



IBM zSystems Tech Bytes

Presented by the Washington Systems Center

z/OS Performance “HOT” Topics

Kathy Walsh

IBM Distinguished Engineer

October 13, 2022

Agenda

- IBM z16™ Processor Information
- IBM zPCR and zBNA Tools Information
- Performance and Capacity Planning Topics
- Interesting APARs
- Other Performance Recommendations
- Addendum
 - Older APARs or Performance Information



IBM z16

Changes to Techdocs

<https://www.ibm.com/support/pages/ibm-techdocs-technical-sales-library>



Techdocs

Welcome to Techdocs, the IBM Technical Sales Library.

Use the search bar below to locate documents within the Techdocs Library.

This site provides access to the Technical Sales Support organization's technical information database. It gives you access to the most current installation, planning and technical support information available from IBM technical sales support, and is constantly updated. These documents cover a wide range of material including TechNotes & Tips, White Papers, Flashes, Presentations and Tools.

Enter Keyword: Select Doctype: Select Date:

IBM z16 At a Glance

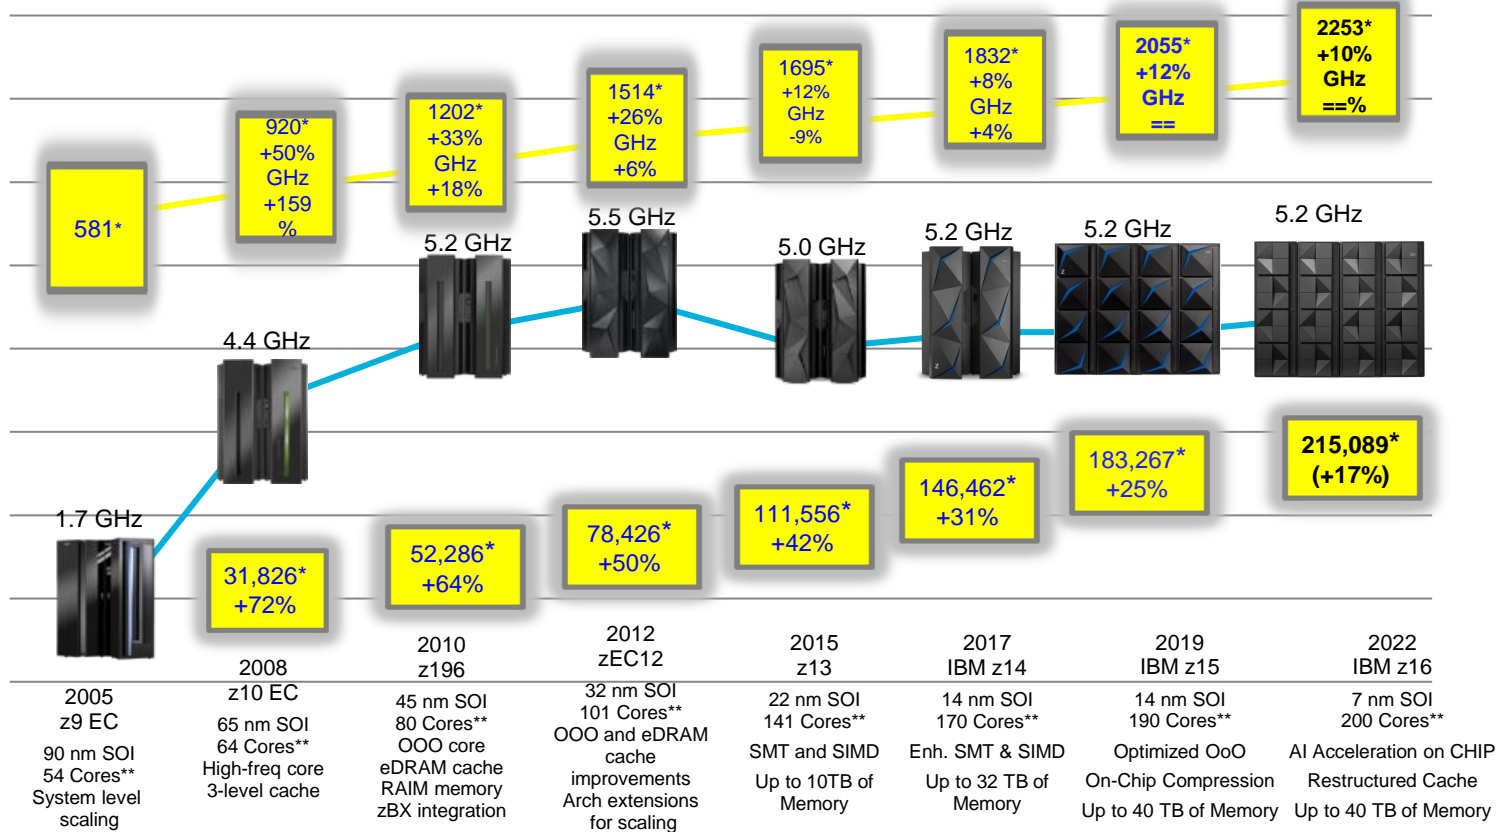
IBM z16

Model A01		
Feature	Customer PUs	Max Memory
Max200**	200	40 TB
Max168**	168	40 TB
Max125	125	30 TB
Max82	82	20 TB
Max39	39	10 TB

- Machine Type: **3931**
- One Model – **A01**
 - One, two, three, or four 19” Frames (A, B, C, IBM Z)
- Five features
 - Max39, Max82, Max125, Max168** and Max200**
- Processor Units (PUs)
 - 48 (57 for Max200) PU cores per CPC drawer
 - Up to 24 standard SAPs per system (up to 8 optional additional SAPs)
 - 2 spares designated per system
 - 85 LPARs
 - Sub-capacity available for up to 39 CPs
- Memory
 - 6 x 2 x 8 channel Reed - Solomon RAIM Memory design
 - System Minimum of 512 GB
 - Up to 10 TB per drawer
 - Up to 40 TB for System and up to 32 TB per LPAR (OS dependent)
 - Virtual Flash Memory
- I/O
 - Up to 48 PCIe+ Gen3 Fanouts -- 2 port @16GBps each and Integrated Coupling Adapters 2 port @ 8 GBps per system

** - NOTES: Max168 and Max200 are factory build only (no CPC Drawers field upgrades)

IBM z16 Continues the CMOS Mainframe Heritage



GHz / PCI*


* PCI Tables are NOT adequate for making comparisons of IBM Z processors. Additional capacity planning required
 ** Number of PU cores for customer use

IBM z16 Full and Sub-Capacity CP Offerings


- Sub-capacity CPs, up to 39 may be ordered (317 total capacity levels).
If more CPs are ordered, all must be full 7xx capacity.
- All CPs on an IBM z16 CPC must be the same capacity (except during Recovery Boost periods).
- All specialty engines are full capacity.
- zIIP to CP ratio – 2:1 and is the same for CPs of any capacity (except during System Recovery Boost periods).

CP Capacity – Relative to Full Capacity Uni

701 = 100% ≈ 2,253 PCI
601 ≈ 66% ≈ 1,496 PCI
501 ≈ 41% ≈ 937 PCI
401 ≈ 12% ≈ 280 PCI



MSU Sub Capacity



z16 Sub-Capacity Ratios

- The IBM z16 Sub-capacity models have different ratios than earlier models
- 6xx and 5xx models are closer in capacity to the 7xx model
 - More use cases for more/slower options with sub-capacity models

Capacity Ratios Relative to 7xx within family		
	z15	z16
401	0.13	0.12
501	0.38	0.42
601	0.56	0.66
701	1.00	1.00

Processor Unit (Core) Locations: Customer, SAP, IFP and Spare

IBM z16		1 st Drawer				2 nd Drawer				3 rd Drawer				4 th Drawer			
Feature	Cust PUs	Cust PUs	SAP	IFP	Spare	Cust PUs	SAP	IFP	Spare	Cust PUs	SAP	IFP	Spare	Cust PUs	SAP	IFP	Spare
Max200	200	47	6	2	2	51	6	0	0	51	6	0	0	51	6	0	0
Max168	168	39	5	2	2	43	5	0	0	43	5	0	0	43	5	0	0
Max125	125	39	5	2	2	43	5	0	0	43	5	0	0				
Max82	82	39	5	2	2	43	5	0	0								
Max39	39	39	5	2	2												

- PUs can be purchased as CPs, IFLs, Unassigned IFLs, zIIPs, ICFs or Additional SAPs
 - zIIP to CP purchase ratio is 2:1
 - Additional SAPs + Permanent SAPs may not exceed 32
 - Any un-configured PU can act as an additional Spare PU
 - CPs and zIIPs initial placement in 1st drawer working up
 - IFLs and ICFs initial placement in highest order drawer working down
- Upgrades available from any lower feature to any higher any models
 - Achieved via Concurrent Drawer Add from Max39 to Max125
 - Upgrade to Max168 and Max200 from any other feature not supported. Max168 and Max200 are factory built only.

z15 vs z16 Hardware and Topology Comparison

z15 CPU 5.2 GHz
 Caches L1 private 128k i, 128k d / core
 L2 private 4 MB i, 4 MB d / core
 L3 shared 256 MB / CP chip
 L4 shared 960 MB / drawer

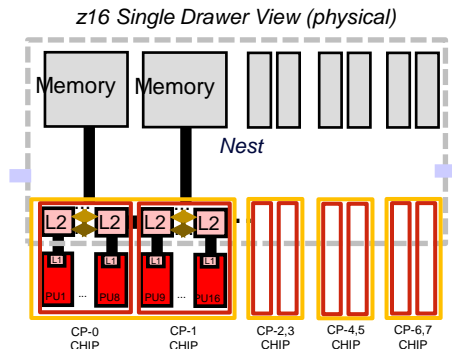
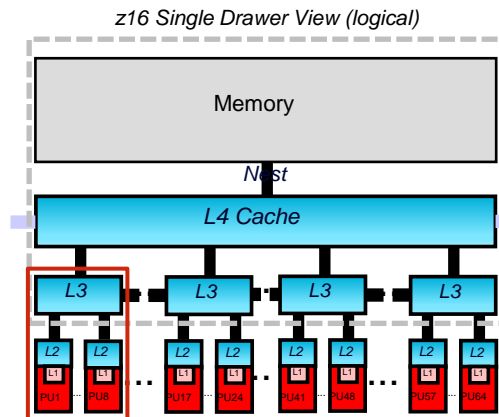
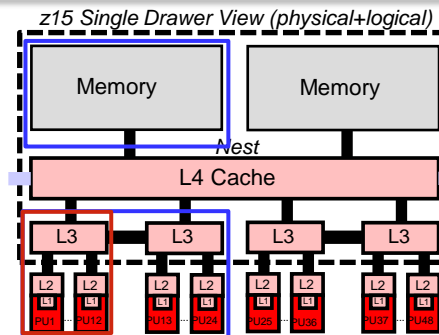
Topology

- 12 cores + 1 L3 / CP chip
- 2 CP chips / cluster
- 2 clusters + 1 L4 (48 engines) / drawer
- 5 drawers / CEC
- Book interconnect: numa star

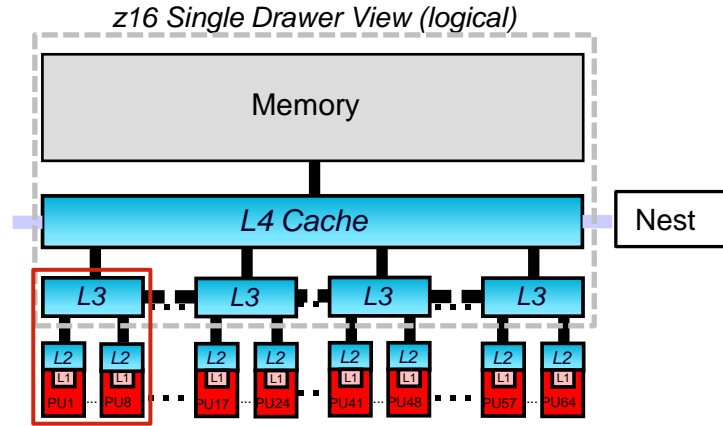
z16 CPU 5.2 GHz
 Caches L1 private 128k i, 128k d / core
 L2 private 32 MB unified / core
 virtual victim L3 up to 7x32 = 224 MB / CP chip
 virtual victim L4 up to 8x32x8 = 2 GB / drawer

Topology

- 8 (core + L3)s / CP chip
- 2 CP chips / DCM
- 4 DCMs (64 engines) / drawer
- 4 drawers / CEC
- Book interconnect: NUMA star



CPU MF Instrumentation – SMF 113 Records



SMF 113 records introduced with the IBM System z10 server on February 26, 2008.
The IBM Best Practice based on z10 lessons learned is to enable the SMF 113 records on all z/OS LPARs

CPU MF Store Into Instruction Stream (SIIS) Information

- CPU MF can be used to help identify potential SIIS timeframes
 - Based on % of certain I Writes / D Writes sourced
 - An LPAR view, identifies when it happens, not who is causing it
 - Identify the program(s) running in the time period, e.g., via zBNA Top Programs
 - Use a hot spot analyzer to find the code segments
 - Remediate the source code to correct the issue
- White paper - <http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102806>

Processor	SIIS Indicator %	Description
zEC12 / zBC12	E130 / B4 * 100%	D Writes sourced with L2 intervention / D Writes
z13 / z13s	E163 / B2 * 100%	I Writes sourced with L3 intervention / I Writes
z14 / ZR1	E164 / B2 * 100%	I Writes sourced with L3 intervention / I Writes
z15	E164 / B2 * 100%	I Writes sourced with L3 intervention / I Writes
z16	E170 / B2 * 100%	I Writes sourced with L2 intervention / I Writes



IBM z16 Integrated Accelerator for AI

Centralized On-chip accelerator shared by all cores



Very low and consistent inference latency



Compute capacity for utilization at scale




Variety of AI models ranging from traditional ML to RNNs and CNNs



Security – provide enterprise-grade memory virtualization and protection



Extensibility with future firmware and hardware updates



IBM Z Flexible Capacity for Cyber Resiliency

On demand, automatable transfer of capacity between IBM z16 machines



Greater Flexibility

Dynamically shift production capacity cross sites between IBM z16 machines

Flexibility and elasticity for DR test, planned maintenance, proactive outage avoidance and actual DR scenarios

Works in conjunction with other temporary record types, e.g., On/Off Capacity On Demand, Tailor Fit Pricing for Hardware, etc.

Complete Customer Control

Remotely transfer capacity – no on-site personnel (IBM or customer) required after initial set up

Flexible duration of capacity transfer, up to 1 year

Fully automated using solutions such as GDPS

Improved Compliance for Disaster Recovery

Simplify compliance and improve confidence DR scenarios including test

Closer mapping between test and production scenarios

No need for CBU records using this solution

System Recovery Boost Options – z16 Enhancements


1	Customer-Identified Middleware Region Start/Restart Boost Boost the z/OS system on which an STC middleware instance is being started/restarted
2	SVC Dump Boost Boost System on which an SVC Dump is being taken CHNGDUMP, DISPLAY DUMP,INFO and DISPLAY DUMP,OPTIONS enhancements
3	HyperSwap® Configuration Load Boost Boost load/re-load processing for making changes to z/OS HyperSwap® policy/configuration

- The new boost use cases share the existing recovery process boost time pool of 30 minutes per LPAR per day (in aggregate) with the earlier z15 use cases

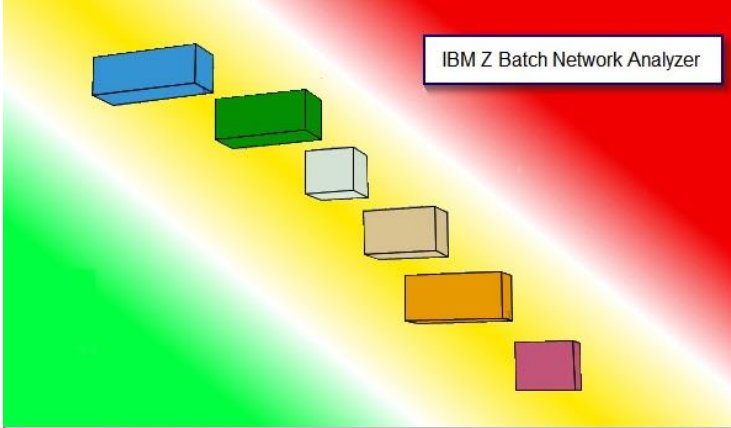

New Versions of zPCR and zBNA

Processor Capacity Reference
for
IBM Z and LinuxONE
zPCR version 9.5c
d3022084cb3784d2eeb0cc14801ea4eeea3f3832

(C) Copyright IBM Corp. 2003, 2022. All rights reserved.
Contains graphics software from KL Group
(C) Copyright KL Group 2000. All rights reserved
E-mail support: zpcr@us.ibm.com
IBM Employee web site: supportcontent.ibm.com/support/pages/node/6354531
IBM BP web site: www.ibm.com/partnerworld/techdocs/prs1796



zPCR



IBM Z Batch Network Analyzer

IBM Z Batch Network Analyzer
V2.4.2 - Build ZBN2207211556 Jul 21, 2022 3:56 PM

zPCR Large Partition Support

- There can be negative impacts on a partition's performance when CPs are in different drawers
 - PR/SM will attempt to put all VH and VM CPs for a partition on the same drawer
- zPCR will warn you when LCPs defined are larger than a drawer

z15 Host = 8561-T01(Max145)/700 with 64 CPs: GP=50 zIIP=10 IFL=3 ICF=1
4 Active Partitions: GP=1 zIIP=1 IFL=1 ICF=1

Include	Partition Identification					Partition Configuration				Capping	
	No.	Type	Name	SCP	Workload	Mode	LCPs	Weight	Weight %	✓	ABS
<input type="checkbox"/>		GP	LPAR 1	z/OS-2.4*	Average	SHR	44	800		<input type="checkbox"/>	
<input type="checkbox"/>		zIIP	LPAR 1	z/OS-2.4*	Average	SHR	6	800		<input type="checkbox"/>	
<input checked="" type="checkbox"/>	1	GP	LPAR 2	z/OS-2.4*	Average	SHR	12	200	100.00%	<input type="checkbox"/>	
<input checked="" type="checkbox"/>		zIIP	LPAR 2	z/OS-2.4*	Average	SHR	2	200	100.00%	<input type="checkbox"/>	

zPCR - Notice ×

Large Partitions Exceeding a Drawer Boundary

On IBM z16 (3931) models, the best practice is a partition's logical CP count should not exceed the number of RCPs in the largest drawer.

Partitions which exceed a drawer boundary have special capacity considerations.

This partition is likely to exceed a drawer boundary.

Immediately contact IBM to review your configuration: zPCR@us.ibm.com

Check to allow inclusion of this partition

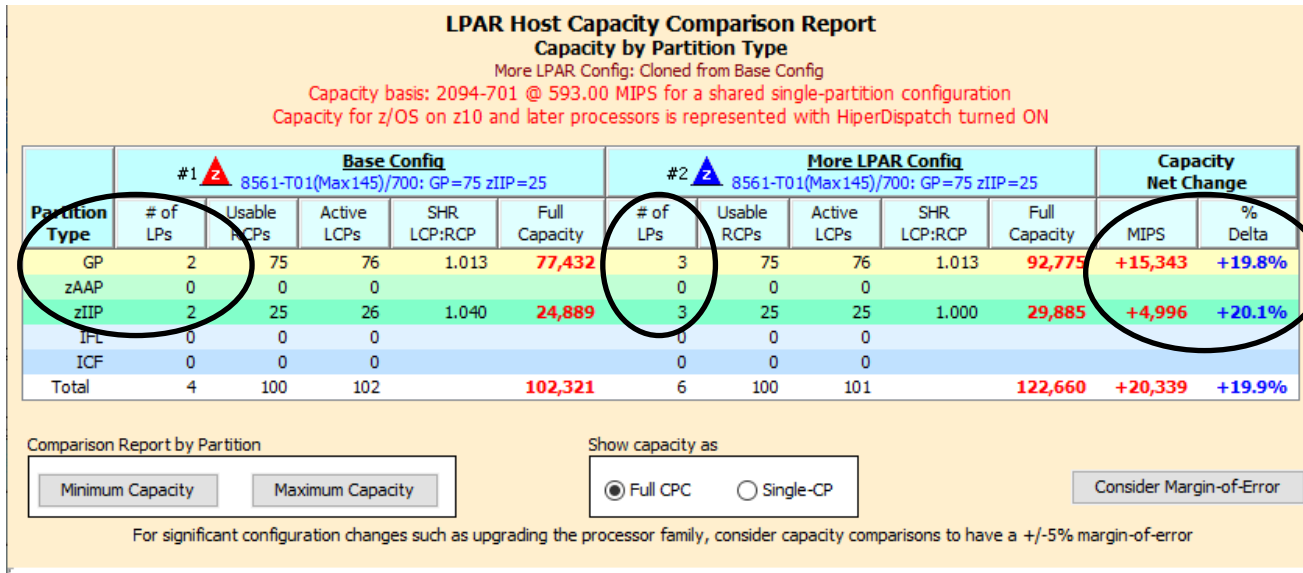
OK

Value of Adding One LPAR

Configuration change modeled:

1. Running two 38-way LPARs with 9 zIIPs changed to
2. Running 2 LPARs with 28-way LPARs with 8 zIIPs and 1 20-way LPAR with 9 zIIPs

Value of Parallel Sysplex with True Data Sharing



zPCR HiperDispatch Window

- Dynamically shows VH/VM/VL CPs for each partition based on weights of all partitions
- Small partitions with share of < 1 CP but with 2 or more CPs will have 1 VM and 1 VL never parked, noted in the tool

Partition Detail Report

Based on LSPR Data for IBM Z Processors
Study ID: Not specified
#1 Configuration #1

Description: Created from EDF VIDTRAIN.MASTER.Z301.EDF for CPC99FFF interval #6
z15 Host = 8561-T01(Max145)/700 with 41 CPs: GP=25 zIIP=10 IFL=1 ICF=5
23 Active Partitions: GP=11 zIIP=10 IFL=1 ICF=5
Capacity basis: 2094-701 @ 1.000 ITRR for a shared single-par
Capacity for z/OS on z10 and later processors is represented with H

Partition Identification							
Include	No.	Type	Name	SCP	Assigned Workload	Mode	Active LCPs
<input checked="" type="checkbox"/>	1	GP	WSCS001	z/OS-2.4*	Average	SHR	4
<input checked="" type="checkbox"/>	2	GP	WSCC001	z/OS-2.4*	Average	SHR	2
<input checked="" type="checkbox"/>	3	GP	WSCZ001	z/OS-2.4*	Average	SHR	1
<input checked="" type="checkbox"/>	4	GP	WSCX001	z/OS-2.4*	Average	SHR	12
<input checked="" type="checkbox"/>	5	GP	WSCW005	z/OS-2.4*	Average	SHR	13
<input checked="" type="checkbox"/>	6	GP	WSCG001	z/OS-2.4*	Average	SHR	3
<input checked="" type="checkbox"/>	7	GP	WSTC003	z/OS-2.4*	Average	SHR	2

Table View Controls

Display zAAP/zIIP/IFL/ICF Associated Partitions

With Parent GP Separate by Pool

Show: GP Pool GP zAAP zIIP

All Partitions Includes Only IFL ICF

Capacity Summary by Pool

CP Pool	Real CPs	LPs	DED LCPs
GP	25	11	
zIIP	10	10	26
IFL	1	1	1
ICF	5	1	5
Totals	41	23	5

HiperDispatch Assignment for Shared Partitions

Study ID: Not specified
#1 Configuration #1

Description: Created from EDF VIDTRAIN.MASTER.X001.EDF for CPC99FFF interval #61: Date=2021-09-09 Time=11:00:00
z15 Host = 8561-T01(Max145)/700 with 41 CPs: GP=25 zIIP=10 IFL=1 ICF=5
23 Active Partitions: GP=11 zIIP=10 IFL=1 ICF=5


Partition Identification										HiperDispatch Configuration					
Include	No.	Type	Name	SCP	Assigned Workload	Mode	LCPs	Weight Percent	Engines by Weight	VHs	VMs	VM %	VLs	Nvr Pk	VL Nvr Pk %
<input checked="" type="checkbox"/>	1	GP	WSCX001	z/OS-2.4*	Average	SHR	11	38.64%	9.66	9	1	66%	1	0	0%
<input checked="" type="checkbox"/>	2	GP	WSCS001	z/OS-2.4*	Average	SHR	6	9.09%	2.27	1	2	64%	3	0	0%
<input checked="" type="checkbox"/>	3	GP	WSCC001	z/OS-2.4*	Average	SHR	2	4.55%	1.14	0	2	57%	0	0	0%
<input checked="" type="checkbox"/>	4	GP	WSCZ001	z/OS-2.4*	Average	SHR	2	4.55%	1.14	0	2	57%	0	0	0%
<input checked="" type="checkbox"/>	5	GP	WSCW005	z/OS-2.4*	Average	SHR	13	32.95%	8.24	7	2	62%	4	0	0%
<input checked="" type="checkbox"/>	6	GP	WSCG001	z/OS-2.4*	Average	SHR	3	4.55%	1.14	0	2	57%	1	0	0%
<input checked="" type="checkbox"/>	7	GP	WSTC003	z/OS-2.4*	Average	SHR	2	1.14%	0.28	0	1	14%	0	1	14%
<input checked="" type="checkbox"/>	8	GP	WSCICP01	z/OS-2.4*	Average	SHR	1	1.14%	0.28	0	1	28%	0	0	0%

Host Summary SMT Benefit LCP Alternatives zAAP/zIIP Load **HiperDispatch** Curate Capacity

For significant configuration changes such as upgrading the processor family, consider capacity comparisons to have a +/-5% margin-of-error

zBNA New IBM z16 Support and Recent Enhancements

- Ability to trim large .dat files for easier analysis
- New filtering support
 - Import a list of Job or DSN Masks from a text file
 - New cross-application filters panel
- New Data Set I/O performance by hour support (SMF 42.6)
- New DFSort Z Sort reports
 - Improved identification of sorts eligible for IBM Z Sort acceleration
- Updates to zHyperLink™ Db2 log write candidate algorithm



zIIP Tuning – Measurable IIPCP time with Low Utilization

- Typical causes:

- Spiky arrival rates
- Not aggressively unparking vertical lows

- First Step:

- Review zIIP weights:

- Increase the relative weight leaving smaller LPARs with less weight, since small LPARs with 2 logicals always have 2 unparked zIIPs
- Change the relative weight slightly in order to get 2 Vertical Mediums vs 1 Vertical Medium if it nets up always unparked logical (1.95 CPs gives 1VH, 1VM, but 2.01 CPs gives 1 VH, 2VM)

- Second Step:

- Validate the IIPCP time can run on a zIIP

- Verify the IIPCP time is not due to Suspend Lock Serialization issues by running with IIPHONORPRIORTY=NO for a short period of time
- If the IIPCP time remain then it is due to lock negotiation and tuning ZIIPAWMT will not help

zIIP Tuning – Measurable IIPCP time with Low Utilization

■ Third Step:

– Adjust IEAOPTxx parm ZIIPAWMT

- Adjust ZIIPAWMT to force the zIIP engine to wait longer before asking for help
- May require iterative tuning – go slow
 - Change to 6400 – double the HD=YES interval
 - Increase to the same value used for HD=NO or 12,000
 - Double the value to 24,000
 - Last increase to 36,000 can be done, but it is recommended to NOT go above the 36,000 value

– Validate over time

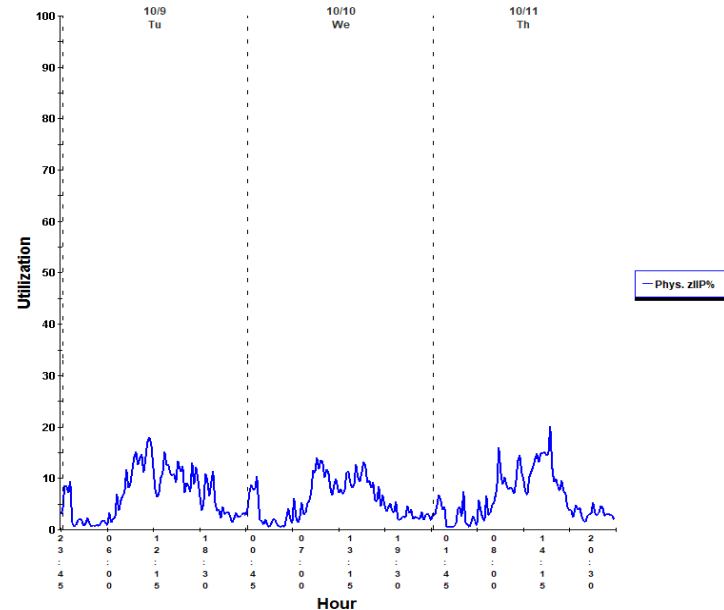

- The "needs help" support will be slower which is more important the busier the zIIP pool becomes, so don't let your zIIP pool get on average above 80% busy with increased ZIIPAWMT settings

– Review again if the LPAR's zIIP work grows to 4 zIIP engines or more by weight

- The LPAR will have at least two WLM Node Affinity Tables (NAT) and the high zIIPAWMT value may impact the ability of the zIIPs in one NAT to ask the cores in the other zIIP NAT to help
- If ≥ 4 zIIP engines by weight per LPAR, **recommend** use default ZIIPAWMT 3200 value, if IIPCP continues raise value to no more than 12000

zIIP SMT Concerns at Low or High Utilizations

- SMT Capacity Factors may show unrealistic values
 - Sampling used to calculate CF and MAX CF values can be impacted by zIIP utilization
 - zIIP busy less than 10% or greater than 90% can record odd values for Max Capacity Factor which impacts CPU service times and capture ratios



zIIP Service Time and SMT MT=2

```

INTERVAL 000.59.59
SERVICE TIME ---APPL %---
CPU 30524.27 CP      786.47
SRB  959.181 IIPCP   2.12
RCT   0.000 IIP     68.40
IIT  43.715 AAPCP   0.00
HST   0.000 AAP      N/A
IIP 3214.679
AAP   N/A
  
```

RMF CPU Report from same interval

SYSTEM ID WSC1				
----- MULTI-THREADING ANALYSIS -----				
CPU TYPE	MODE	MAX CF	CF	AVG TD
CP	1	1.000	1.000	1.000
<u>IIP</u>	2	1.306	1.142	1.463

$68.4 / 100 * \text{interval } 3599 = 2461.716 \text{ seconds}$
 $2461.716 \text{ seconds} * \text{Max CF } 1.306 = 3215.001 \text{ zIIP seconds}$

Values may change every interval

z/OS Correlator SMF Records



- SMF 98.1
 - Functions:
 - Defined in SMFPRMxx via new keyword HFTSINTVL(ss) | NOHFTSINTVL
 - Write interval based on seconds (5, 10, 15, 30, 60) and the recommended interval is 5 seconds
 - Supervisor Exploitation
 - Type 98 subtype 1 records contain performance information for the z/OS supervisor component about the workload and its significant jobs. It includes metrics such as utilization, concurrency, efficiency, contention, and queuing
 - Subsystem Exploitation
 - Db2
 - CICS, IMS - Requires Correlator
 - White Paper: <https://www.ibm.com/support/pages/node/6437547>
 - **APAR OA62268** – IBM z/OS Workload Interaction Correlator (WIC) Entitlement Enhancement
 - z/OS V2R4 with RMF or V2R5 with Advanced Data Gatherer (which is entitled when RMF is licensed) are entitled to the IBM z/OS WIC at no additional charge

APAR PH34378 → OA61811, OA62502, PH41024

■ PH34378

- With high performance DBATs bound with RELEASE DALLOCATE
- WLM enclave created for single transaction will be retained for potentially up to **200 transactions**
- Ended transaction for WLM has been the ending of the enclave, not the transactions in the enclave
 - With PH34378, reported response times can vary greatly from one enclave to the next
 - Service classes with these types of high performance DBATs need to be converted from a response time goal to a velocity goal

APARs OA61811, OA62502, PH41024

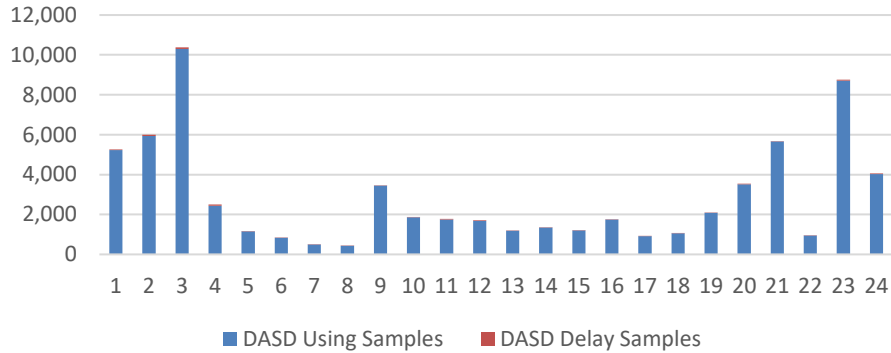
- OA61811 – WLM Enhancement
 - WLM provided extension to IWM4EDEL (Delete an Enclave) to provide the actual number of transactions and aggregated response time of the transactions that were processed during the use of the enclave
- PH41024 – DB2 12 Support for WLM APAR OA61811
 - Gives DB2 the ability to provide the additional data required for OA61811
- OA62502 – z/OS Data Gatherer Support
 - Additional data collected in support of OA61811 will now be collected in RMF 72 subtype 3 and RMF MON III table ERBRCDG3
 - New fields in subtype 3 records will contain actual transaction counts and aggregated response time information
 - RMF Postprocessor report is not updated
 - Does not change requirement to move to velocity goal

Workload Manager I/O Priority Queuing



- I/O typically has the same priority as CPU dispatch priority
 - Setting parameter to YES allows for I/O to have higher priority if workload missing goal due to I/O delays, mostly IOSQ
 - Changes WLM velocity calculation by adding I/O using and delay samples
 - Leads to inflated achieved velocities: workload may have CPU delay but still achieving velocity goal due to dominance of I/O Using samples and few or no I/O delay samples
 - Result is very high velocity goals are needed in policy and real CPU delay detection is slower
- **Recommendation: In WLM Policy set I/O PRIORITY MANAGEMENT: NO**
 - Allows WLM to be more responsive to CPU delays

Workload Manager I/O Priority Queuing – Importance 1 Service Class



DASD Using and Delay Samples



Using Samples




Workload Manager I/O Priority Queuing – Importance 5 Service Class



Capture Ratio and First Reference Page Faults

- First Reference Page Faults happen when a program obtains virtual storage and then first touches the page
 - Page is not backed by a real storage frame, so a page fault is taken to drive Dynamic Address Translation
 - This is uncaptured CPU time
- Values above 100,000 per second should be considered problematic
- Primary cause of low capture ratio seen in multiple sites
- Other causes of low capture ratios include
 - Fragmented storage pools (most common)
 - Can use SMF30ICU to identify in zBNA
 - High page fault rates
 - Suspend /Spin lock contention
 - Long internal queues



Addendum

- Older information which should still be understood, or make you go Hmmm.
- APARs which are still interesting, even though they are older.





- Session Evaluation link is provided in the Chat for this session.
- Please fill out a session evaluation as it does help us greatly!
- Survey: <https://ibm.biz/ztechbytes-09-wsc>



Notices and disclaimers

- © 2022 International Business Machines Corporation. No part of this document may be reproduced or transmitted in any form without written permission from IBM.
- **U.S. Government Users Restricted Rights — use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM.**
- Information in these presentations (including information relating to products that have not yet been announced by IBM) has been reviewed for accuracy as of the date of initial publication and could include unintentional technical or typographical errors. IBM shall have no responsibility to update this information. **This document is distributed “as is” without any warranty, either express or implied. In no event, shall IBM be liable for any damage arising from the use of this information, including but not limited to, loss of data, business interruption, loss of profit or loss of opportunity.** IBM products and services are warranted per the terms and conditions of the agreements under which they are provided.
- IBM products are manufactured from new parts or new and used parts. In some cases, a product may not be new and may have been previously installed. Regardless, our warranty terms apply.”
- **Any statements regarding IBM's future direction, intent or product plans are subject to change or withdrawal without notice.**
- Performance data contained herein was generally obtained in a controlled, isolated environments. Customer examples are presented as illustrations of how those
- customers have used IBM products and the results they may have achieved. Actual performance, cost, savings or other results in other operating environments may vary.
- References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business.
- Workshops, sessions and associated materials may have been prepared by independent session speakers, and do not necessarily reflect the views of IBM. All materials and discussions are provided for informational purposes only, and are neither intended to, nor shall constitute legal or other guidance or advice to any individual participant or their specific situation.
- It is the customer’s responsibility to insure its own compliance with legal requirements and to obtain advice of competent legal counsel as to the identification and interpretation of any relevant laws and regulatory requirements that may affect the customer’s business and any actions the customer may need to take to comply with such laws. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the customer follows any law.

Notices and disclaimers

- Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products about this publication and cannot confirm the accuracy of 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. IBM does not warrant the quality of any third-party products, or the ability of any such third-party products to interoperate with IBM's products. **IBM expressly disclaims all warranties, expressed or implied, including but not limited to, the implied warranties of merchantability and fitness for a purpose.**
- The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents, copyrights, trademarks or other intellectual property right.
- IBM, the IBM logo, ibm.com and [names of other referenced IBM products and services used in the presentation] are trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at: www.ibm.com/legal/copytrade.shtml

Various Performance and Interesting APARs

- OA62937 – RMF Post Processor CHAN Activity Report Incorrectly Reports Channels at 100% Utilization
 - Problem can occur since there is no serialization between the channel path cycle gatherer and interval gatherer
 - With APAR channel path cycle gatherer module is changed to skip a cycle if interval gathering is running concurrently
- OA63167 – Not all processors will be enabled for I/O when switching from HD=YES to HD=NO and with CPENABLE(0,0)
 - With CPENABLE set to 0,0 all online processors should be enabled for I/O, but parked vertical low processors will not be enabled for I/O when switching from HiperDispatch YES to NO
 - No PTF supplied, circumvention is to ensure there are no parked vertical low CPs, either through temporary weight change or varying off vertical low CPs

Other APARs

- OA62145 – RMF Reporting 100% MVS Busy but MON III delay has zero CPU Delays
 - Additional symptom is higher-than-normal CPU consumed by WLM address space
 - Occurs at the end of System Recover Boost zIIP Boost period on z15
- OA62171 – Problem described by APAR OA62145 exposes a secondary problem which can result in a loop
 - Only occurs on z15 where a system boost has occurred
 - With fix for OA62145 problem is likely to be short lived which will not be externally detectable, but can increase likelihood of program check in IEAVEWDI and resulting in a dump

New Data Set Types Support: Basic and Large Format



- ***Basic and large format encryption (non-extended format DASD data sets)***
 - **z/OS V2.3 and up, SMS-managed only**
 - **Access using BSAM and QSAM APIs**
 - Transparent to application except for **DASD space calculations** (due to new 8-byte block prefix for encrypted basic and large data sets)
 - **Access using EXCP, then application changes are required**
 - The EXCP program must account for an 8-byte prefix on each block
 - The EXCP program must encrypt the data before writing and decrypt after reading
 - IBM will provide new macro for encryption/decryption so result will be compatible with the access methods

Modified physical format for encrypted basic and large format data sets: block prefix

System Determined Block size (SDB) considers the prefix length when deriving the optimal user block size

BEST PRACTICE: Use encrypted sequential extended format data sets

Various Performance APARs

▪ **OA60426: JES2 Delay due to Contention on JES2 Checkpoint lock**

- JES2Clxx and JES2EDS address spaces are assigned a service class of SYSTEM and dispatching priority of x'FF' due to HIPRI attribute being specified
- JES2 is defined at Service Class SYSSTC x'FE' so runs at lower priority
 - Most impact is felt when the LPAR has low n-way and is part of a JES2 MAS

▪ **OA60377: POOR VSAM RLS PERFORMANCE DUE TO HOLDING LOCAL LOCKS TOO OFTEN**

- Each batch RLS Rec Mgmt request for a task calls MMF which invokes WLM with an option which causes obtains / releases of the local lock
- Change is made to obtain the local lock ONLY for the first batch task-mode Rec Mgmt request
- SMFIOCNT processing to count I/O blocks for SMF30 recording is changed to be done once every 2000 RLS Rec Mgmt requests, not every 20 requests

SMF Updates

■ APAR OA56924 – IOS Support

- z/OS support to update display commands and provide SMF records in support of fibre channel endpoint security
- Update SMF 124 subtype 1 with link diagnostic and four new subtypes
 - Subtype 2 - Endpoint security information
 - Subtype 3 - Endpoint security authentication key update
 - Subtype 4 - Endpoint security encryption key update
 - Subtype 5 - External key manager event record

■ APAR OA59126 – SMF 30 records

- Added fields to track user created data spaces
- High water mark of number of in-use data spaces created in problem state, and user key during the job step
- Total number of data spaces created in problem state, and user key during the job step

WLM Availability Recommendation

- Need to maintain a backup copy of the WLM Service Definition
 - Outside of the Couplexx data sets
 - Use the WLM ISPF Administrative Application
 - OA60011: IWMAM079 SAVING WLM POLICY AS XML

```
File Utilities Notes Options Help
-----
Functionality LEVEL032      Definition Menu      WLM Appl LEVEL035
Command ==> _____

Definition data set . . . : none

Definition name . . . . . WLM      (Required)
Description . . . . . WSC Default Service Definition

Select one of the following options.
__ 1. Policies                12. Tenant Resource Groups

Save to...

Data set name . . _____
Save format . . . 1  1. XML format (Dataset must be sequential)
                   2. ISPF tables (Dataset must be partitioned)

F1=Help  F2=Split  F5=KeysHelp  F9=Swap  F12=Cancel
```