



Cheryl's Hot Flashes #19

Cheryl Watson
February 29, 2008, Session 2509
Watson & Walker, Inc.
www.watsonwalker.com

home of **Cheryl Watson's TUNING Letter, CPU Chart,
BoxScore, and GoalTender**



Thank you, Walt Caprice!

- In memoriam – teacher, mentor, friend to all...we will miss him.
- Journal, Guestbook and Pictures
 - <http://www.caringbridge.org/visit/waltcaprice>
- In lieu of flowers the family is requesting donations to be sent to:

Delaware Humane Association
701 A Street
Wilmington, DE 19801
<http://www.dehumane.org>



Agenda

- Survey Questions
- Important Fixes
- New Function Maintenance
- zPCR Analysis
- Some Goodies!
- Last SHARE
- EWCP Requirements
- This SHARE

Survey Questions – Hardware

- Current Server Type (now or within next 12 months)
 - z800, z900, z890 (23% at last SHARE)?
 - z990 (23%)?
 - z9-BC (24%)?
 - z9-EC (38%)?
 - Z10?
 - Older Hardware (1%)?
- Using zAAP Processors (32)?
- Using zIIP Processors (40)?
- Have Used On/Off Capacity on Demand (21)?
- Using Variable WLC Pricing (31)?
- Doing Heavy Cryptographic Work (9)?
- Who is exploiting the full 65,520 cylinders of a 3390-54?

Survey Questions – Software

- Operating System
 - z/OS 1.7 (80%)? (EOS is September 2008)
 - z/OS 1.8 (23%)?
 - z/OS 1.9 (0)?
 - Older (22%)?
- Using WebSphere on z/OS (26%)
- Using VSAM RLS for CICS (16)?
- Have used zPCR (15)?
- Running WLM-managed initiators (25%)?
- Using or planning to use HyperPAV (19%)?
- Using or planning to use zLinux?
- Using z/OS 1.9 SMF Logger?

Important Fixes

WLM APARs



- **OA22607** – Increased SMF30CPT for WLM A.S.
 - High CPU in address space, especially when running many DB2 subsystems and many CICS regions with high maxtask values
 - Applies to z/OS 1.8; also to z/OS 1.7 with APAR OA12935 applied
 - Not marked HIPER, but WLM was using 3-5% of a processor
 - Thanks to **Alex McFarlane** of Technology Services Group (subsidiary of The Bank of New York Mellon)

- **OA23230** – Overlay of Common Storage
 - SMF99 subtype 11 buffer overlaid when WLM Group Capacity is defined and a new LPAR is dynamically allocated to the CEC
 - Can cause abends and even IPL, but not marked HIPER
 - Closed on 21Feb2008 for z/OS 1.8 and 1.9
 - Option – turn off Group Capacity or don't add LPAR

Important Fixes

RMF APARs



- **OA22958** – Low SSCHRT in WKLD Activity Report
 - If PEND time and control unit queue times are zero, then the I/O is not counted, so resulting measurements could show abnormally low SSCHRT counts and high response times for DASD devices
 - Applies to z/OS 1.7 to 1.9 (closed 19Dec2007)
 - Always been this way, but faster systems are increasing the likelihood of this occurring
- **OA23195** – Missing Samples or No Data in RMF Mon III
 - Delay in gathering XCF data causes missing data in RMF Monitor III
 - Applies to z/OS 1.9 (closed 7Dec2007)
 - Application of PTF improves performance of XCF data gathering

Thanks to **Jerry Urbaniak** of Acxiom

Important Fixes

HIPER APARs



- **OA23376** – Missing Records Using QSAM PUT to EF DSN
 - The last couple of blocks before a CLOSE are not getting written out when using QSAM PUT writing to an extended format dataset
 - Applies to z/OS 1.6 to 1.9 (closed 17Jan2008)
 - Occurs after applying fix for **OA16854**
 - Thanks to **Jerry Urbaniak** of Acxiom
- **OA23592** – SMF30UCT & SMF89UCT Contain Invalid Values!
 - Outrageously high (e.g. 25 times the total installed MSUs on the floor) MSU values for DB2
 - Applies only to z/OS 1.8 (closed 17Jan2008)
 - Read APAR for instructions on contacting IBM about measured usage charges (MULC).
 - Thanks to **Brian Peterson** of Blue Cross Blue Shield of Minnesota

New Function APARs

- **OA19421** (z/OS 1.6-1.8, 1Apr2007) - *Increased CPU, Elapsed Time, and CF Accesses Using RLS 64 Bit Buffering.* Accessing data sets using 64-bit buffering for VSAM RLS may increase CPU usage, elapsed time, and Coupling Facility (CF) and DASD accesses. Advertised as new function, but appears to be a performance APAR.
- **OA19439** (DFSMS for z/OS 1.7+, 27Mar2007) - *New Function.* As an HFS grows in size, it can become increasingly fragmented, causing performance problems. This APAR allows you to determine if fragmentation is occurring.
- **OA20077** (DFSMS for z/OS 1.7+, 19Mar2007) - *New Function - Add Two Fields to SMF Type 21 Record - Device Read and Write Bytes Processed.* This new function adds two counters to the SMF Type 21 demount record - SMF21DBR (device bytes read) and SMF21DBW (device bytes written).

New Function APARs

- **OA19194** (z/OS 1.7-1.8, 14Sep2007) - *New Function - AMATERSE, Alias TRSMAIN, Is a Supported Service Aid Function.* This maintenance provides a new service aid program that may be used to compress data being sent to IBM support for diagnostic purposes.
- **OA21635** (z/OS 1.6+, 20Sep2007) - *New Function - Recalibrate SYNC/ASYNC Thresholds.* A heuristic algorithm is used by z/OS to optimize host processor utilization by intentionally converting coupling facility requests from synchronous to asynchronous execution when past performance shows that requests are taking a long time to complete. Specifically, the conversion thresholds associated with both simplex lock and simplex list and cache requests have been increased.
- **PK37032** (Communications Server for z/OS 1.8+, 31Jul2007) - *Support For FTP Batch Pipes.* This maintenance adds new function support in the FTP server for Batch Pipe subsystems. Batch Pipes offers a way to connect jobs so that data from one job can move through processor storage to another job without going to DASD or tape. It may allow two or more jobs that formerly ran serially to run concurrently.

New Function APARs

- **OA15791** (DFSMSdss 1K0+, 6Dec2007) - *New Function Space Efficient FlashCopy SEFLC*. This new function allows a customer to perform a FlashCopy to a target that is a space efficient volume, which is intended to use only the amount of storage needed by the copy. Physical space for space efficient volumes is allocated on a track by track basis from a repository in the storage subsystem.
- **OA21126/OA22927/OA21124** (z/OS 1.7-1.9, 7Dec2007) - *New Function - XRC Multiple Readers*. This provides support for the XRC Multiple Readers function of D/T2107. It offers XRC session performance improvements by enabling parallel offload of recordsets from primary storage subsystem cache sidefiles. It also provides the capability for system data mover reader tasks to use parallel access volume aliases of base XRC utility devices.
- **OA22255** (z/OS 1.7-1.9, 27Dec2007) - *In Support Of New Function*. Users of some system logger health checks see exception conditions for log streams after they have been potentially addressed. The health checks report on the most recent 16 conditions since system logger started. The only way to clear these conditions is to restart system logger. The system logger health checks are modified to give installations a method to stop reporting exceptions for log streams conditions that have been addressed without having to restart system logger.

zPCR Analysis

- zPCR is provided at no charge to customers from the WSC – www.ibm.com/support/techdocs/atmastr.nsf/Webindex/PRS1381 (or search for 'zPCR Download').
- It provides an analysis of expected usage when hardware is changed, the LPAR configuration is changed, or the workload mix is changed.
- Every installation should be using this – **EVERY ONE!**
- Hardest problem is to determine what type of workload you have, although most people fall into the LOIO workload.
- **APAR OA22424** (z/OS 1.8-1.9) expected (31Mar2008) which writes workload characteristics (e.g. I/O density) in the SMF Type 23 (Statistics) record. After you apply APAR, please contact **Gary King** to get data collection program (see session #2532 for more info).

zPCR Analysis

- In our latest newsletter, we just described some interesting analysis on adding zIIPs and zAAPs
- Conclusion - Simply adding a zIIP or zAAP can cause some overhead on your general CPs
 - Adding a zIIP to a 701 can lose 6% of CP
 - Adding a zIIP to a 708 loses about 1.6% of the other CPs
 - The point – use the zIIP or zAAP to remove work from the other CPs so that you won't notice the slight overhead incurred

Table 2 – Switching Effect of zIIP Processors

Processor	Relative Capacity Per General Processor								
	0	1	2	3	4	5	6	7	8
2094-701	1.274	1.211							
2094-702	1.252	1.209	1.165						
2094-703	1.230	1.197	1.159	1.122					
2094-704	1.209	1.181	1.149	1.115	1.080				
2094-705	1.188	1.165	1.136	1.105	1.074	1.043			
2094-706	1.167	1.157	1.122	1.094	1.065	1.036	1.008		
2094-707	1.148	1.130	1.107	1.081	1.055	1.028	1.002	0.975	
2094-708	1.128	1.112	1.091	1.068	1.044	1.019	0.994	0.969	0.944

Some Goodies!

- Publications (Ways to Access)
 - z/OS Library
<http://www-03.ibm.com/systems/z/os/zos/bkserv/>
 - IBMLink Publications – (my favorite)
 - <http://www.elink.ibm.com/publications/servlet/pbi.wss?CTY=US>
 - IBM Basic Skills Information Center –
<http://publib.boulder.ibm.com/infocenter/zoslnctr/v1r7/index.jsp>
 - Installation –
<http://www-03.ibm.com/systems/z/os/zos/installation/>
 - Google! (add 'site:ibm.com')

Some Goodies!

- CMG Publications

- Proceedings contain full peer-reviewed technical papers on capacity planning, performance, systems management, and many other topics
- CMG has just announced public access to conference proceedings from 1997 to 2005
- Registration is required, but membership isn't required
- Go to www.cmg.org/proceedings for these public documents

- Sign up for MeasureIT – another publication for members and non-members
- Latest issue contains part 3 of an excellent series on SAS 9 Tips by **Phil Mason**
- Go to www.cmg.org and select 'Publications', then 'MeasureIT'

Some Goodies!

- White Papers on WSC Techdocs
 - **WP101157** – z/OS Large Memory Performance Studies (11Dec2007)
 - **WP101130** - z/OS System Management Facilities (SMF) Recording with MVS Logger (12Nov2007)
 - **WP100988** - Capacity Planning for zIIP-Assisted IPsec (Updated 7Sep2007)

- IBM Manuals
 - **SG24-7316-00** - Introduction to the New Mainframe: z/VM Basics (26Nov2007)
 - **SG24-7473-00** - DB2 9 for z/OS Performance Topics (11Sep2007)

Some Goodies!

- Interesting Paper
 - “*Common z/OS Challenges that Can be Resolved*”
 - Found with Google, then traced back to IBM Education Assistant
 - Go to <http://publib.boulder.ibm.com/infocenter/ieduasst/stgv1r0/index.jsp>.
 - Select ‘z/OS Operating System’, then select ‘z/OS V1.7’
 - Even though it’s for V1.7, most sections are applicable to any release
 - Topics include ASM, RSM, CSA Tracker, JES2 Exit Changes, etc.

Last SHARE: Capacity Planning

- Majority of capacity plans use service units
- Service units aren't consistent
 - Small LPARs on a large CEC appear to take less time (based on the SU/Sec value assigned), but in fact take more CPU time due to overhead from other LPARs
 - Service unit usage increases as the CEC utilization increases, especially for memory hogs
 - Service units per second are constant when using IRD, even as LPs are varied on and offline
 - Variation used to be within 5%, now it's closer to 20-30%
 - **We have submitted a requirement that all LPARs on a CEC use the CEC value of SU/Sec instead of the LPAR view of SU/Sec, or at least have the data element available for users.**

Help EWCP With Requirements

- Please help us out and vote
- Go to www.share.org, select Member Center, then select Requirements, and get a logon id
- Indicate interest in EWCP project, then review requirements and add discussion not and/or vote (3 are ready for voting, 3 are ready for discussion)
- Need votes for:
 - [SSEWCP05012](#) - WLM must consider I/O delay when starting and stopping Initiators
 - [SSEWCP064198](#) - Add SMF timestamp in TOD format to all SMF records
 - [SSEWCP07001](#) - Provide CEC Service Unit Rate in SMF & RMF Records

This SHARE: z/OS 1.9 Info Center

- New z/OS 1.9 Information Center (Beta)
 - Presented in session 0102 by **Geoff Smith**
 - New documents will be in html format, enabling more hits with search engines like Google
 - Go to - <http://publib.boulder.ibm.com/infocenter/zos/v1r9/index.jsp>
 - Send feedback with 'Contact z/OS' link at bottom of page
 - I ***LOVE*** it!

This SHARE: SMF Logger

- Several Sessions on SMF Logger
 - See sessions #2853 (**Peter Relson**), #2850 (**Skip Robinson**), and #2530 (**Marianne Hammer**)
 - Handles higher SMF rates with no overhead; looks great especially when you don't have contention on your CF
 - Can allow you to turn on more record types
 - Be careful of duplicate data – there is no CLEAR function
 - Needs sysprog to turn on/off dynamically
 - Be sure that all APARs are applied
 - I'd like to get an idea of current SMF activity – would you help?
 - I'll include results at next SHARE
 - Email (cheryl@watsonwalker.com) or fax (941-827-9493) a copy of any output from IFASMFDP (which shows types of records, number of records, record size, and interval), preferably for full day (or at least including peak periods)


This SHARE: HiperDispatch

- Several Sessions on HiperDispatch for z10
 - See sessions #2540 (**Horst Sinram**), #2838 (**John Eells**), #2832 (**Harv Emery**), and #2875 & #2532 (**Gary King**)
 - Decreases CPU time for CPU-intensive work by keeping tasks near their cache (APARs available for z/OS 1.8 +)
 - This is expected to reduce variability in job CPU times
 - **John Eells** had a great example of comparison – see the next slide from his presentation 2838...
 - L1, L1.5, and L2 are on same MCM, with up to 4 CPs on each MCM
 - L1 & L1.5 are unique to single CP, but L2 shared between CPs on the MCM
 - Default is turned off, but you can turn it on dynamically
 - Note - LSPRs were run with HiperDispatch turned on
 - Look for white paper on HiperDispatch at WSC soon
 - This is SOOO exciting!

Hiperdispatching—Why

- **Cache and memory latency on a hypothetical server**

- **L1 Cache** – 1 machine cycle – Lake Buena Vista, FL — 0 miles
- **L1.5 Cache** – 4 machine cycles – Bay Hill, FL—4 miles
- **L2 Cache** – variable, 10's of machine cycles – Kennedy Space Center—40 miles
- **Real memory** – ~ 600 machine cycles – New Orleans, LA—600 miles



2509 - Cheryl's Hot Flashes #19 - www.watsonwalker.com

See You in San Jose!



- **Email:** technical@watsonwalker.com
- **Web site:** www.watsonwalker.com