



z/OS Performance “HOT” Topics

Kathy Walsh

IBM

August 10, 2015

Session Number: 17861



#SHAREorg



SHARE is an independent volunteer-run information technology association that provides education, professional networking and industry influence.

Copyright (c) 2015 by SHARE Inc. Except where otherwise noted, this work is licensed under <http://creativecommons.org/licenses/by-nc-sa/3.0/>



Agenda



■ Performance and Capacity Planning Topics

- Introduction of z Systems z13 Processor
- Overview of SMT
- CPU MF and HIS Support
- zPCR Latest Status
- New RSM and WLM APARS
- New SMF Support
- Blocked Workloads
- Change in EXCP counts for zFS

▶ Addendum

- Older APARs or Performance Information



United States [change]

[Home](#) [Solutions](#) [Services](#) [Products](#) [Support & downloads](#) [My IBM](#)

Welcome Kathy Walsh [Not you?] [IBM Sign in]

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.

Techdocs Library

- Flashes
- Presentations & tools
- Technotes & tips
- FAQs
- White papers
- Solution scenario profiles
- Customer support plans
- Sizings
- Auxiliary Material
- Search Techdocs
- Techdocs feedback

Related links

- Redbook publications
- IBM Software Support Handbook

New to Techdocs?

Is this your first visit to **Techdocs** (the Technical Sales Library)?

→ [Learn more](#)

Returning to Techdocs?

Looking for what's new in the **Techdocs Library**?

→ [Latest updates](#)

Need Technical Support?


Looking for support resources or other documents and tools?

→ [Support & downloads](#)

z13 Overview



z13



- Machine Type
 - 2964
- 5 Models
 - N30, N63, N96, NC9 and NE1
- Processor Units (PUs)
 - 39 (42 for NE1) PU cores per CPC drawer
 - Up to 24 SAPs per system, standard
 - 2 spares designated per system
 - Dependant on the H/W model - up to 30, 63, 96, 129, 141 PU cores available for characterization
 - 85 LPARs, increased from 60
 - Sub-capacity available for up to 30 CPs
 - 3 sub-capacity points
- Memory
 - RAIM Memory design
 - System Minimum of 64 GB
 - Up to 2.5 TB GB per drawer
 - Up to 10 TB for System and up to 10 TB per LPAR (OS dependant)
 - LPAR support of the full memory enabled
 - 96 GB Fixed HSA, standard
 - 32/64/96/128/256/512 GB increments
 - Flash Express

Session: z13 Performance, Thurs 1:45

IBM z13 versus zEC12 Hardware Comparison




zEC12

— CPU

- 5.5 GHz
- Enhanced Out-Of-Order

— Caches

- L1 private 64k i, 96k d
- L2 private 1 MB i + 1 MB d
- L3 shared 48 MB / chip
- L4 shared 384 MB / book



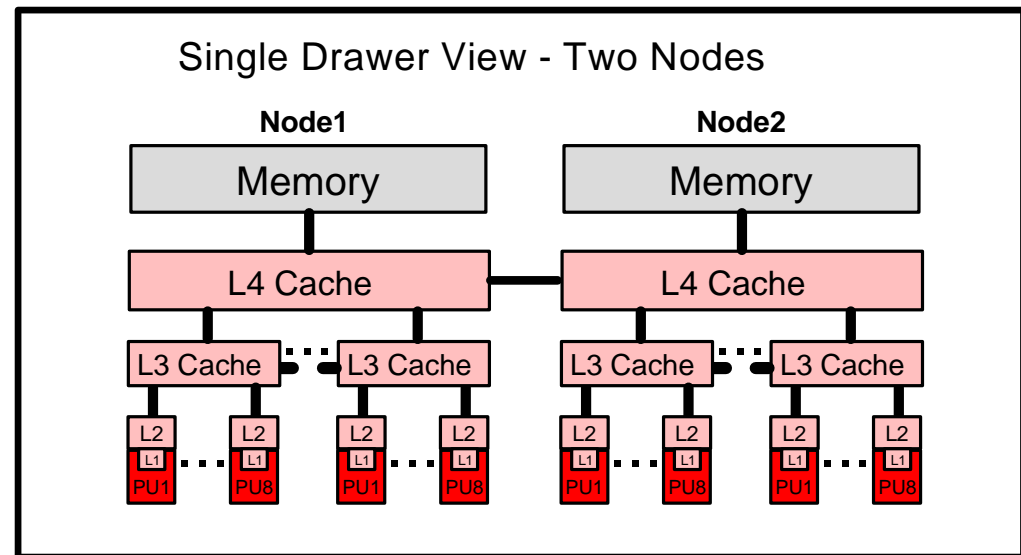
z13

— CPU

- 5.0 GHz
- Major pipeline enhancements

— Caches

- L1 private 96k i, 128k d
- L2 private 2 MB i + 2 MB d
- L3 shared 64 MB / chip
- L4 shared 480 MB / node
- plus 224 MB L3 NIC Directory



z13 Processor and Memory Assignment and Optimization



- Default processor assignments by POR, MES adds, and On Demand activation:
 - Assign IFLs and ICFs to cores on chips in “high” drawers working down, CPs and zIIP in low drawers working up
 - Objective: Keep partitions using IFLs and ICFs “away” from z/OS partitions on CPs and zIIPs in different drawers if possible

- PR/SM assigns available memory and logical processors at activation
 - Logical Processor specified in Image Profile assigned a *core* if Dedicated or if shared a “*home*” drawer, node and chip.
 - If Logical Processor becomes a HiperDispatch “Vertical High”, the Shared Logical Processor is assigned a specific core

 - Ideally assign all memory in one drawer with the processors if everything “fits”, with memory striped across drawers with processors if memory or processors must be split

z13 Processor and Memory Assignment and Optimization



- PR/SM optimizes resource assignment when triggered
 - Triggers: Partition activation or deactivation or significant processor entitlement changes, dynamic memory increases or processor increases or decreases (e.g. by CBU) or MES change
 - Optimization
 - Examine partitions in priority order by size of “processor entitlement” (dedicated processor count or shared processor pool allocation by weight)
 - Changes Logical Processor “home” drawer/node/chip assignment
 - Moves Logical Processor to different chips, nodes, drawers (LPAR Dynamic PU Reassignment)
 - Relocates partition memory to active memory in a different drawer(s) using Dynamic Memory Relocation (DMR)
 - If available but inactive memory hardware is present (e.g. hardware driven by Flexible or Plan Ahead) in a drawer where more active memory would help: activate it, reassign active partition memory to it, and deactivate the source memory hardware, again using DRM
 - PR/SM can use all memory hardware but concurrently enables no more memory than the client has paid to use

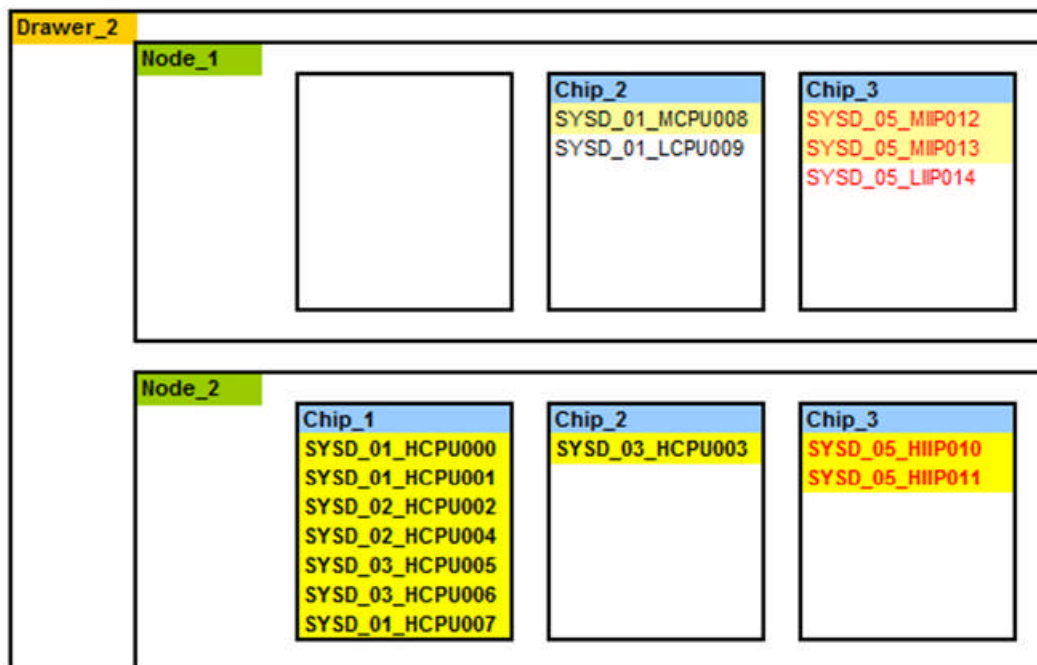
SMF 99 Subtype 14 – HiperDispatch Topology



- SMF 99 Subtype 14 contains HiperDispatch Topology data including:
 - Logical Processor characteristics: Polarity (VH, VM, VL), Affinity Node, etc.
 - Physical topology information
 - zEC12: Book / Chip
 - Z13: Drawer / Node / Chip
- Written **every 5 minutes** or when a Topology change occurs
- Topology Changes:
 - Configuration change or weight change
 - Driven by IRD weight management
 - Record provides a “Topology Change” indicator to show when the topology changed
- Recommend: Collect SMF 99 subtype 14s for all LPARs on the CEC
 - Record has a single LPAR scope so need all LPARs to get total picture
- New *WLM Topology Report* available to process SMF 99.14 records
 - http://www.ibm.com/systems/z/os/zos/features/wlm/WLM_Further_Info_Tools.html#Topology



Example: SMF 99 z13 Topology Report



SSSS_nn_PTYPnnn

SSSS - SMFID

nn - Affintiy Node

P - Polarity (H M L)

TYP - **CPU** | **IIP** | IFL | ICF

nnn - Core ID



CPU Measurement Facility

- Available on all System z processors since the z10
- 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
- z/OS HIS started task
 - Gathered on an LPAR basis
 - Writes SMF 113 records
- 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
- In z/OS V2.2:
 - HIS will no longer require USS definitions
 - Modify HIS (F HIS) command is restructured



Session 17313: SMF 113 Processor Cache Counter Measurements – Wed 11:15

Session 17556: 2015 CPU MF Update – Thurs 10:00

Session 17665: Processor Reporting: RMF and Hardware Instrumentation Services (HIS) – Fri 10:00

Looking for z13 Migration “Volunteers”



- Want to validate / refine Workload selection metrics

Looking for “Volunteers”

(3 days, 24 hours/day, SMF 70s, 72s, 113s per LPAR)

“Before” and “After”

Production partitions preferred

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

Simultaneous Multithreading (SMT)



- Simultaneous multithreading allows instructions from one or two software threads to execute on a zIIP or IFL processor core
- SMT helps to address memory latency, resulting in an overall capacity* (throughput) improvement per core
 - Thread performance (instruction execution rate per thread) may be faster running in single thread mode
 - SMT is not available for CPs so LSPR ratings do not include it
- Capacity improvement is variable depending on workload. For AVERAGE workloads the estimated capacity* of a z13:
 - zIIP is 38% greater than a zEC12 zIIP
 - IFL is 32% greater than a zEC12 IFL
- SMT exploitation: z/OS V2.1 + PTFs in an LPAR for zIIPs and z/VM V6.3 + PTFs for IFLs



*Capacity and performance ratios are based on measurements and projections using standard IBM benchmarks in a controlled environment. 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.

Simultaneous Multithreading (SMT)



- IBM z13 supports two instruction streams (threads) per core
- Threads share core resources:
 - In time: Address translator, instruction execution units, pipeline slots, ...
 - Cache misses provide opportunities for other thread
 - Processor ensures fairness between threads
 - In space: Data and instruction caches, branch tables, TLBs, ...
 - A thread can't necessarily execute instructions instantly and must compete and win use of desired core resources shared between threads
 - Each thread has its own state and can do most things a core can do
 - E.g., take interruptions, start I/O, load wait PSW, signal other threads
- READY TO RUN threads share core
 - Threads NOT READY TO RUN are still unproductive while resolving cache miss
 - Core resources are productive when either READY TO RUN Thread is executing

Exploiting SMT on z13 with zIIPs



- z/OS V2R1 SMT APARs must be applied
 - OA43366 (BCP), OA43622 (WLM), OA44439 (XCF)
 - z/OS manages threads according to the SMT Mode
- Define a LOADxx PROCVIEW CORE|CPU
 - Setting is for the life of the IPL
 - PROCVIEW CORE on z13 enables SMT support
- New IEAOPTxx parameter to control zIIP SMT mode
 - MT_ZIIP_MODE=1|2
 - MT_ZIIP_MODE=2 for 2 active threads
 - When PROCVIEW CPU is specified the zIIP MT Mode is always 1
 - Update IEAOPTxx to change mode and issue SET OPT=xx command
 - Requires HD=YES, dynamic switching to HD=NO is not allowed
- New LOADxx and IEAOPTxx controls ONLY available on z/OS V2R1 and higher



RMF CPU Activity – zIIP CPU Activity

C P U A C T I V I T Y

CPU		z/OS V2R1		SYSTEM ID SYSD		DATE 02/25/2015			
MODEL		CPC CAPACITY 4452		RPT VERSION V2R1 RMF		TIME 10.45.00			
H/W MODEL		CHANGE REASON=NONE		SEQUENCE CODE 00000000008DA97		HIPERDISPATCH=YES			
---CPU---		TIME %				--- MT % ---		LOG PROC	
NUM	TYPE	ONLINE	LPAR BUSY	MVS BUSY	PARKED	PROD	UTIL	SHARE	%
A	IIP	100.00	88.46	75.43	0.00	89.45	79.13	100.0	HIGH
B	IIP	100.00	81.96	68.84	0.00	87.79	71.95	100.0	HIGH
C	IIP	100.00	59.18	49.67	0.00	88.09	52.13	53.2	MED
D	IIP	100.00	31.78	27.10	0.00	88.09	27.99	53.2	MED
E	IIP	100.00	11.13	9.68	0.00	91.11	10.14	0.0	LOW
TOTAL/AVERAGE			54.50	43.26		88.91	48.27	306.4	
----- MULTI-THREADING ANALYSIS -----									
CPU	TYPE	MODE	MAX CF	CF	AVG TD				
	CP	1	1.000	1.000	1.000				
	IIP	2	1.384	1.225	1.585				



RMF CPU Activity – Workload Activity

W O R K L O A D A C T I V I T Y

z/OS V2R1 SYSPLEX WSCZPLEX DATE 02/25/2015 INTERVAL 04.59.999 MODE = GOAL PAGE 1
 RPT VERSION V2R1 RMF TIME 10.45.00

POLICY ACTIVATION DATE/TIME 02/24/2015 17.37.33

----- SERVICE POLICY

REPORT BY: POLICY=WLM2013
 WSC WLM 2013 policy

-TRANSACTIONS-		TRANS-TIME	HHH.MM.SS.TTT	--DASD I/O--	---SERVICE---	SERVICE TIME	---APPL %---	--PROMOTED--	----STORAGE----
AVG	96.73	ACTUAL	22	SSCHRT 202.7	IOC 800326	CPU 1012.133	CP 7.89	BLK 0.000	AVG 25289.43
MPL	96.73	EXECUTION	22	RESP 3.0	CPU 75322K	SRB 2.206	AAPCP 0.00	ENQ 0.000	TOTAL 1677310
ENDED	58207	QUEUED	0	CONN 0.5	MSO 0	RCT 0.005	IIPCP 2.88	CRM 0.000	SHARED 7812.45
END/S	194.02	R/S AFFIN	0	DISC 0.0	SRB 164186	IIT 0.341		LCK 1.684	
#SWAPS	122	INELIGIBLE	0	Q+PEND 0.7	TOT 76286K	HST 0.000	AAP N/A	SUP 0.000	-PAGE-IN RATES-
EXCTD	0	CONVERSION	0	IOSQ 1.8	/SEC 254288	AAP N/A	IIP 238.72		SINGLE 0.0
AVG ENC	30.40	STD DEV	15			IIP 991.030			BLOCK 0.0
REM ENC	0.00				ABSRPTN 2629				SHARED 0.0
MS ENC	0.00				TRX SERV 2629				HSP 0.0



Service Time is in Normalized SMT-1 time

$$\text{IIP APPL \%} = 991.030 * 100 / (299.999 * 1.384) = 238.69\%$$

RMF CPU Activity – PDR Changes



PARTITION DATA REPORT

z/OS V2R1 SYSTEM ID SYSD START 02/25/2015-10.45.00 INTERVAL 000.04.59
 RPT VERSION V2R1 RMF END 02/25/2015-10.50.00 CYCLE 1.000 SECONDS


MVS PARTITION NAME	UOSP02	NUMBER OF PHYSICAL PROCESSORS	96
IMAGE CAPACITY	742	CP	36
NUMBER OF CONFIGURED PARTITIONS	85	IFL	48
WAIT COMPLETION	NO	IIP	12
DISPATCH INTERVAL	DYNAMIC		

----- PARTITION DATA -----					-- LOGICAL PARTITION PROCESSOR DATA --					-- AVERAGE PROCESSOR --	
NAME	S	WGT	NUM	TYPE	-----DISPATCH	TIME DATA-----	LOGICAL PROCESSORS	---	PHYSICAL PROCESSORS	---	
					EFFECTIVE	TOTAL	EFFECTIVE	TOTAL	LPAR MGMT	EFFECTIVE	TOTAL
UOSP02	A	240	5	IIP	00.13.33.129	00.13.37.513	54.21	54.50	0.12	22.59	22.71
UOSP01	A	700	11	IIP	00.00.01.307	00.00.01.342	0.04	0.04	0.00	0.04	0.04
PHYSICAL						00.00.00.564			0.02		0.02
					-----	-----			-----	-----	-----
TOTAL					00.13.34.437	00.13.39.421			0.14	22.62	22.76

Session 17636: RMF: The Latest and Greatest – Mon 3:15 PM
Session 17632: CF Activity Report Review – Tue 1:45 PM



Early SMT Testing



HonorPriority=NO

HonorPriority=YES

zPCR Latest Status



- Why use zPCR
 1. LSPR Processor Capacity Ratios Tables
 2. LPAR Configuration Capacity Planning

- Version 8.7c (2/16/2015)
 - The IBM z Systems (z13) processor family has been added, with 231 General Purpose models (90 sub-capacity and 141 full-speed) and 141 IFL models
 - LSPR data is now based on z/OS 2.1
 - New support added for Absolute Capping
 - On z13 processors SMT can be activated for z/OS running on zIIPS and z/VM running on IFLs. SMT is currently supported only by z/OS 2.1 and z/VM 6.3
 - In Advanced-Mode, the number of definable LPAR configurations has increased from 7 to 10

Session 17563: zPCR Capacity Sizing Lab - Part 1 Intro – Wed 10:00 AM
Session 17553 zPCR Capacity Sizing Lab - Part 2 Lab - Wed 6:00 PM

SMT Support in zPCR



Partition Detail Report
zPCR V8.6

IBM Confidential

Partition Detail Report

Based on LSPR Data for IBM System z Processors
 Study ID: Sample zPCR Advanced-Mode Study
 #2 Planned zNext 2964-707
 Description: XYZ Production on zNext 2964-707
zNext/700 Host = 2964-N30/700 with 13 CPs: GP=7 zIIP=1 IFL=4 ICF=1
9 Active Partitions: GP=3 zIIP=2 IFL=3 ICF=1
 Capacity basis: 2094-701 @ 593.00 MIPS for a shared single-partition configuration
 Capacity for z/OS on z10 and later processors is represented with HiperDispatch turned ON

Include	Partition Identification					Partition Configuration				Capping		Partition Capacity	
	No.	Type	Name	SCP	Workload	Mode	LCPs	Weight	Weight %	✓	ABS	Minimum	Maximum
<input checked="" type="checkbox"/>	1	GP	LP-01	z/OS-2.1	Average	SHR	7	700	53.85%	<input type="checkbox"/>		5,553	10,312
<input checked="" type="checkbox"/>	2	GP	LP-02	z/OS-2.1	Average	SHR	4	400	30.77%	<input type="checkbox"/>		3,193	5,931
<input checked="" type="checkbox"/>		zIIP	LP-02	z/OS-2.1	Average	SHR	1	200	50.00%	<input type="checkbox"/>		770	1,540
<input checked="" type="checkbox"/>	3	GP	LP-03	z/OS-2.1	Avg-High	SHR	3	200	15.38%	<input type="checkbox"/>		1,472	4,099
<input checked="" type="checkbox"/>		zIIP	LP-03	z/OS-2.1	Avg-High	SHR	1	200	50.00%	<input type="checkbox"/>		731	1,461
<input checked="" type="checkbox"/>	4	IFL	LP-04	z/VM	Average/LV	SHR	4	400	64.00%	<input type="checkbox"/>		3,936	6,149
<input checked="" type="checkbox"/>	5	IFL	LP-05	Linux	Average/L	SHR	2	200	32.00%	<input type="checkbox"/>		1,968	3,075
<input checked="" type="checkbox"/>	6	IFL	LP-06	Linux	Low-Avg/L	SHR	1	25	4.00%	<input type="checkbox"/>		259	1,617

Table View Controls

Display zAAP/zIIP/IFL Partitions

With Associated GP Separate by Pool

Show GP Pool Specialty Pools

All Partitions GP zAAP zIIP

Includes Only IFL ICF

Capacity Summary by Pool

CP Pool	RCPs	Partitions	LCPs	SHR LCP:R...	Capacity
GP	7	3	14	2.000	10,218
zAAP	None				n/a
zIIP	1	2	2	2.000	1,501
IFL	4	3	7	1.750	6,162
ICF	1	1	1	All DED	1,349
Totals	13	9	24		19,229

Add SMT Benefit to Capacity Results

Host Summary Modify SCP/Workload LCP Alternatives zAAP/zIIP Loading

For significant configuration changes, capacity comparisons should be considered to have a +/-5% margin-of-error
 Upgrading the processor family is considered a significant configuration change

Input fields have white background; Single-click a "selection field" for drop-down list; Double click a "key-in field" to open.

SMT Support in zPCR



Partition Detail Report

Graph Documentation

zPCR V8.6

IBM Confidential

Partition Detail Report
 Based on LSPR Data for IBM System z Processors
 Study ID: Sample zPCR Advanced-Mode Study
 #2 Planned zNext 2964-707
 Description: XYZ Production on zNext 2964-707
zNext/700 Host = 2964-N30/700 with 13 CPs; GP=7 zIIP=1 IFL=4 ICF=1
9 Active Partitions: GP=3 zIIP=2 IFL=3 ICF=1
 Capacity basis: 2094-701 @ 593.00 MIPS for a shared single-partition configuration
 Capacity for z/OS on z10 and later processors is represented with HiperDispatch turned ON

Include	Partition Identification					Partition Configuration				Capping		SMT Benefit	Partition Capacity	
	No.	Type	Name	SCP	Workload	Mode	LCPs	Weight	Weight %	✓	ABS		Minimum	Maximum
<input checked="" type="checkbox"/>	1	GP	LP-01	z/OS-2.1	Average	SHR	7	700	53.85%	<input type="checkbox"/>		5,553	10,312	
<input checked="" type="checkbox"/>	2	GP	LP-02	z/OS-2.1	Average	SHR	4	400	30.77%	<input type="checkbox"/>		3,193	5,931	
<input checked="" type="checkbox"/>		zIIP	LP-02	z/OS-2.1	Average	SHR	1	200	50.00%	<input checked="" type="checkbox"/>		962	1,925	
<input checked="" type="checkbox"/>	3	GP	LP-03	z/OS-2.1	Avg-High	SHR	3	200	15.38%	<input type="checkbox"/>		1,472	4,099	
<input checked="" type="checkbox"/>		zIIP	LP-03	z/OS-2.1	Avg-High	SHR	1	200	50.00%	<input checked="" type="checkbox"/>		913	1,827	
<input checked="" type="checkbox"/>	4	IFL	LP-04	z/VM	Average/LV	SHR	4	400	64.00%	<input type="checkbox"/>		4,723	7,379	
<input checked="" type="checkbox"/>	5	IFL	LP-05	Linux	Average/L	SHR	2	200	32.00%	<input type="checkbox"/>		2,361	3,690	

Table View Controls

Display zAAP/zIIP/IFL Partitions
 With Associated GP Separate by Pool

Show GP Pool Specialty Pools
 All Partitions GP zAAP zIIP
 Includes Only IFL ICF

Capacity Summary by Pool

CP Pool	RCPs	Partitions	LCPs	SHR LCP:R...	Capacity
GP	7	3	14	2.000	10,218
zAAP	None				n/a
zIIP	1	2	2	2.000	1,876
IFL	4	3	7	1.750	7,394
ICF	1	1	1	All DED	1,349
Totals	13	9	24		20,836

Hide SMT Benefit from Capacity Results Clear all SMT Benefit values when Hiding Global SMT Benefit Show SMT SCP Restrictions

Host Summary Modify SCP/Workload LCP Alternatives zAAP/zIIP Loading

For significant configuration changes, capacity comparisons should be considered to have a +/-5% margin-of-error
 When displaying default SMT Benefit value, margin-of-error is +/-10%; Larger SMT Benefit values, margin-of-error will be greater
 Upgrading the processor family is considered a significant configuration change

Input fields have white background; Single-click a "selection field" for drop-down list; Double click a "key-in field" to open.
 Capacity Values include SMT Benefit for one or more zIIP and/or IFL partitions

Global SMT

zPCR V8.6

SMT Benefit Set by Partition Type

✓	On Return, Set All Partitions of Type	SMT Benefit
<input checked="" type="checkbox"/>	zIIP	25%
<input checked="" type="checkbox"/>	IFL	20%

Restore zPCR Defaults

Tamper Resistant SMF



- Use digital signatures to detect a change, addition/removal of an SMF record from a group of records
 - Increases the value of SMF data by making it verifiable
 - Applications recording to SMF can transparently leverage this support
 - Industry standard encryption
- What is a digital signature?
 - A way to ensure the source and validity of data
 - The signer will first hash the data and then encrypt the hash with their private key – The encrypted hash is the signature
 - The consumer of the data can hash the same data and decrypt the signature to obtain the signer's hash
 - The hashes will then be compared – When these values match then the data is considered verified
 - Result – detection and deterrence of data corruption

Tamper Resistant SMF



- The SMF data is signed on the way to System Logger
 - As each record is written to the logstream it is hashed
 - Running hash maintained per unique SMF type/subtype
 - Periodically, the hash will be digitally signed and the signature data will be recorded to the logstream as a signature record
 - On the global interval a signature is created for all data hashed during the interval and recorded to the logstream
 - These operations are performed with the private key
- IFASMF DL understands signature records and will optionally move them with the records of an associated SMF type/subtype
- IFASMF DP can verify a set of SMF records has not been tampered with when signature records are available.
 - This operation performed with the public key
- Need to create a public/private key pair via ICSF
 - SMF does not care about the type of key (clear or secure) as long as the available hardware can support it

New SMF Support for BCPii



- Authorized z/OS application that can:
 - Monitor status or capacity changes
 - Obtain configuration data related to CPC or image
 - Image and CPC Commands (like IPL, (de)Activate, OOCoD etc.)
 - Change temporary capacity
 - Query and update LPAR settings
 - Set activation profiles

- BCPii now cuts SMF 106 whenever an application issues an API (HWISET or HWICMD) which successfully modifies hardware resources
 - Provides an audit trail
 - Subtype 1 (HWISET) SYS(TYPE(106(1)))
 - Subtype 2 (HWICMD) SYS(TYPE(106(2)))

- SHARE Top 50 requirement SSMVSE12018

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 z/OS V2.2 fixed 1M pages will be unconditionally included in the available frame count regardless of whether the INCLUDE1MAFC value is specified or not

z/OS V2.2 RSM Scalability Enhancements

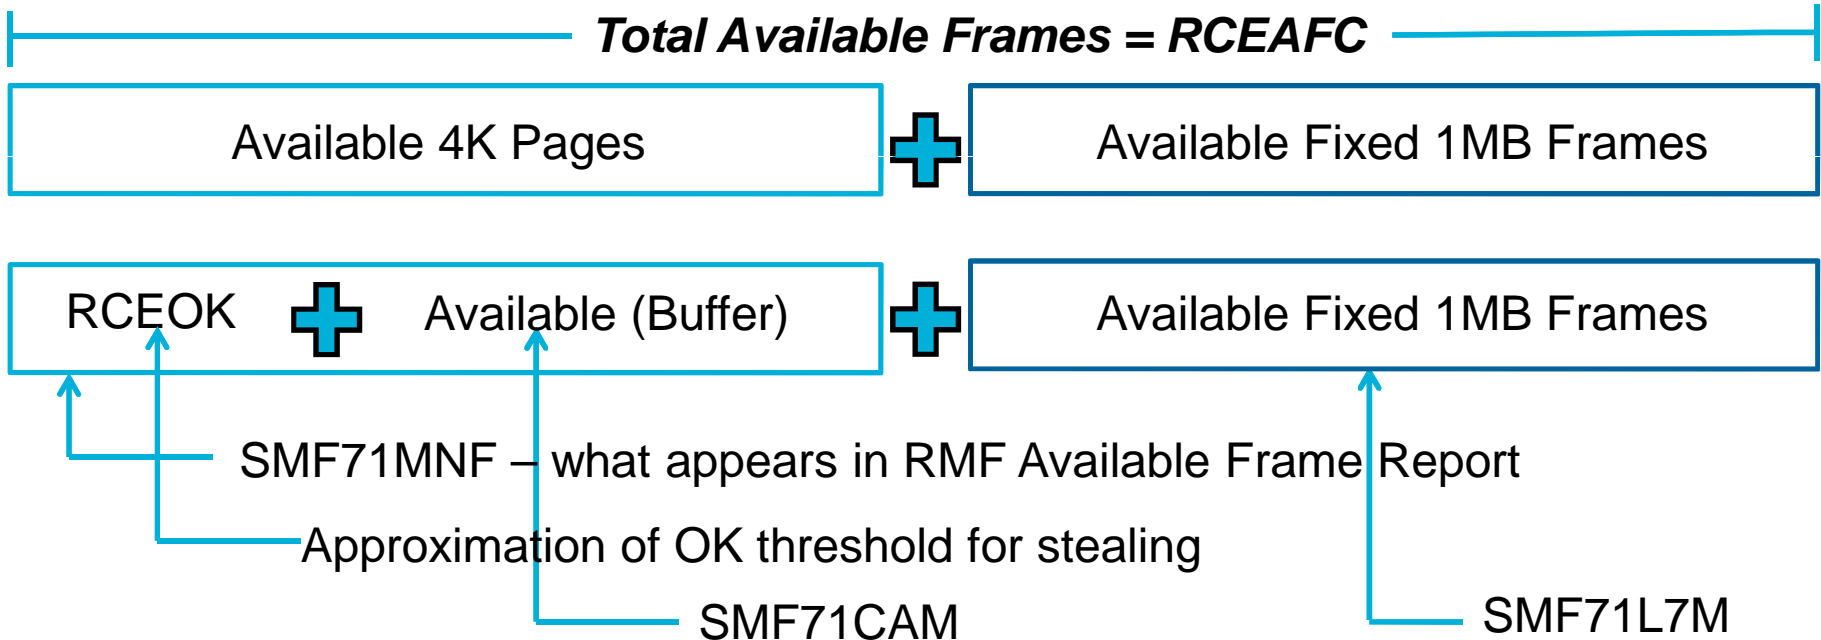


- RSM was modified to support large amounts of real storage:
 - Initialization of real storage is more efficient
 - Management of 1MB frames is improved
 - Improvements to configuring storage online and offline
 - Steal processing can find eligible frames more efficiently
- RSM was modified to increase concurrency both at the system and application levels
 - Multiple page faults can be dealt with concurrently within an address space or with common storage
 - Page fixing and unfixing can occur concurrently within an address space or with common storage
 - Getmain/Freemain of storage can occur concurrently with page fixing/freeing and page faults
 - Less contention on available frame queues
- Rolled Down to z/OS V2.1 via APARs OA44207 and OA44436

Available Frame Calculations



- Available frames are tracked in RCEAFC
 - New IEASYSxx parm INCLUDE1MAFC changes RCEAFC calculation by including Available 1MB Frames in count
 - Min, Max, and Average values are tracked – Minimums shown below



z/OS 2.2 will include Available Fixed 1MB Frames in both SMF71MNF and SMF71CAM

New Hiperdispatch APAR – OA47968



- HiperDispatch performance improvements of the park/unpark algorithm
 - The HiperDispatch park/unpark algorithm now considers the processor topology information of the vertical low (VL) processors
 - For IBM z13 and above hardware only the HiperDispatch park/unpark algorithm now considers the average VL processor utilization when unparking VLs due to available processor free capacity (white space)

Session 17637: WLM Update for z13, z/OS 2.1 and 2.2 – Wed 8:30 AM

Session 17312: WLM - Effective Setup and Usage of WLM Report Classes – Wed 1:45 PM

Session 17322: WLM in One Page – Fri 11:15 AM

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 heavy 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

New WAS APAR – PI33798



- Reduce contention on CML lock under heavy load
- WLM in z/OS 2.1 added 2 new parameters SERVCLS and INSERVCLS to the IWM4ECRE enclave create macro
- When Websphere (Liberty profile or Traditional WAS) detects it is running on z/OS 2.1 or later, it will invoke IWM4ECRE with the new SERVCLS and INSERVCLS parameters
 - Initially IWM4ECRE is invoked with no service class token and it returns a service class token for the specified classification values
 - Websphere saves the service class token for the specified classification values and later on any requests with the same classification values Websphere will look and find a service class token and pass this to IWM4ECRE
 - Passing a service class token on enclave create, enables WLM to skip classification
 - The classification part of enclave create is where WLM gets the CML
 - Once service class tokens for the different types of classification being done have been obtained there will be less contention on the CML lock which provides performance benefits

z/OS 2.2: New zFS Support for EXCP Counts




- z/OS 2.1 and Prior: zFS updates EXCP counts for every cache page touched by the current operation on the user task
 - Number of cached pages really touched not consistent even for single system cases
 - In directory operations, pages touched is determined by current layout of the directory and what was cached in memory
 - Also if forwarding directory operations to owners, the client did not touch pages so if some directory operations were run on a client instead of an owner the EXCP counts may be way off
- In z/OS 2.2 changes are made:
 - File reads and writes operate the same - number of cached pages touched because clients do direct file IO so this is an exact count of user cache pages touched
 - For directory updates and reads, zFS will determine the minimum number of meta cache pages required for the operation and use that number
 - Ensures client and owner are consistent and the accounting is simple and reasonably accurate
 - Example:
 - A file create requires 2 anode table pages and a directory page to be updated, so we simply use the value 3 for write count for a file create all
 - Ensures consistency whether or not the application is run on the owner or the client

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.



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.



RSM APAR - OA44207 – New Function

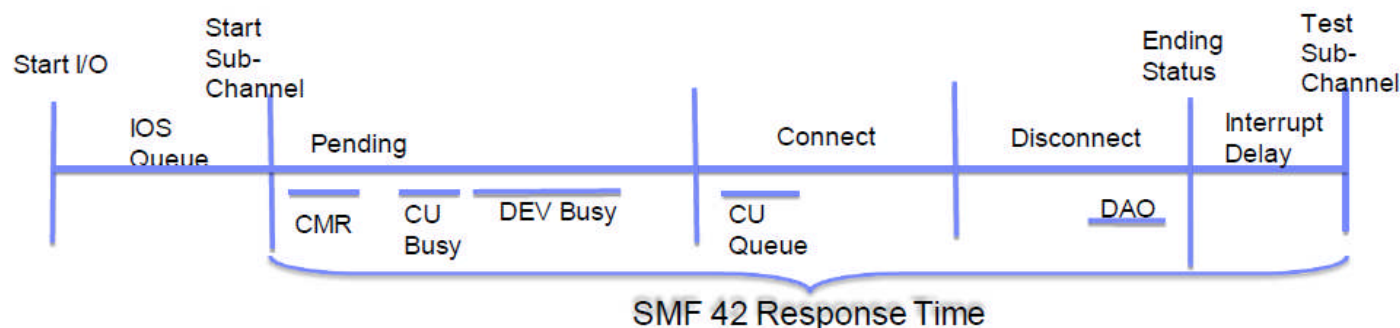
- Relieve scalability issues for LPARs running with large amounts of real storage between 256GB to 1 TB, or with a large LFAREA
 - Reduces amount of time RSM runs disabled during RSM initialization
 - Real storage used to contain the PFTE for the LFAREA is moved to a new area to help prevent depletion of preferred storage below the bar
 - Size of PLAREA is reduced when the LFAREA is very large allowing more preferred storage above the bar
 - New structures are defined for managing LFAREA frames which will reduce amount of time RSM runs disabled
 - Improvements to the Config STOR, offline command
 - RSM frame steal is improved by skipping over areas not eligible for steal

RSM APAR - OA44436 – New Function

- Support for LPARs running with large amounts of real storage
 - Automatic reconfigurable storage reduction when an RSU parm may cause insufficient preferred frames for IPL. MSG IAR0004I issued.
 - Indicator bit in RCE to determine when all online real frames at IPL have been initialized and available for use
 - Support for OA45505 to allow SRM to gather more granular global steal stats



SMF 42 Control Unit & Channel Measurement Fields



- Response time is 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

Single Instruction Multiple Data (SIMD) Vector Processing



- Single Instruction Multiple Data (SIMD)
 - A type of data parallel computing that can accelerate code with integer, string, character, and floating point data types
- Providing optimized SIMD math & linear algebra libraries to minimize effort on the part of middleware/application developers
- Providing compiler built-in functions so software applications can leverage as needed (e.g. for use of string instructions)
- OS / Hypervisor Support:
 - z/OS: 2.1 SPE available at GA
 - Linux: IBM is working with its Linux Distribution partners to support new functions/features
 - No z/VM Support for SIMD
 - LOADxx MACHMIG can be used in z/OS to disable SIMD at IPL time
 - Compiler exploitation
 - IBM Java
 - XL C/C++ on zOS
 - XL C/C++ on Linux on z
 - Enterprise COBOL
 - Enterprise PL/I

[MASS - Mathematical Acceleration Sub-System](#)
[ATLAS - Automatically Tuned Linear Algebra Software](#)


Workloads		
Java 8 on z/OS	C/C++ Compiler built-ins for SIMD operations (z/OS and Linux on z Systems)	MASS & ATLAS Math Libraries (z/OS and Linux on z Systems)
SIMD Registers and Instruction Set		



IBM z13: SIMD – Single Instruction Multiple Data

Hardware for exploiting data-parallelism

- Large uniform data-set that needs the same operation performed on each element
- Can offer dramatic speedup to data-parallel operations (matrix ops, string processing, etc)



4x less iterations in the SIMD loop

```
//SISD C example, adding two arrays
for (i=0;i<128;i++)
{
  c[i] = a[i] + b[i];
}
```

```
//SIMD C example, adding two arrays
for (i=0;i<32;i++)
{
  vec_add(c[i*4], a[i*4], b[i*4]);
}
```



IBM z13: SIMD – Single Instruction Multiple Data

IBM z13 running Java 8 on z/OS Single Instruction Multiple Data (SIMD) vector engine exploitation

java.lang.String exploitation

- compareTo
- compareToIgnoreCase
- contains
- contentEquals
- equals
- indexOf
- lastIndexOf
- regionMatches
- toLowerCase
- toUpperCase
- getBytes

java.util.Arrays

- equals (primitive types)

String encoding converters

For ISO8859-1, ASCII, UTF8, and UTF16

- encode (char2byte)
- decode (byte2char)


Auto-SIMD

- Simple loops
(eg. Matrix multiplication)

Primitive operations are between 1.6x and 60x faster with SIMD

(Controlled measurement environment, results may vary)

Application Serving – SSL-Enabled DayTrader3.0

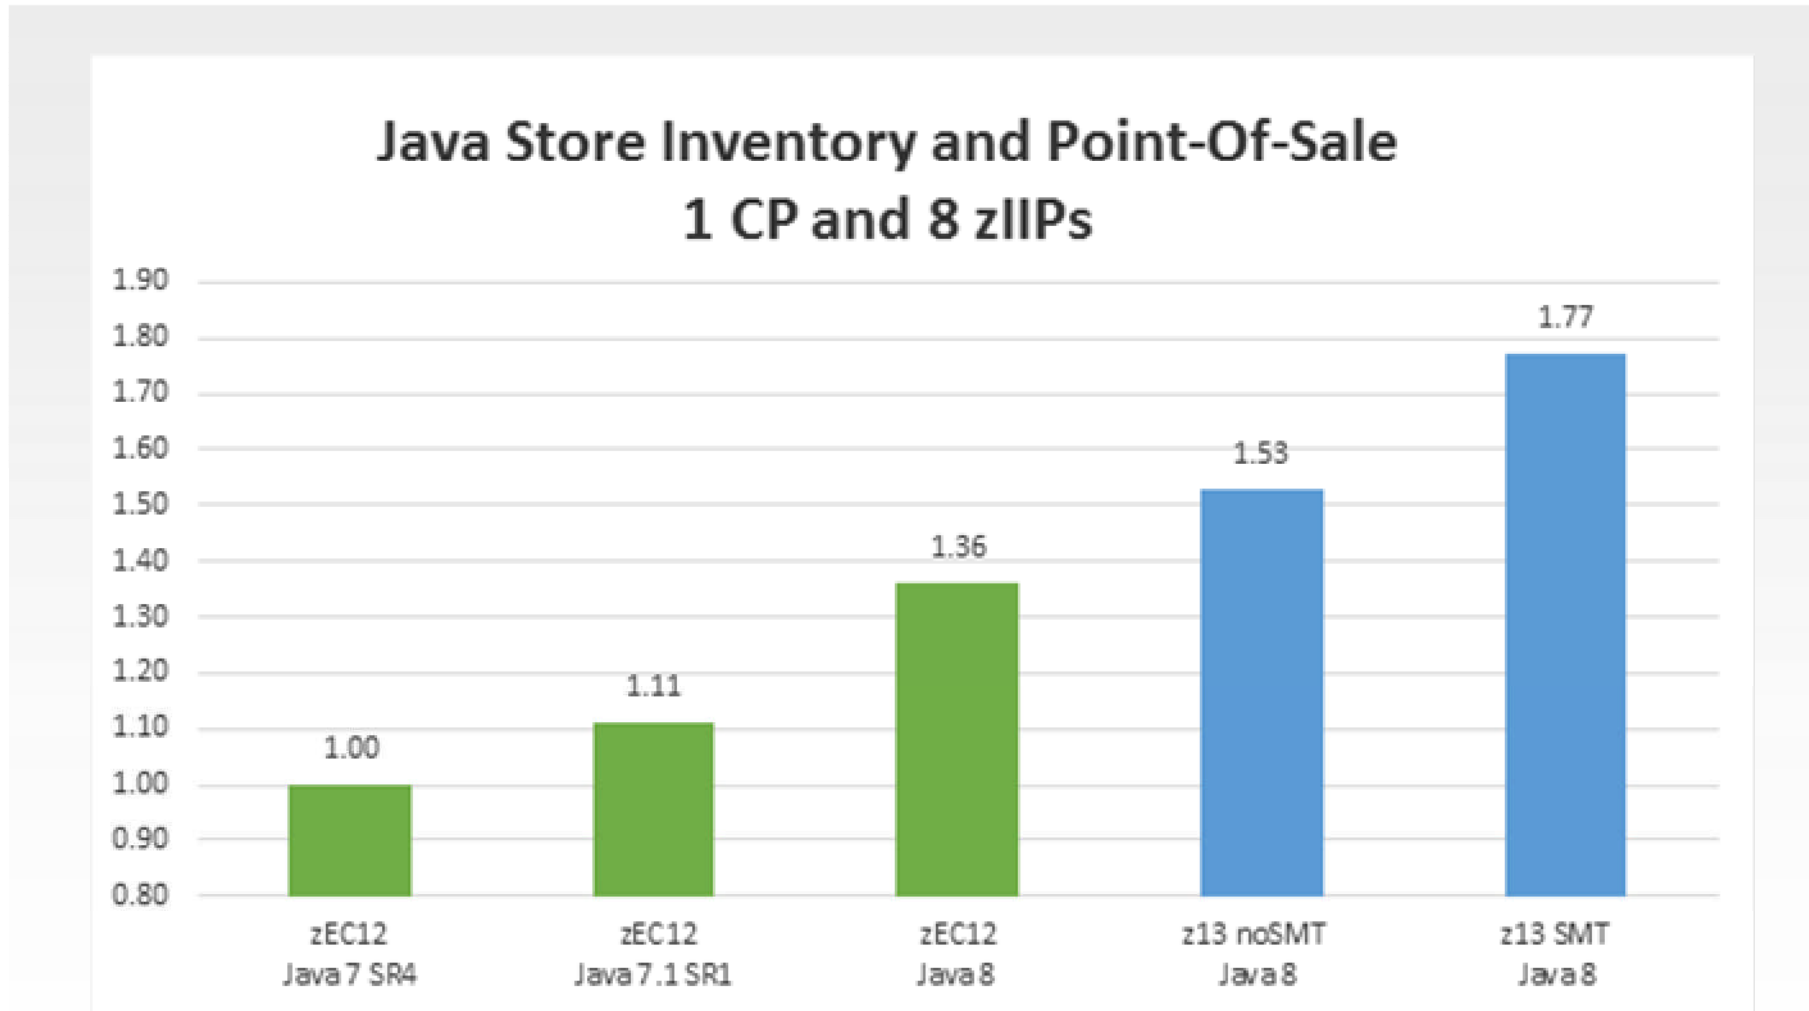


2.62x improvement in throughput with IBM Java 8 and IBM z13

(Controlled measurement environment, results may vary)



Java Store Inventory and Point-of-Sale App with IBM Java 8 and z13



1.77x improvement in throughput with IBM Java 8 and IBM z13

(Controlled measurement environment, results may vary)