  03.41.33.940  03.45.17.340

WSC1        A   100                 NO          10   IIP   00.27.55.074  00.27.59.165      18.63     18.67
WSC2        A   200                 NO           2   IIP   00.01.14.519  00.01.15.979       4.14      4.22
WSC3        A   300                 NO          15   IIP   01.00.20.522  01.00.32.002      26.84     26.93
WSC5        A   200                 NO           1   IIP   00.00.04.683  00.00.04.937       0.52      0.55
WSC6        A   200                 NO           1   IIP   00.00.03.072  00.00.03.315       0.34      0.37
WSC7        A   200                 NO           1   IIP   00.00.17.676  00.00.18.319       1.97      2.04
*PHYSICAL*      1200                             30        00.00.18.310
                                                         ------------  ------------
  TOTAL                                                  01.29.55.548  01.30.32.029
```

8.3% * 22 CPs = 1.83 CPs

© 2014 IBM Corporation

LPAR Cluster Report

© 2014 IBM Corporation

LPAR Cluster Report

- LPAR Cluster - An LPAR Cluster is a set of 1 or more logical partitions resident on the same physical server and in the same sysplex running z/OS

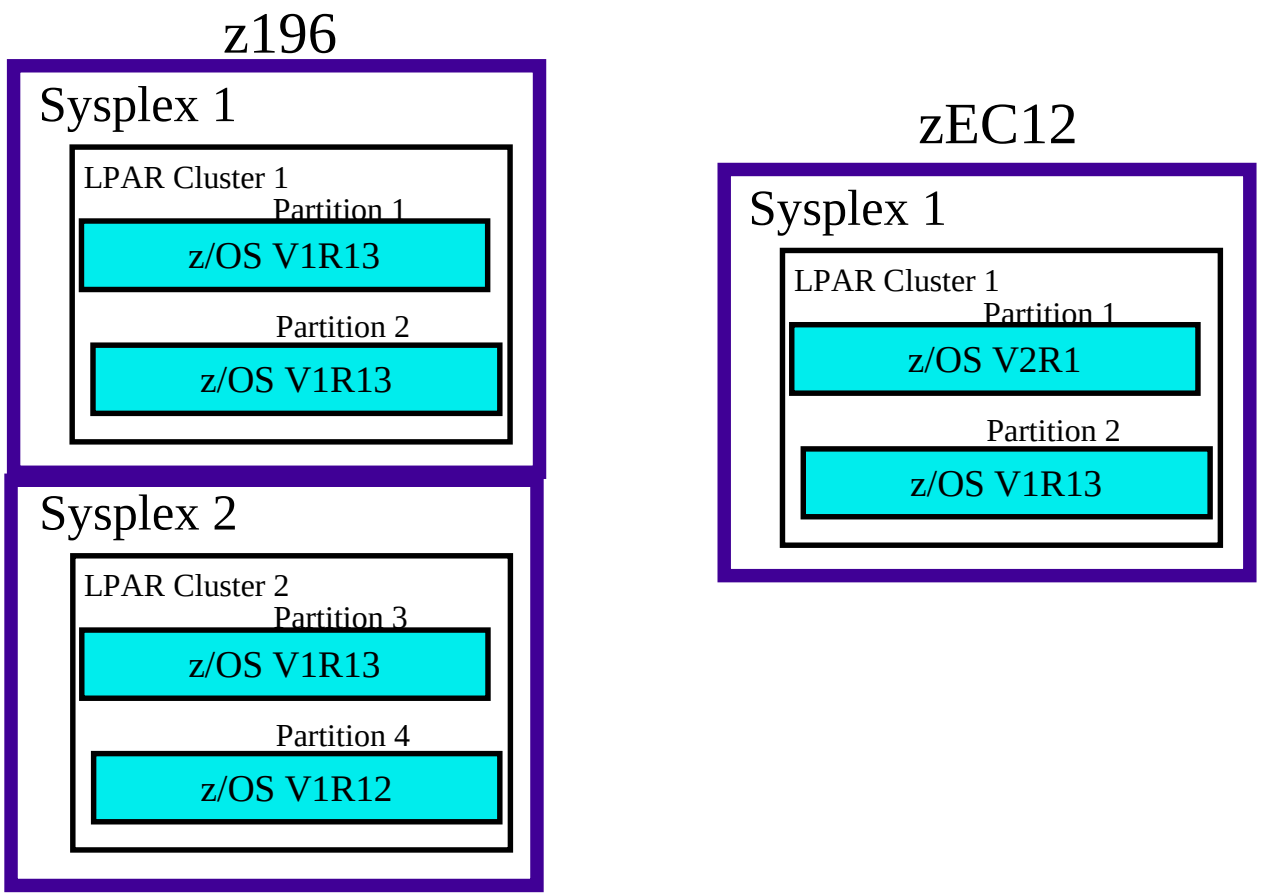

© 2014 IBM Corporation

LPAR Cluster Report

```
                                                L P A R  C L U S T E R  R E P O R T
  z/OS V1R12    SYSTEM ID WXYZ                   DATE 02/14/2014    INTERVAL 14.59.982
                CONVERTED TO z/OS V1R13 RMF      TIME 11.30.00      CYCLE 1.000 SECONDS

                              ------ WEIGHTING STATISTICS ------   ---- PROCESSOR STATISTICS ----
                              --- DEFINED ---  ---- ACTUAL -----   ---- NUMBER ---    -- TOTAL% --
CLUSTER  PARTITION SYSTEM  INIT  MIN  MAX  AVG  MIN %  MAX %   DEFINED  ACTUAL   LBUSY  PBUSY
SYSPLEXG AB1       ABCD     400  300  680  680    0.0   92.0        23    16.0   86.23  68.98
         YZ1       WXYZ     580  300  680  300   90.5    0.0        19     8.0   76.06  30.42
-----------------  ----------  ----------  ----------------  ------------------------------
                   TOTAL     980                                    42           162.3  99.41
```

- Watch the difference between Defined INIT and Actual AVG
  - Indicates how much weight movement is occurring
- Actual Min% and Max% tells how often the LPAR weight was within 10% of the min or max definition
- Sum of LPAR weights in a cluster will always equal the TOTAL value from the initial weights. IRD will not steal from one cluster to give capacity to another

© 2014 IBM Corporation

Group Capacity Report

© 2014 IBM Corporation

Group Capacity

- Manage CPU for a group of z/OS LPARs on a single CEC
  - Limit is set to total usage by all LPARs in group
    - Members which don't want their share will donate to the other members
  - Independent of sysplex scope and IRD LPAR cluster
  - Works with defined capacity limits on an LPAR
    - Target share will not exceed defined capacity
  - Works with WLM LPAR CPU management (IRD)
  - Can have more than one group on a CEC but an LPAR may only be a member of one group
  - LPARs must share engines and specify WAIT COMPLETION = NO
- Capacity groups are defined on the HMC Change LPAR Group Controls panels
  - Specify group name, limit in MSUs, and LPARs in the group
  - Members can be added or removed dynamically

© 2014 IBM Corporation

RMF Group Capacity Enhancement

MVS PARTITION NAME	WSC9
IMAGE CAPACITY	729
NUMBER OF CONFIGURED PARTITIONS	15
WAIT COMPLETION	NO
DISPATCH INTERVAL	DYNAMIC

GROUP NAME	WSCGRP1
LIMIT	257
AVAILABLE	100

G R O U P C A P A C I T Y R E P O R T

GROUP-CAPACITY NAME	LIMIT	PARTITION	SYSTEM	-- MSU -- DEF	ACT	WGT	---- CAPPING ---- DEF	WLM%	ACT%	- ENTITLEMENT - MINIMUM	MAXIMUM
WSCGRP1	257	WSC9	WSC9	0	42	87	NO	50.0	50.0	102	257
		WSCC	WSCC	0	115	132	NO	45.8	24.1	154	257
			TOTAL		157	219					

© 2014 IBM Corporation

Group Capacity Report

G R O U P C A P A C I T Y R E P O R T

GROUP-CAPACITY NAME	LIMIT	PARTITION	SYSTEM	-- MSU -- DEF	ACT	WGT	---- CAPPING ---- DEF	WLM%	ACT%	- ENTITLEMENT - MINIMUM	MAXIMUM
WSCGRP1	257	WSC9	WSC9	0	42	87	NO	50.0	50.0	102	257
		WSCC	WSCC	0	115	132	NO	45.8	24.1	154	257
			TOTAL		157	219					

- Capping WLM%
  - Percentage of time when WLM considers to cap the partition
- Capping ACT%
  - Percentage of time when capping actually limited the usage of processor resources for the partition
- Entitled minimum is calculated from LPAR weights of LPARs in the group
  - Percent share of WSC9 is its weight (87) divided by total weight (219) or (87/219) = .397
  - Group limit is 257, so WSC9 minimum entitlement is .397 * 257 = 102
- Entitled maximum is lower of either group capacity, or individual LPAR defined capacity

© 2014 IBM Corporation

Additional Information

- ★RMF Report Analysis - additional field descriptions, SC33-7991
- ★RMF Users Guide, SC33-7990
- ★RMF Performance Management Guide, SC33-7992
- ★PR/SM Planning Guide, SB10-7033 (online)

© 2014 IBM Corporation

Questions

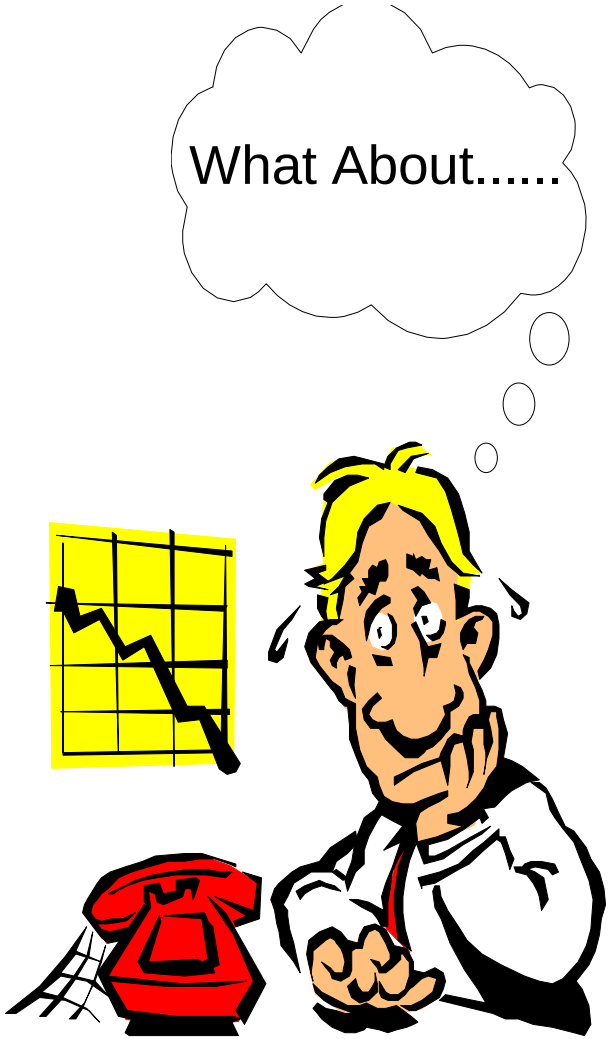

© 2014 IBM Corporation

Trademarks

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

Not all common law marks used by IBM are listed on this page. Failure of a mark to appear does not mean that IBM does not use the mark nor does it mean that the product is not actively marketed or is not significant within its relevant market.

Those trademarks followed by ® are registered trademarks of IBM in the United States; all others are trademarks or common law marks of IBM in the United States.

For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml:

*, AS/400®, e business(logo)®, DBE, ESCO, eServer, FICON, IBM®, IBM (logo)®, iSeries®, MVS, OS/390®, pSeries®, RS/6000®, S/30, VM/ESA®, VSE/ESA, WebSphere®, xSeries®, z/OS®, zSeries®, z/VM®, System i, System i5, System p, System p5, System x, System z, System z9®, BladeCenter®

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.

© 2014 IBM Corporation

Notice Regarding Specialty Engines (e.g., zIIPs, zAAPs and IFLs):

Any information contained in this document regarding Specialty Engines ("SEs") and SE eligible workloads 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). 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 (“AUT”).

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.

© 2014 IBM Corporation