


z/OS Performance “HOT” Topics

Kathy Walsh
IBM

August 4, 2014
Session Number: **15841**



#SHAREorg



1

Agenda

- Performance and Capacity Planning Topics
 - New Performance White Papers
 - Instructions Counts
 - CPU MF and HIS Support
 - zPCR Latest Status
 - zBNA Latest Status
 - zEDC Compression
 - Mobile Workload Pricing
 - New RSM and WLM APARS
 - New SMF 42 Support
 - Websphere WOLA and WLM
- ▶ Addendum
 - Older APARs or Performance Information

www.ibm.com/support/techdocs

Techdocs - the Technical Sales Library

This site provides access to the Technical Sales Support organization's technical information databases. It gives you access to the most current installation, planning and technical support information available from IBM pre-sales support, and is constantly updated. You can browse or search these databases by date, document number, product, platform, keywords, etc.

New to Techdocs? Take a look at our [detailed introduction](#), which describes the document categories available (those listed on the navigation area on the left side of this page).

Rather than browse these categories, as a convenience you may enter a search of the full **Techdocs** database, or of any category you wish, here:

Search: Allow word variants
 for: "Fuzzy" search
 Hits: Order by:
 Include docs updated: [Help for Search](#)

Also available: our [Advanced search](#), where you can select documents based on various assigned document attributes.

White Paper: Checklist for Establishing Capacity Groups

- **WP102437**

- Describes the HMC procedure for setting up a Group Capacity profile and assigning LPARs to become members of the group
- Gives step by step directions including all of the screen shots needed to complete the task including updating the Running System
- Shows how to validate via RMF the Capacity Group definitions are active

White Paper: Managing zOS on a Uni-processor - V2

- **WP102437**

- Discusses managing work on a uniprocessor system including the impact of looping or CPU-intensive tasks
- Availability features such as blocked workload and ENQ promotion

Benefits of Exploiting Large Memory for DB2 Buffer Pools

- Three papers will be in the series, two available now:
 1. Advantages of Configuring More Memory for DB2 Buffer Pools
 - **WP102464**
 - <http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP102464>
 2. Performance Report on Exploiting Large Memory for DB2 Buffer Pools with SAP®
 - **WP102461**
 - <http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP102461>
 3. Benefits of very large buffers from a DB2 perspective (target date: 4Q2014)
- Concept:
 - DB2 buffer pools and associated memory were scaled at generally high (>80%) CPU utilization. Buffer pools were grown by as little as 14GB to 100's of GB
 - Performance benefits are highly workload dependent, many clients should see measureable, single digit percentage reduced CPU time and double digit percentage better response time
 - Large CPU time reductions are due mainly to avoiding hardware cache disruptions, which are directly tied to the task switching associated with I/O operations. Thus, eliminating synchronous database IO through configuration of larger buffer pools can have a direct influence on CPU time used.

z/OS 2.1 - Instruction Counts

- Provide a more consistent metric which does not see as much CPU variability due to:
 - Impacts of hardware caching
 - LPAR configurations
 - Software stack
 - Workload interactions – dispatch rate
- New support provides Instruction Counts in SMF 30 records
 - Supported on z10 processors and later which supports CPUMF
 - Similar to the SMF 30 CPU time fields
 - SMFPRMxx new keyword `SMF30COUNT|NOSMF30COUNT`
 - Requires CPUMF basic and extended counters in HIS be active
- Early use of the metric is showing more variability due to impact of interrupts on recording

15836: Why is the CPU Time so Variable, : Tuesday, 3:00 PM

15847: System z Processor Consumption Analysis, Tuesday, 1:30 PM

Need z/OS 2.1 Instruction Count “Volunteers”

- We want to determine Instruction Count vs CPU time consistency at various utilizations

Looking for “Volunteers”

- SMF data from 3 days from 2 different “like” weeks (e.g. Mon-Wed)
–24 hours/day, SMF 30s, 70s, 72s, 113s per LPAR
- z/OS 2.1 with “SMF30COUNT” enabled and CPU MF Counters enabled
- Production partitions only
- Preferred Customer Profile: (any one type of workload)
 - Websphere, CICS running with Transaction goals, DDF workload, or Batch jobs that run repeatedly thru day (Utility job)

If interested send note to jpburg@us.ibm.com,

No deliverable will be returned

Benefit: Opportunity to ensure your data is used to influence analysis

CPU Measurement Facility

- Available on all System z processors since the z10
- Facility provides hardware instrumentation data for production systems
- Two Major components
 - Counters
 - Cache and memory hierarchy information
 - SCPs supported include z/OS and z/VM
 - Sampling
 - Instruction time-in-CSECT
- New z/OS HIS started task
 - Gathered on an LPAR basis
 - Writes SMF 113 records
- New z/VM Monitor Records
 - Gathered on an LPAR basis – all guests are aggregated
 - Writes new Domain 5 (Processor) Record 13 (CPU MF Counters) records
- Minimal overhead

zPCR Latest Status

■ Version 8.5c (6/16/2014)

- For all EC processor models, algorithms concerning **GP+zAAP** or **GP+zIIP** partitions have been improved. Switching overhead is now based on the number of books occupied by GP/zAAP/zIIP CPs rather than the total number of books installed
- For these partition combinations on multi-book processor configurations, capacity results may improve slightly depending on how the GP/zAAP/zIIP CPs are actually distributed across the configured books

■ Version 8.5b (5/30/2014)

- New **LSPR Table Control** window is now displayed next to the **LSPR Table** window
- Provides the ability to specify what processors will appear in the **Multi-Image** and **Single-Image LSPR Table** windows
- Up to 20 specific processor models (GP, IFL, or both) can be added to a **Favorites** list.

■ Version 8.5 (1/31/2014)

- LSPR data continues to be based on z/OS 1.13
- Tool updated for Java7
 - Java7 and all tools installed by a user must be updated at the same time
 - Impacts zPCR, zBNA, and zSoftCap
- Removes requirement for Windows “Administrator Authority”
 - Still need Admin Authority to remove old version

Example of Using zPCR Favorites Support

LSPR Capacity Ratio Table

Workload Graph Help

z/OS-1.13 LSPR Data (07/23/2013)

LSPR Multi-Image Capacity Ratios
General Purpose CPs
 Values are applicable for z/OS; representative of z/VM and Linux
 Capacity basis: 2094-701 @ 559.792 MIPS for a typical multi-partition configuration
 Capacity for z/OS on z10 and later processors is represented with HiperDispatch turned ON

Processor	Features	Flag	MSU	LSPR Workload Category				
				Low	Low-Avg	Average	Avg-High	High
zEnterprise 196/700								
2817-701	1W	=	150	1,195	1,199	1,202	1,176	1,151
2817-702	2W	=	281	2,325	2,298	2,272	2,192	2,117
2817-703	3W	=	408	3,431	3,370	3,311	3,178	3,055
2817-704	4W	=	531	4,513	4,415	4,320	4,134	3,964
2817-705	5W	=	650	5,576	5,435	5,300	5,062	4,845
2817-706	6W	=	766	6,620	6,430	6,251	5,962	5,699
2817-707	7W	=	879	7,645	7,403	7,175	6,835	6,526
2817-708	8W	=	988	8,652	8,352	8,072	7,682	7,328
2817-709	9W	=	1,091	9,640	9,278	8,943	8,503	8,105
2817-710	10W	=	1,191	10,611	10,183	9,788	9,300	8,858
2817-711	11W	=	1,286	11,565	11,066	10,609	10,072	9,587
2817-712	12W	=	1,381	12,501	11,929	11,407	10,822	10,294
2817-713	13W	=	1,473	13,420	12,770	12,181	11,549	10,979
2817-714	14W	=	1,562	14,323	13,592	12,932	12,254	11,643
2817-715	15W	=	1,648	15,209	14,394	13,662	12,937	12,286
2817-716	16W	=	1,731	16,080	15,177	14,371	13,601	12,909
2817-717	17W	=	1,816	16,945	15,956	15,076	14,260	13,528
2817-718	18W	=	1,899	17,805	16,730	15,778	14,916	14,144
2817-719	19W	=	1,983	18,660	17,500	16,476	15,569	14,756
2817-720	20W	=	2,064	19,510	18,266	17,171	16,217	15,365
2817-721	21W	=	2,144	20,355	19,027	17,862	16,863	15,969
2817-722	22W	=	2,224	21,194	19,784	18,550	17,504	16,570
2817-723	23W	=	2,306	22,029	20,537	19,234	18,142	17,168
2817-724	24W	=	2,388	22,859	21,285	19,915	18,777	17,762

Processor models in table = 1,064; In this view = 953; Currently selected = 1

Provisional Reference-CPU | Workload Categories | Copy Selected to Favorites | Table Controls

Provisional Reference-CPU is active
 Select multiple processors with **Ctrl+LeftClick** or **Shift+LeftClick**; For flag explanation, position mouse on indicator

LSPR Table Control

Settings Help

Processors Displayed

- All Families
- Selected Families
- Favorites

Selected Families

- z9-BC
- z10-BC
- z114
- zBC12
- z9-EC
- z10-EC
- z196
- zEC12

Favorites

1	2817-708
---	----------

Move Selected Processor

⬆ ⬇ ⬅ ⬇

zPCR Favorites Support

Workload Graph Help

zPCR V8.5c

z/OS-1.13 LSPR Data (07/23/2013)

LSPR Multi-Image Capacity Ratios
Favorite CPs

Values are applicable for z/OS; representative of z/VM and Linux
Capacity basis: 2094-701 @ 559.792 MIPS for a typical multi-partition configuration
Capacity for z/OS on z10 and later processors is represented with HiperDispatch turned ON

Processor	Features	Flag	MSU	LSPR Workload Category				
				Low	Low-Avg	Average	Avg-High	High
2827-7xx I10	10W IFL	=		13,552	12,829	12,179	11,460	10,821
2827-7xx I11	11W IFL	=		14,751	13,926	13,188	12,399	11,699
2827-7xx I12	12W IFL	=		15,924	14,994	14,166	13,308	12,548
2827-7xx I13	13W IFL	=		17,085	16,046	15,126	14,200	13,381
2827-506	6W	=	409	3,571	3,441	3,320	3,142	2,982
2827-507	7W	=	470	4,129	3,968	3,820	3,613	3,427
2827-508	8W	=	530	4,679	4,486	4,308	4,073	3,863
2827-509	9W	=	588	5,221	4,993	4,784	4,524	4,290

Processor models in table = 1,064; In this view = 8; Currently selected = 8

Provisional Reference-CPU Workload Categories Copy Selected to Favorites Table Controls

Global Reference-CPU is active; double click any processor row to set it as a Provisional Reference-CPU
Select multiple processors with **Ctrl+LeftClick** or **Shft+LeftClick**; For flag explanation, position mouse on indicator

Settings Help

Processors Displayed

All Families

Selected Families

Favorites

Selected Families

z9-BC z9-EC

z10-BC z10-EC

z114 z196

zBC12 zEC12

Favorites

1	2827-7xx I10
2	2827-7xx I11
3	2827-7xx I12
4	2827-7xx I13
5	2827-506
6	2827-507
7	2827-508
8	2827-509

Move Selected Processor

⏪ ⏩ ⏴ ⏵

■ **System z Capacity Planning Opportunities:**

- Per thread (engine) speed improvements for CMOS CPs is slowing dramatically
 - Every CMOS platform is facing this issue
 - Future capacity gains will be by adding more CPs rather than much faster CPs
 - Enhances need for parallel operation and more reliance on parallel sysplex

- Availability of subcapacity models continues to grow
 - Provide capacity as more, slower processors increasing parallelism
 - Especially useful in environment with large number of LPARs
 - Additional capacity can be acquired in smaller increments
 - Receive benefit since Specialty CPs run at full n-way speed

**Impact of these trends will most likely be seen first in the
Batch Window**

IBM System z Batch Network Analyzer (zBNA)

- IBM System z Batch Network Analyzer
 - A free, “as is” tool to analyze batch windows
 - Available to Customers, Business Partners and IBMers
 - PC based, providing graphical and text reports
 - Includes Gantt charts and support for Alternate Processors (what if scenarios)



Releases

- zBNA V1.3
 - Support for SMF 42 records to understand the “Life of a Data Set”
- zBNA V1.4.2
 - Support for SMF 14/15 records to understand zEDC compression candidates
- zBNA V1.4.5
 - Support to allow creation of CSV files from most zBNA panels

Available Now on Techdocs

<https://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS5132>

Why use zBNA?

- Perform "what if" analysis and estimate CPU upgrade effect on batch window
- Identify job time sequences based on a graphical view
- Filter jobs by attributes like CPU time / intensity, job class, service class, etc.
- Review the resource consumption of all the batch jobs
- Drill down to the individual steps to see the resource usage
- Identify candidate jobs for running on different processors
- Identify jobs with speed of engine concerns (top tasks %)
- Identify by job which datasets are used and the datasets performance info
- For any dataset identify every job in the time window using the data set
- Identify top zEDC compression candidates and understand the zEDC card capacity required

15207: zBNA Tool – Because Batch is Back, Wednesday, 10:00 AM

15837: zBNA User Experience, Wednesday, 4:15 PM

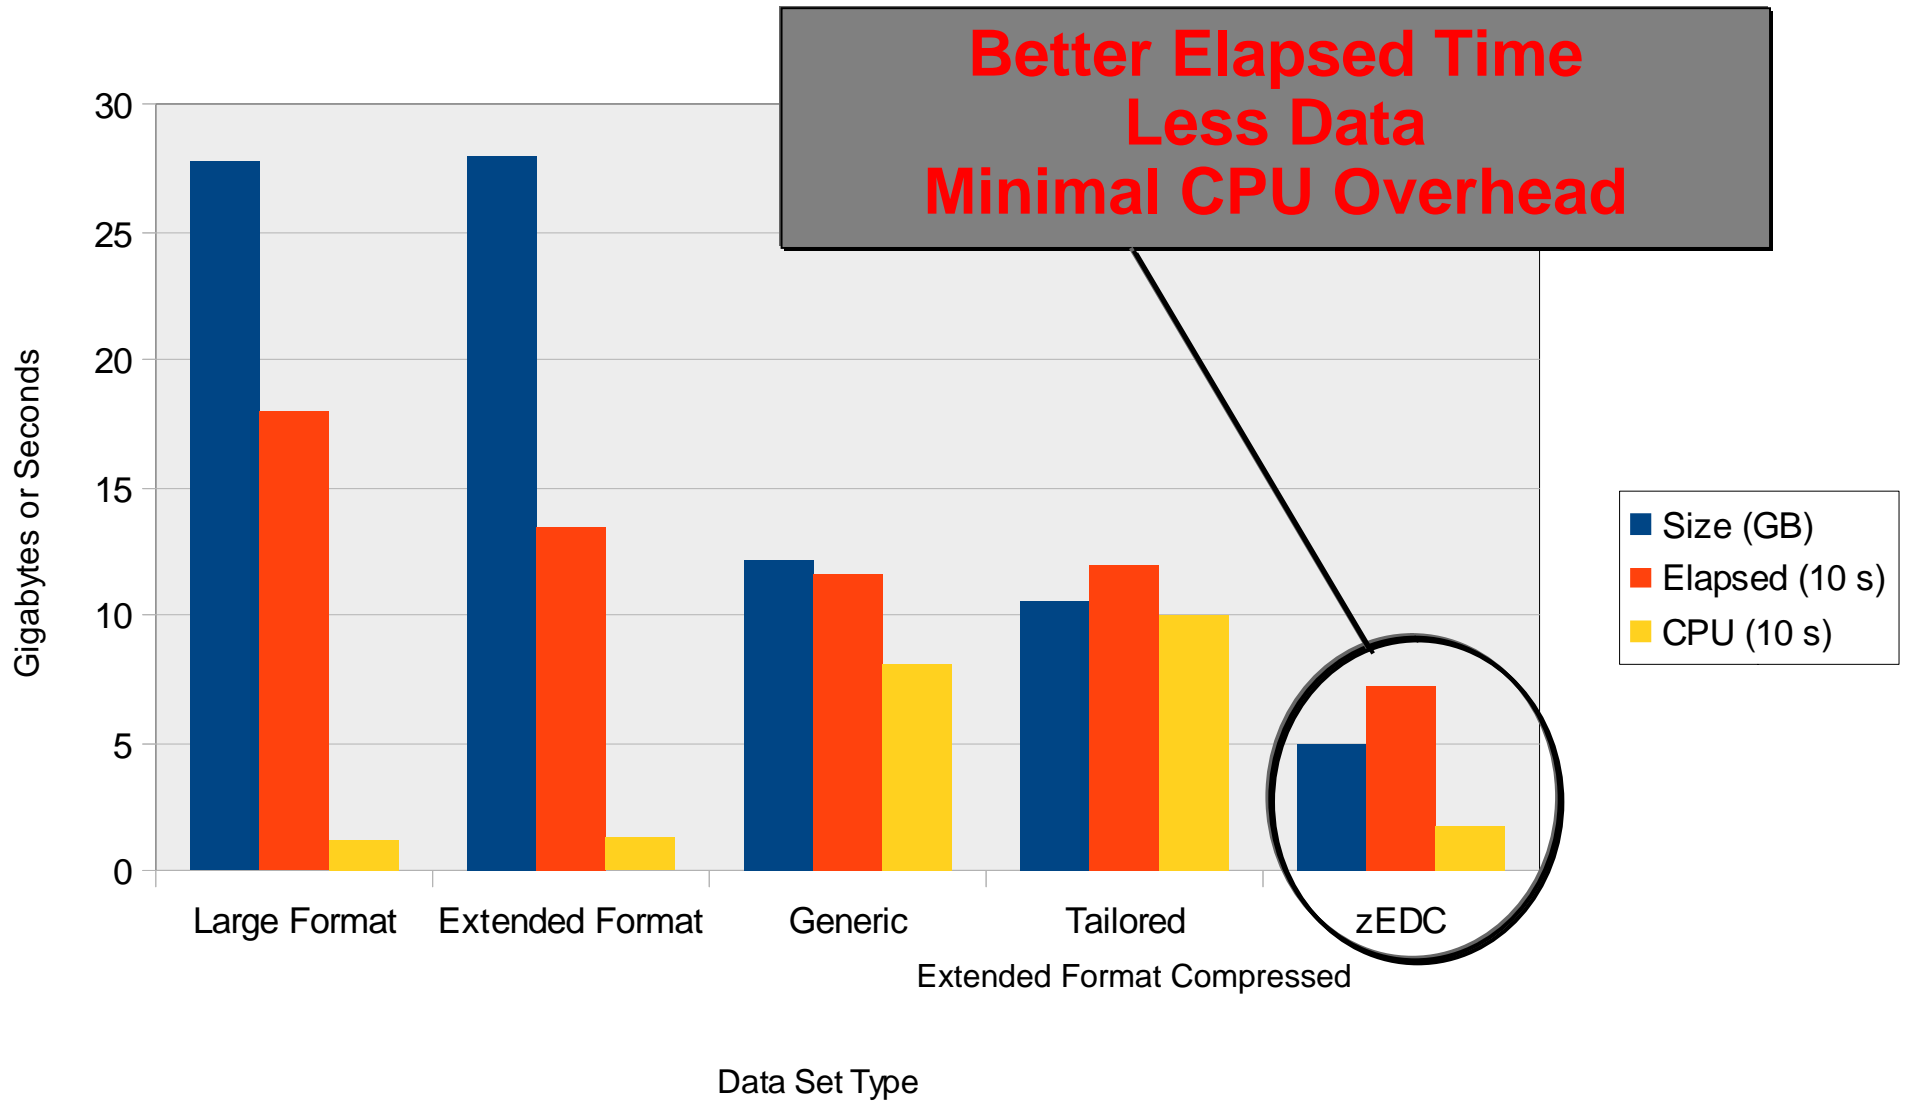
15671: zBNA Hands-on Lab, Thursday, 4:15 PM

zEC12 - Data Compression Express - zEDC

- The z Enterprise Data Compression (zEDC) Express offering provides a low-cost data compression to z/OS system services and applications
 - Implemented as a Peripheral Component Interconnect Express (PCIe) device
 - The compression function is provided via FPGA firmware
 - Can install up to 8 devices in a single processor
 - Each device is sharable by up to 15 LPARs
- Exploiters will see the following benefits
 - Increased performance for reading and writing compressed data
 - Reduced disk space
- Exploiters:
 - SMF Logstreams
 - Extended Format BSAM and QSAM data sets
 - IBM SDK for z/OS, Java Technology Edition, V7 R1
 - IBM Sterling Connect:Direct for z/OS Standard Edition V5.2
 - IBM Encryption Facility for z/OS
 - z/VM 6.3 support for guest exploitation of zEDC



QSAM/BSAM zEDC – Value!



Disclaimer: Based on projections and/or measurements completed in a controlled environment. Results may vary by customer based on individual workload, configuration and software levels.

RMF Support for zEDC Express

- New support in RMF will provide information on PCI Express based functions
 - zEnterprise Data Compression (zEDC) capability using zEDC Express
 - RDMA (Remote Direct Memory Access) over Converged Enhanced Ethernet
- Information added via new subtype to SMF 74 (74.9) called PCI Express Based Function Activity
- Written by RMF Monitor III
 - General PCIE Activity – both zEDC and RDMA activity
 - Hardware Accelerator Activity
 - Hardware Accelerator Compression Activity
 - Types of data provided:
 - I/O Queue Time
 - I/O Execution Time
 - Number of compressed bytes in and out
 - Number of decompressed bytes in an out
 - Device drive buffer statistics

15710: RMF: The Latest and Greatest, Monday, 4:30 PM

15724: Understanding z/OSMF for the Perf Mgmt SysProg, Tuesday, 10:00 AM

15713: Mining Gold from RMF Mon 3 – XML Batch Reporting Facility, Wed, 8:30 AM

Mobile Workload Pricing for z/OS

■ Benefits

- Improves the cost of growth for mobile transactions processed in System z environments such as CICS, IMS, DB2, MQ and WAS
- Mobile Workload Pricing (MWP) for z/OS enhances Sub-Capacity pricing
 - Mitigates the impact of Mobile on MLC charges where higher transaction volumes cause a spike in processor utilization
 - Normalizes the rate of transaction growth
- No infrastructure changes required, no separate LPARs needed
 - It is an enhanced way of reporting sub-capacity MSUs
 - System runs as it always has, workload execution is not altered

■ Hardware requirements

- Available to all enterprises running a zEC12 or zBC12 server (actual mobile work may run on any zEnterprise machine including z196 and z114)


■ MWP Announcement Letters – 2014-05-06

- | | | | |
|-------------|-----------------------------|-------------|---------------|
| • AP14-0219 | Asia Pacific | • JP14-0279 | Japan |
| • A14-0429 | Canada | • LP14-0279 | Latin America |
| • ZP14-0280 | Europe, Middle East, Africa | • 214-223 | United States |

Mobile Workload Pricing Reporting Process


- New Mobile Workload Reporting Tool (MWRT) – available 06/30/2014
 - A new Windows-based Java tool to report sub-capacity MSUs and make adjustments to reported LPAR MSUs based on Mobile transaction data
 - Standard SCRT methodology plus new feature to adjust for Mobile workload impact
 - New tool will replace SCRT for customers who take advantage of Mobile Workload Pricing
- Customers must track mobile transactions and produce a file showing mobile CPU consumption each month
 - Record mobile transaction data, including CPU seconds, on an hourly basis per LPAR
 - Load the resulting data file into MWRT each month (IBM-specified CSV format)
 - Run MWRT and submit the results to IBM each month (Replaces SCRT process)
- MSU adjustments and monthly peak calculation for billing
 - MWRT will **subtract 60%** of the reported Mobile MSUs from a given LPAR in each hour, adjusting the total LPAR MSU value for that hour
 - This will function like a partial “off-load” from a software pricing perspective
 - When LPAR value is adjusted, all software running in the LPAR benefits from lower MSUs
 - Tool will calculate monthly MSU peak for a given machine using the adjusted MSU values

Example: Sample LPAR MSU values by hour



- SCRT calculates the Rolling 4-Hour Average (R4HA) MSU peak
 - All workloads are included

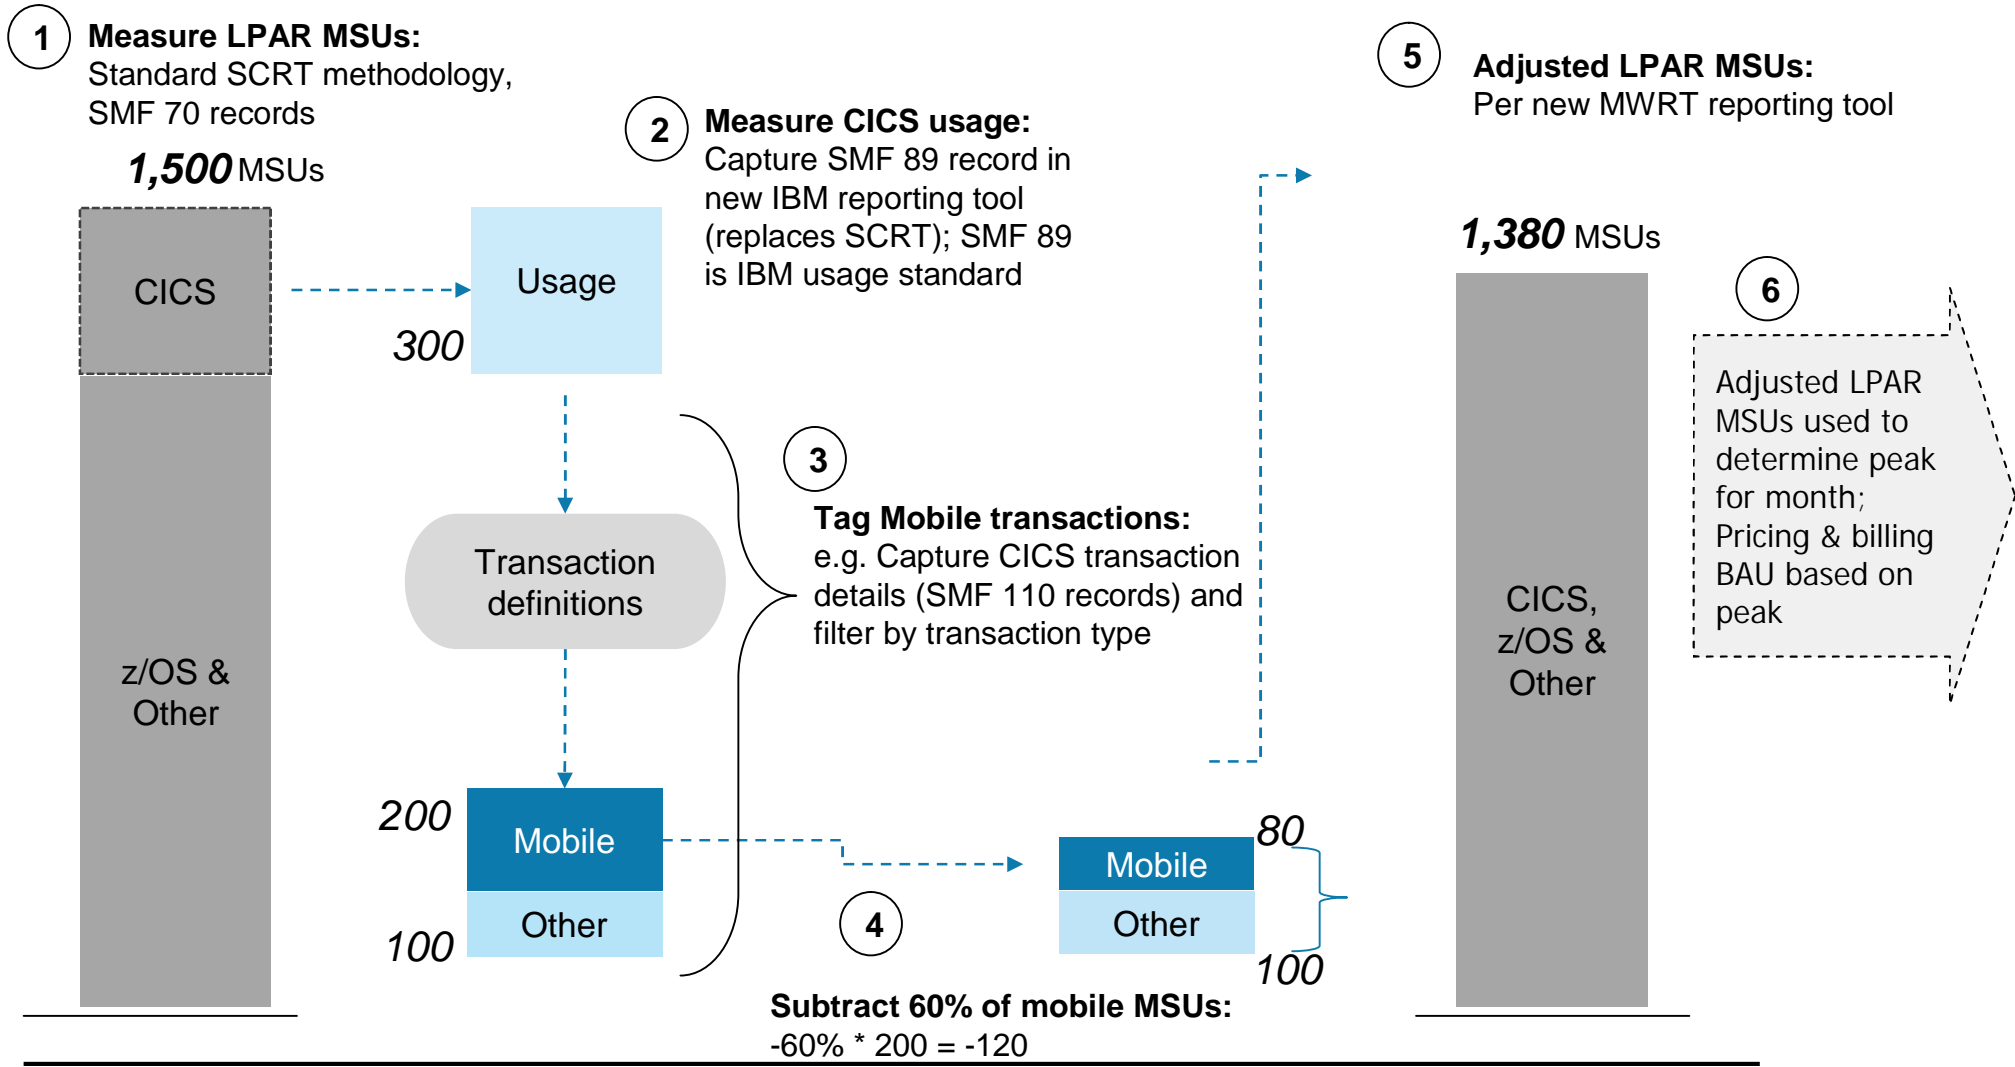
LPAR MSU values adjusted for mobile contribution



- MWRT removes 60% of the Mobile workload, interval-by-interval
 - Non-mobile workload is unchanged
 - Billing for the month is based upon the newly calculated R4HA curve after the mobile workload has been reduced

***Provides benefit when Mobile workloads contribute to monthly peak MSUs;
 Off-peak MSU adjustments will not affect MSUs used for billing.***

Example: reducing Mobile impact to LPAR peak



LPAR MSUs for billing (Standard)

z/OS/Other 1,500
CICS 1,500

LPAR MSUs for billing (Adjusted)

z/OS/Other 1,380
CICS 1,380

Figures are for illustrative purposes only.

Identifying Mobile Transaction Workload

- Customers are responsible for processing their mobile transaction data into a predefined format to be loaded into MWRT for each sub-capacity reporting period.
 - IBM must approve the data gathering methodology.
- The data must consist of **general purpose processor CPU seconds** for each mobile transaction program summarized by hour by LPAR for all machines processing mobile transactions.
 - Detailed instructions, including CSV file format, available in the MWRT user's guide.
- Mobile Workload Pricing Defining Programs:

5655-S97	CICS TS for z/OS V4	5635-A04	IMS V13
5655-Y04	CICS TS for z/OS V5	5655-DSQ	IMS DB VUE V12
5722-DFJ	CICS VUE V5	5655-DSM	IMS DB VUE V13
5635-DB2	DB2 V9 for z/OS	5655-L82	WS MQ for z/OS V6
5605-DB2	DB2 10 for z/OS	5655-R36	WS MQ for z/OS V7
5615-DB2	DB2 11 for z/OS	5655-W97	WS MQ for z/OS V8
5697-P12	DB2 VUE V9	5655-VUE	WS MQ VUE V7
5697-P31	DB2 10 VUE	5655-VU8	WS MQ VUE V8
5697-P43	DB2 11 VUE	5655-N02	WebSphere App Server for z/OS V7
5635-A02	IMS V11	5655-W65	WebSphere App Server for z/OS V8
5635-A03	IMS V12		

Worksheet example

IBM MWP contract supplement Customer Worksheet

The information from this worksheet will be used to prepare the official contract Supplement (Z126-6628) to the IBM Mobile Workload Pricing Addendum (Z126-6300).

Mobile Application name:	Bank Account Mobile App
MWP Defining Program(s):	CICS, IMS

Mobile Application Details

Details regarding the Mobile Application specified above are in the table below. Fill out one row of information in the table for each MWP Defining Program used by the Mobile Application named above.

A. MWP Defining Program	B. Data source	C. Method to distinguish mobile transactions from other workload	D. Client process for capturing and processing the mobile data
<i>e.g., CICS</i>	<i>e.g., SMF 110</i>	<i>e.g., All transaction types originating from mobile devices (e.g. iPhone app traffic) have been assigned unique names and routed to a specific region.</i>	<i>e.g., Filter mobile transaction types by name from all transactions using a SAS program, and sum the general purpose processor CPU seconds by hour for the affected LPARs.</i>
<i>e.g., IMS</i>	<i>e.g., IMS Logs</i>	<i>e.g., Enable IMS Transaction Level Statistics to produce 56FA log records for the mobile transactions; All mobile transactions contain the word "mobile" in the <u>LUName</u>.</i>	<i>e.g., Extract the mobile transaction details from the IMS log records using IMS PA and sum the general purpose processor CPU seconds by hour for the affected LPARs.</i>

New RSM APAR - OA41968

- New IEASYSxx LFAREA parameter INCLUDE1MAFC
 - LFAREA=(64M,INCLUDE1MAFC)
 - Specifies the 1 MB pages are to be included in the available frame count (RCEAFC)
- RSM changes to:
 - Performs less paging when there is an abundance of available fixed 1M pages
 - More often break up fixed 1M pages to satisfy 4K page demand
 - Attempt to coalesce broken up fixed 1M pages when there is fixed 1M page demand, no guarantee coalescing will be successful
- RMF APAR in Support - OA42510
 - RMF PTFs must be applied prior to specifying INCLUDE1MAFC
 - RMF uses the RCEAFC to generate some of their reports and not applying OA42510 may lead to incorrect RMF reports
- Application programs:
 - Can check the RCEINCLUDE1MAFC bit to determine if the installation specified INCLUDE1MAFC in their LFAREA specification
 - When using STGTEST SYSEVENT to get information about the amount of storage available in the system if INCLUDE1MAFC is specified, available fixed 1M pages are included in this amount
- In a future IBM z/OS release, fixed 1M pages will be unconditionally included in the available frame count regardless of whether the INCLUDE1MAFC value is specified or not

New WLM APAR - OA44526

- BLWLINTHD enhancements
- New support for blocked workloads
 - Allows lower threshold to be set
 - Defaults remain the same
- Useful for all online environments, with little to no batch workload, and use of DB2
 - Helps prevent CPU starved workloads from holding locks which impact higher priority work
- Use RMF Workload Activity Report to measure the amount of blocked workload activity

WORKLOAD ACTIVITY

```
--PROMOTED--  
BLK      3.240  
ENQ      0.000  
CRM      0.000  
LCK      0.000  
SUP      0.000
```

Check and understand why there are CPU times in any service classes

15214: WLM Update for z/OS 2.1 & 1.13, Monday, 3:00 PM

15217: Route Me, WLM – WLM Functions for Dynamic Routing, Tuesday, 11:30 AM

14745: WLM – Performing a Quick WLM Performance Checkup, Wednesday, 3:00 PM

New SMF 42 Subtype 5 and 6 Support

- OA44322 and OA44319 - OPEN
 - Planned for 3rd Quarter 2014
- Updates for SMF records
 - SMF 42 subtype 5 - Storage Class / VTOC / VVDS I/O
 - SMF 42 subtype 6 - DASD Data Set Level I/O Statistics
- New Metrics Provided
 - Read and Write counts for High Performance FICON (zHPF)
 - Improved Resolution for Average Times reported
 - 1 microsecond units Vs 128 microsecond units
 - Command Response Delay and Device Busy Delay
 - Subsets of Pend time

Note: These statements represent the current intention of IBM. IBM reserves the right to change or alter the plans in the future. IBM development plans are subject to change or withdrawal without further notice. Any reliance on this statement of direction is at the relying party's sole risk and does not create any liability or obligation for IBM.

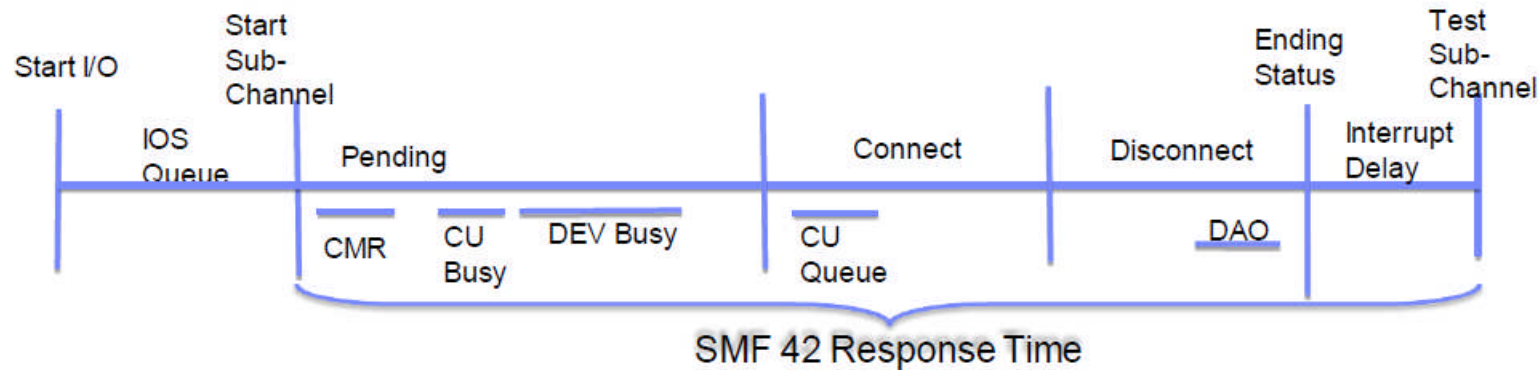
New SMF 42 Subtype 6 Support - WSC Example

					"New" zHPF Metrics	
					zHPF Reads	zHPF Writes
DSN	VOLSER	ADDR	ION	RDT	HRD	HWR
WSCID34.ISPF.ISPPROF	ZOS048	9E05	1	1	0	0
WALSH.CPSTOOLS.JCL	ZOS050	9E06	1	1	0	0
SYS1.PARMLIB	OSPPGE	9C03	2	2	0	0
SYS1.BROADCAST	OSPPGE	9C03	1	1	0	0
MCCOX.EAVTST.DATASET	EAVOL1	8000	1	1	0	0
MATERIA.OMPRSYSA.ENV	ZOS048	9E05	2	2	0	0
JPBURG.SCM.RMFOUT	ZOS049	9E06	10	10	0	0

"New" 1 Microsecond Unit Times				"Old" 128 Microsecond Unit Times				"New" / "Old" Total Resp Ratio
Total Response	Connect	Pend	Disc	Total Response	Connect	Pend	Disc	
R1U	C1U	P1U	D1U	Resp	Con	Pend	Disc	
MS	MS	MS	MS	MS	MS	MS	MS	
10.69	10.44	0.23	0.00	10.62	10.37	0.13	0.00	1.01
11.48	11.24	0.22	0.00	11.39	11.14	0.13	0.00	1.01
0.27	0.10	0.14	0.00	0.13	0.00	0.00	0.00	2.07
0.23	0.07	0.13	0.00	0.13	0.00	0.13	0.00	1.78
0.47	0.28	0.14	0.00	0.38	0.26	0.13	0.00	1.21
6.55	0.19	0.15	6.19	6.40	0.13	0.13	6.14	1.02
0.33	0.14	0.15	0.00	0.26	0.00	0.13	0.00	1.30

Note: These statements represent the current intention of IBM. IBM reserves the right to change or alter the plans in the future. IBM development plans are subject to change or withdrawal without further notice. Any reliance on this statement of direction is at the relying party's sole risk and does not create any liability or obligation for IBM.

SMF 42 Control Unit & Channel Measurement Fields

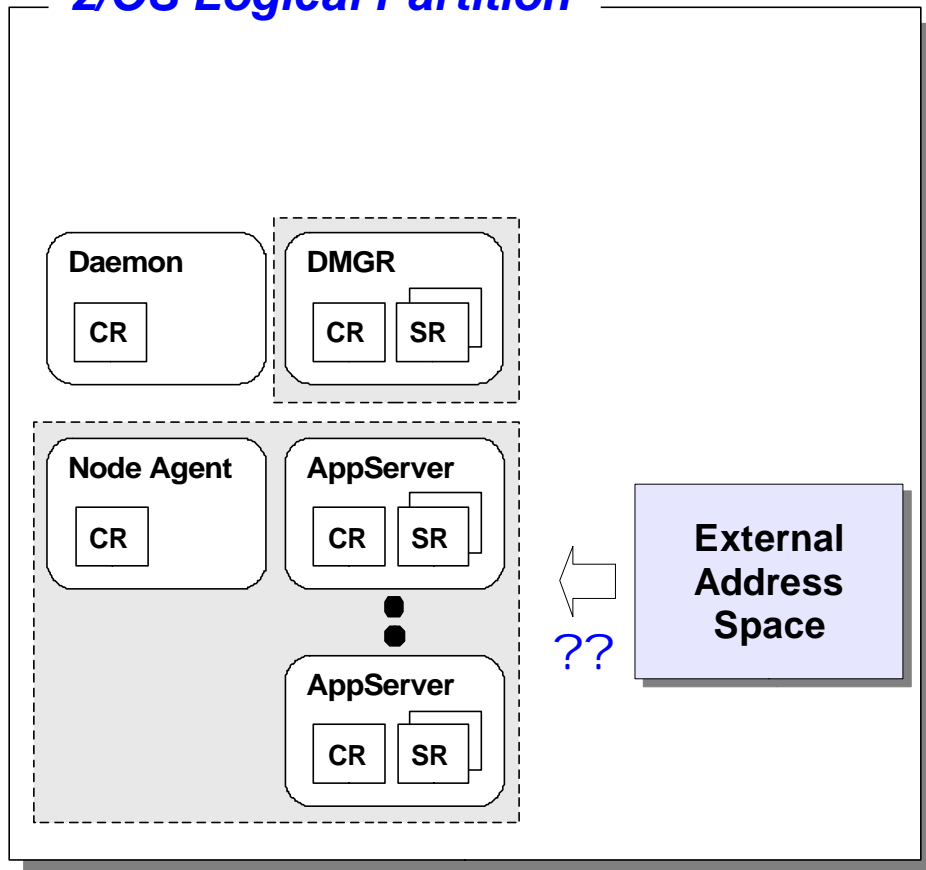


- *Response time is the elapsed time from Start Subchannel to TSCH; CPU dispatching can play a role.*
- *Service time = Pend + Connect + Disconnect*
- *Pend time includes channel selection overhead (not explicitly reported), along with CMR delay, Device Busy delay, CU Busy delay, but not CU Queue delay.*
- *Command Response delay: Starts when the channel sends the 1st command until it receives a response from the controller. One round trip through the fabric.*
- *Device Busy delay measures time associated with an initial status of “device busy” delay time. Example: device reserved by another system*
- *CU Queue delay is time spent queued in the Controller. example: extent conflicts*
- *Device Active Only - Channel End to Device End on last command when not issued at the same time. Subset of Device-Defer Time. example: writes to Synchronous PPRC Secondary*

WebSphere Optimized Local Adapters (WOLA)

It started out as a way to allow program access *into* WAS for high transaction rate batch programs. Other solutions existed, but they all had limitations:

z/OS Logical Partition



Inbound to WAS?

As more and more solutions are built based on Java EE, there is a growing desire to access them by batch, CICS and IMS programs

MQ or Web Services?

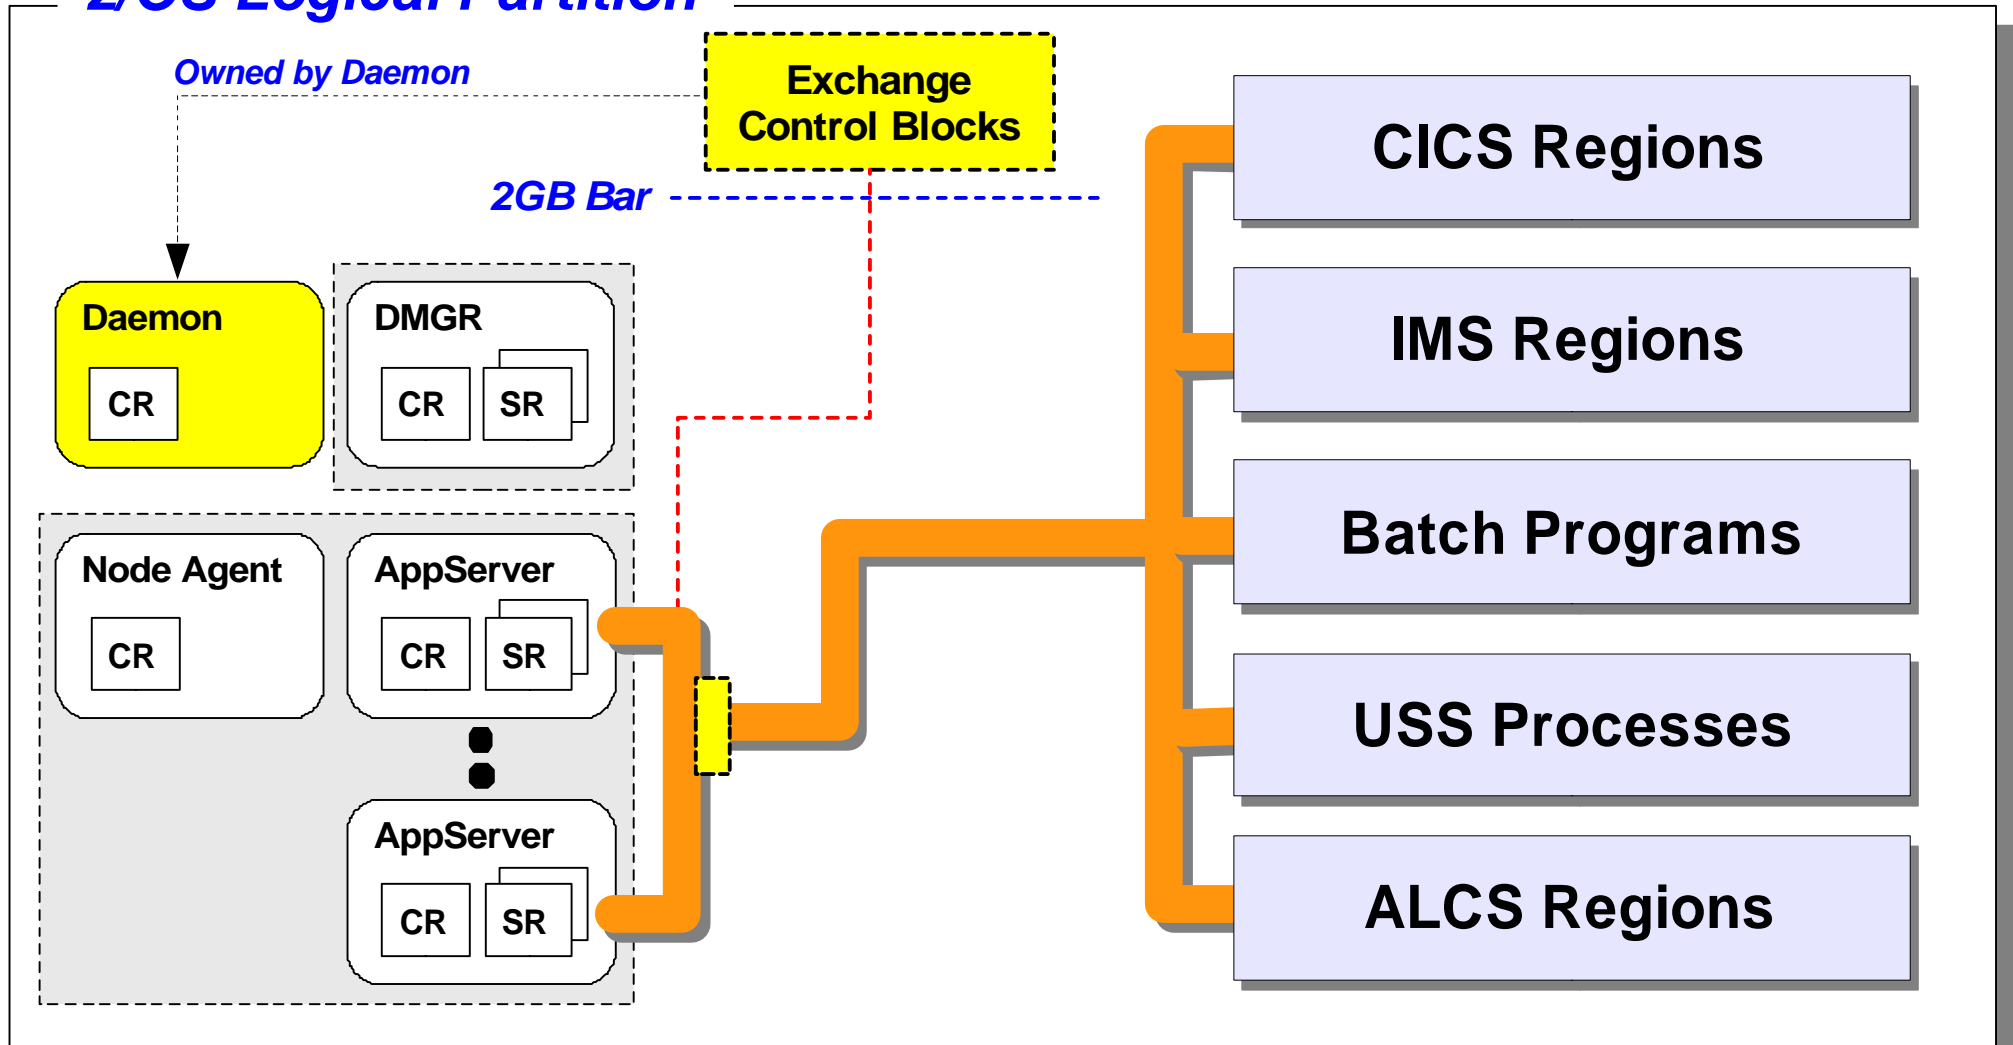
Both are very good technologies and have their role. But for very high throughput and low overhead, each has their drawbacks.

**Something else was needed ...
something *very fast* with as
little overhead per exchange
as possible**

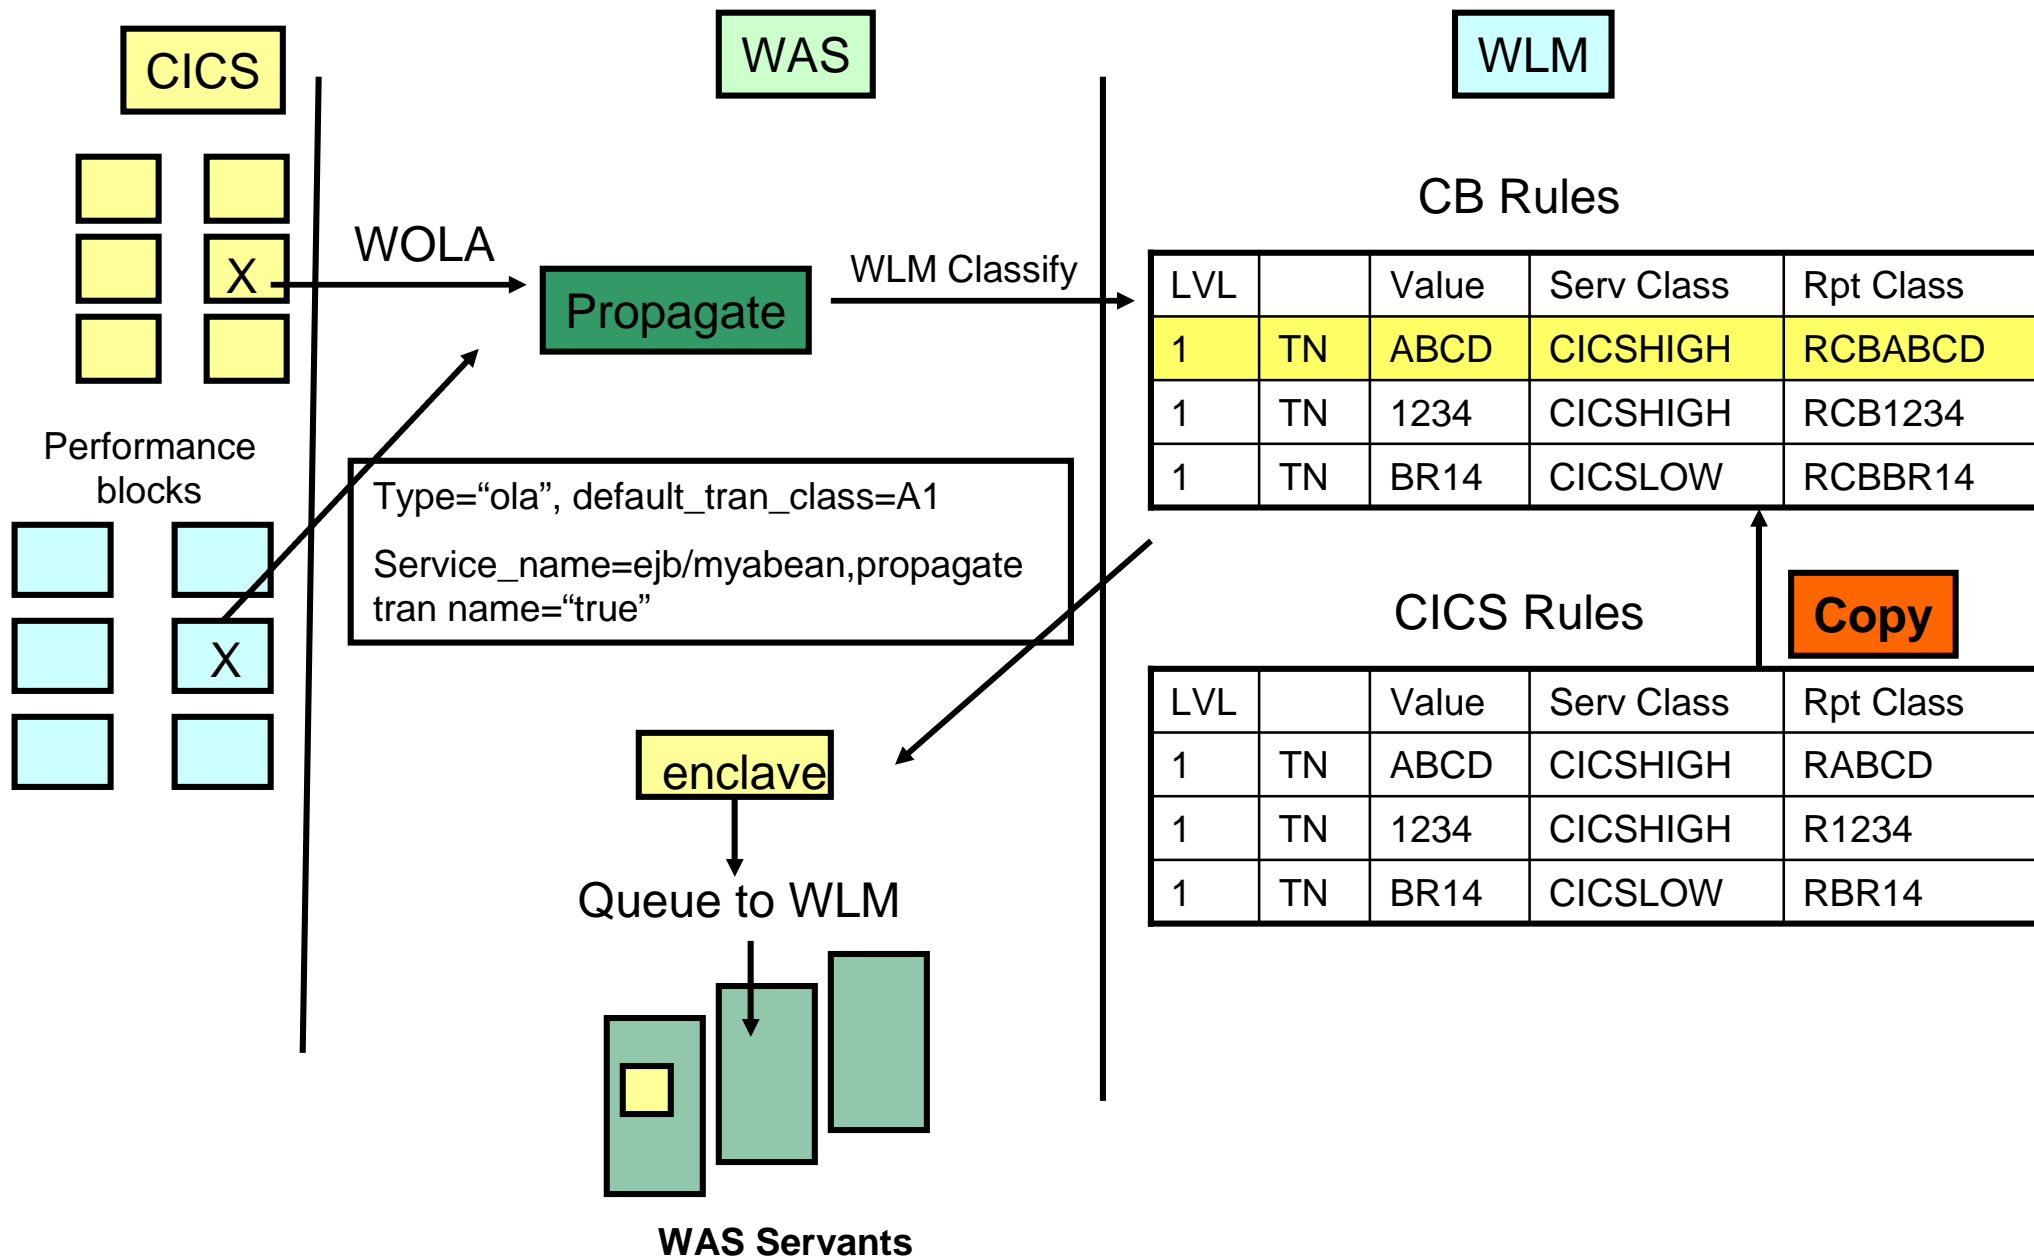
The Motivation Behind WOLA

Build on the Local Comm function and build a bi-directional technology. Outside into WAS, and WAS out to external address space

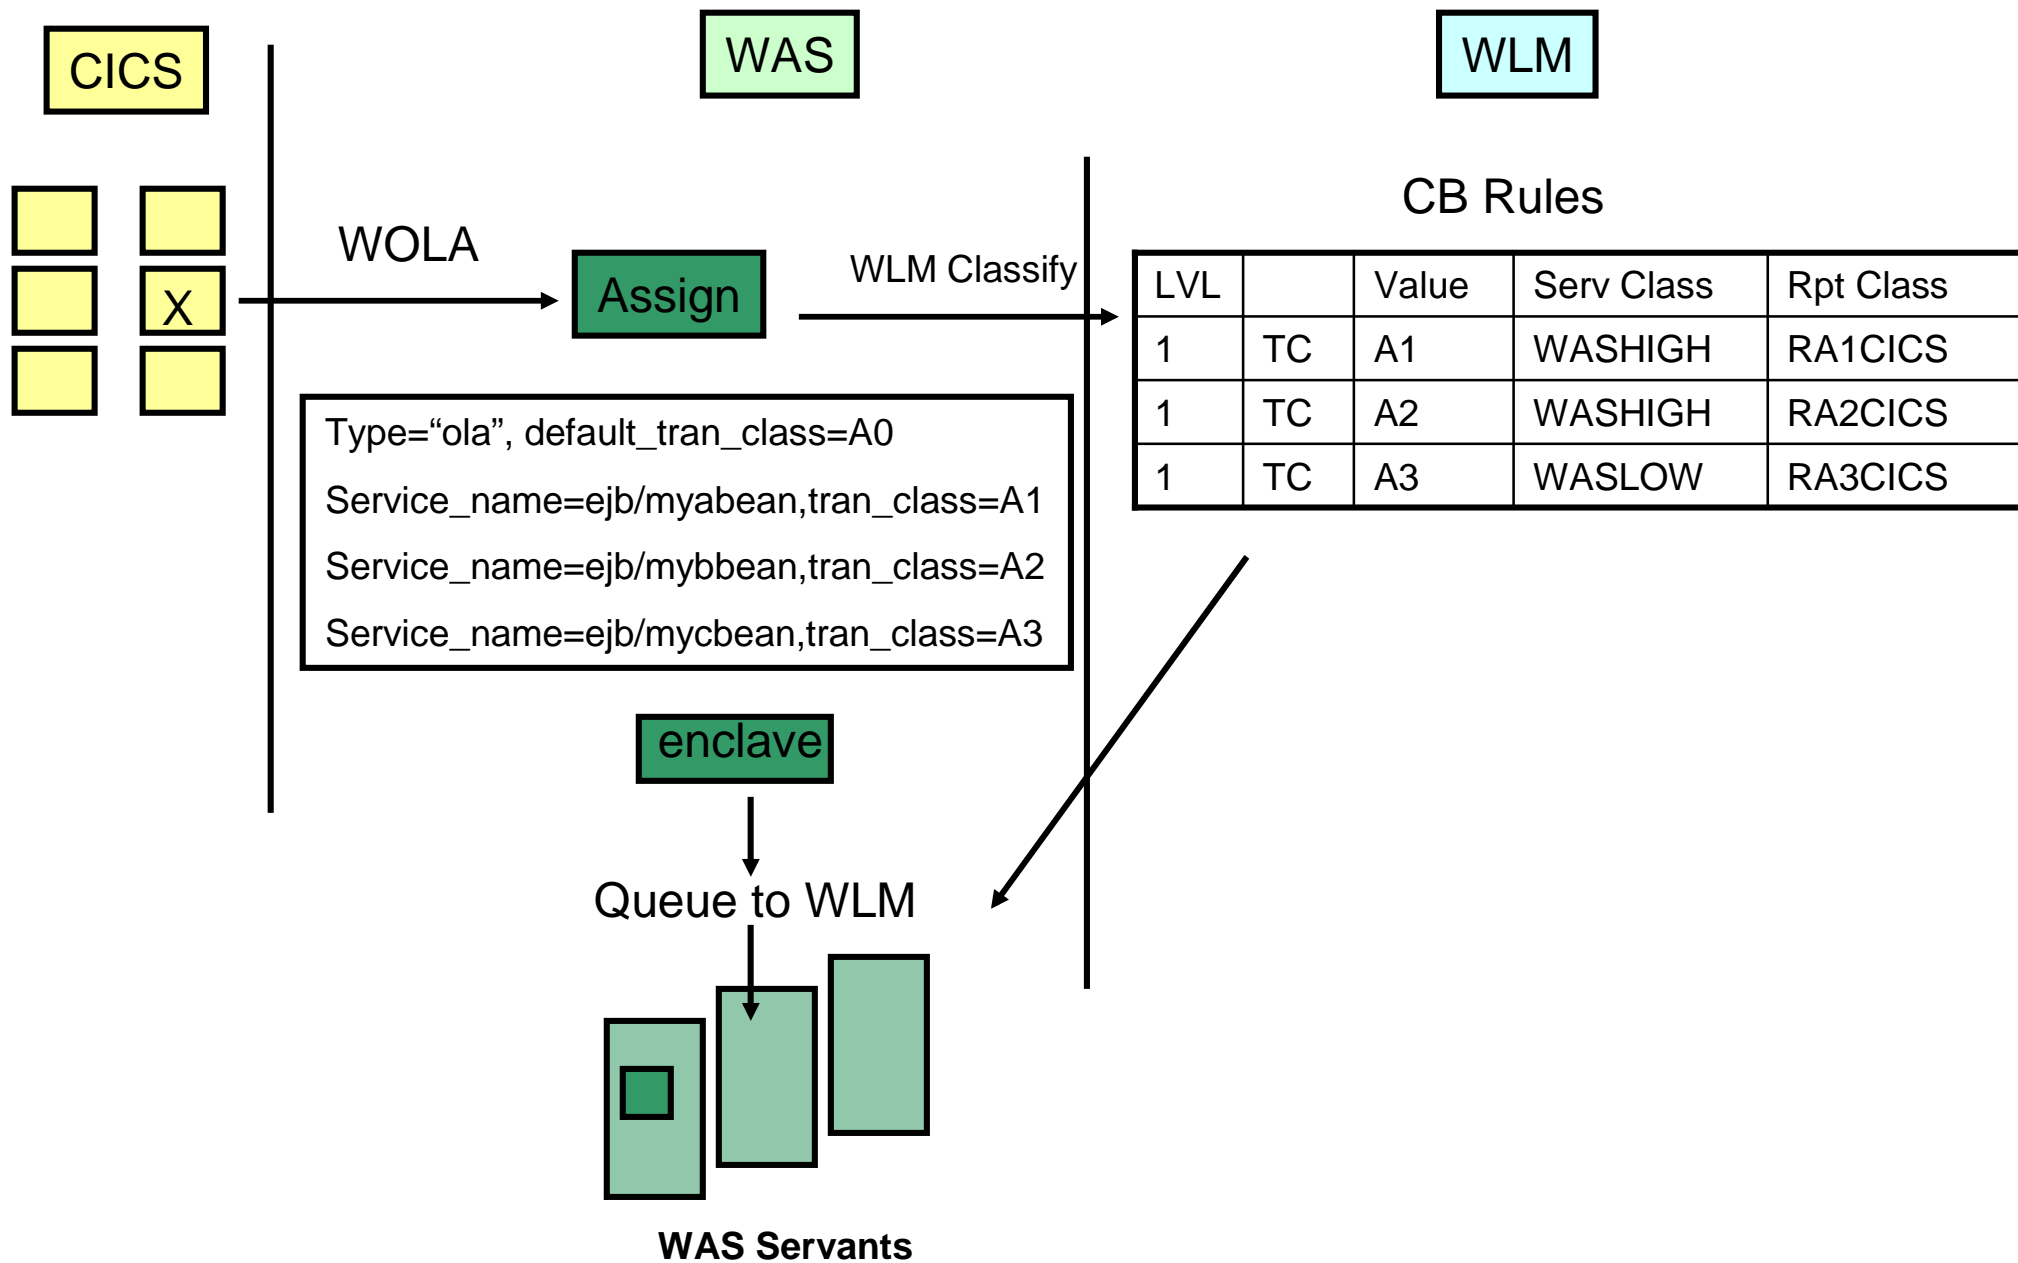
z/OS Logical Partition



WLM and WOLA - Propagate



WLM and WOLA - Assign



WLM and WOLA

■ Propagate

- Allows CICS transaction and the EJB to be viewed as a single transaction
 - A WAS owned enclave with the CICS service class is created to do the work
 - The CICS transaction is suspended until the enclave ends
- Have to be careful the number of servants per application environment are high enough to support the propagated service classes
 - If work arrives using more unique service classes than the MAX number of servants in the AE the work in those service classes wait forever
- Mixing the transaction counts and resource profiles in Client service classes (CICS, Batch, IMS) across these environments

■ Assign

- CICS transaction sends a request and finds a matching entry in the WAS WLM transaction class mapping file and uses the tran class to classify the work to a CB service class


■ More Info:

http://www14.software.ibm.com/webapp/wsbroker/redirect?version=phil&product=was-nd-zos&topic=tdat_olawlm



Addendum

- Older information which should still be understood, or make you go HmMMM.
- APARs which are still causing issues, even though they are old.



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:

*BladeCenter®, DB2®, e business (logo)®, DataPower®, ESCON, eServer, FICON, IBM®, IBM (logo)®, MVS, OS/390®, POWER6®, POWER6+, POWER7®, Power Architecture®, PowerVM®, S/390®, Sysplex Timer®, System p®, System p5, System x®, System z®, System z9®, System z10®, Tivoli®, WebSphere®, X-Architecture®, zEnterprise®, z9®, z10, z/Architecture®, z/OS®, z/VM®, z/VSE®, zSeries®

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 registered 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.

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.

PR/SM Absolute Capping

- zEC12 GA2 and zBC12 allows specification of an “**absolute capping limit**”
 - Expressed in terms of 1/100ths of a processor (0.01 to 255.0)
 - Specified independently from the LPAR weight
 - The shared partition's processing weight still dictates the logical partition priority compared to other shared logical partitions
 - Most effective for absolute caps higher than the capacity the partition's relative weight would deliver
 - Absolute capping is not recommended to be set below the capacity the logical partition's weight would deliver
 - Insensitive to capacity changes or LPAR (de)activations
 - Specified per processor type in image profile and partition controls panel
- Unlike initial capping may be used concurrently with defined capacity and/or group capacity management
 - The absolute capacity becomes effective before other capping controls
 - WLM/SRM recognizes new cap, e.g. for routing decisions
- Requires APAR OA41125
 - PTFs for z/OS 1.12 and z/OS 1.13

PR/SM Absolute Hardware Capping

Customize Image Profiles: SSYS : SOSP15 : Processor

Group Name: <Not Assigned>

Logical Processor Assignments

Dedicated processors

Select	Processor Type	Initial	Reserved
<input checked="" type="checkbox"/>	Central processors (CPs)	3	1
<input type="checkbox"/>	System z application assist processors (zAAPs)	0	0
<input checked="" type="checkbox"/>	System z integrated information processors (zIIPs)	2	0

Not Dedicated Processor Details for:

CPs zIIPs

CP Details

Initial processing weight: 100 (1 to 999) Initial capping

Enable workload manager

Minimum processing weight: 0

Maximum processing weight: 0

Absolute Capping: None Number of processors (0.01 to 255.0) 1.23

Buttons: Cancel, Save, Copy Profile, Paste Profile, Help

New Coupling Thin Interrupts – DYNDISP=THIN

Goal: Expedite the dispatching of the CF partition

- Generate Coupling Thin Interrupts to wake up and dispatch a shared engine in a timely fashion to service work as opposed to having the engine wait for PR/SM to perform its processing
- Once the CF image gets dispatched, the existing “poll for work” logic in both z/OS and CFCC can be used to locate and process the work
- CF will give up control when work is exhausted or when LPAR kicks it off the shared engine
- Requires:
 - CFCC Level 19
 - z/OS 2.1 or z/OS 1.12 / z/OS 1.13 with PTFs
- White Paper: Coupling Thin Interrupts and Coupling Facility Performance in Shared Processor Environments
 - <http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP102400>

15203: Coupling Technology Overview and Planning, Thursday, 8:00 AM

DYNDISP Modes

CF Polling	Dynamic CF Dispatching	Coupling Thin Interrupts
DYNDISP=NO	DYNDISP=YES	DYNDISP=THIN
<p>LPAR Time Slicing</p>	<p>CF Time based algorithm for CF engine sharing</p>	<p>CF releases shared engine if no work left to be done</p>
<ul style="list-style-type: none"> –CF does not “play nice” with other shared images sharing the processor –CF controls processor long after work is exhausted 	<ul style="list-style-type: none"> –CF does its own time slicing –More effective engine sharing than polling –Blind to presence or absence of work to do –No Interrupt Available 	<ul style="list-style-type: none"> –Event Driven Dispatching –CF relies on generation of thin interrupt to dispatch processor when new work arrives –Now the most effective use of shared engines across multiple CF images

z/OS Performance **HOT** Topics

Session: 14022

Kathy Walsh
IBM Corporation

