



Cheryl's Hot Flashes #21

Cheryl Watson
Watson & Walker, Inc.

March 6, 2009
Session 2509

www.watsonwalker.com

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

Agenda

- Survey Questions
- z/OS 1.10 SMF Goodies
- User Experiences
 - Cooling Problem
 - PTF Performance
 - High DB2 V8 zIIP Times
- MTTR
- Workload Characterization
- Miscellaneous Notes
- Politics

Survey Questions – Hardware

- Current Server Type (now or within next 12 months)
 - z800, z890?
 - Z900?
 - z990 (14%)?
 - z9-BC (25%)? z9-EC (35%)?
 - Z9-BC (0%)? z10-EC (20%)
 - Using HiperDispatch?
- Using zAAP Processors (22%)?
- Using zIIP Processors (30%)?
- Using IFLs?
- Have Used On/Off Capacity on Demand (20%)?
- Using Variable WLC Pricing (27%)?

Survey Questions – Software

- Operating System
 - z/OS 1.7 (25%)? (EOS was September 2008)
 - z/OS 1.8 (33%)?
 - z/OS 1.9 (42%)?
 - z/OS 1.10 (0%)?
 - Earlier than z/OS 1.7?
- Using WebSphere on z/OS (33%)
- Have used zPCR (12%)?
- Using or planning to use HyperPAV (25%)?
- Using or planning to use zLinux (39%)?
- Using z/OS 1.9 SMF Logger (0%)?

Cooling Recovery

- Prevent Outage if Cooling Problem
 - On the System z machines (z890, z990, z9s, and z10s), IBM has provided a reliability safeguard to keep the systems up and running if the CPs get too hot because there is a problem with the cooling system.
 - In this situation, the machine slows the processor speed down by 5-14%, issues a message to the operator, and calls home to IBM to notify them of a hardware problem.
 - This is all goodness, because you would rather just slow the machine down than take an outage because a chip burned out.
 - Because this is expected to be a very short time, z/OS changes the SRM speed constant, but not the MSU value.
 - So a job that took 10 seconds of CPU time before the cooling problem could take 10.5 seconds after the slow-down. The service units of the jobs remain the same, but the MSUs could be higher than they should be.
 - In general, all of this should be okay.

Cooling Recovery (cont.)

- **Users Complained**

- Users on the MXG Listserver (www.mxg.com) complained about slowdowns that lasted for several days, or happened multiple times.
- The warning message appears innocuous: *IWM063I WLM Policy was Refreshed Due to a Processor Speed Change.*
- One user saw that the speed constant was changed by 25%, so assumed that the capacity was dropped by 25%.
- I published these results in my Tuning Letter 2008 No. 2, and some people got excited.
- My bad – I forgot to check with WSC.

- **Kathy Walsh Explained**

- The machine was slowed down by its 5-14%, but an MCL update was missing, so the speed constant was calculated incorrectly as 25%. An MCL update corrects that.
- Another installation had a problem because the installation did not allow the system to call home and report the hardware problem.
- See her session 2500.

Cooling Recovery (cont.)

- To Do
 - Automate a message detection so that you are aware of when this occurs.
 - Make sure that MCL is applied.
 - Review new WSC Flash (expected soon).
 - If this occurs during the peak MSU period of the month, indicate on SCRT report that an exception occurred so that you don't get overcharged.

PTF Performance Issue

- RSU0901
 - LE PTF UK41392 (fixing a PE PTF) contains ACTION HOLD warning about high CPU
 - APPLY on a z9-BC (R02) took 45 minutes CPU time and produced 14 million lines of output to SYSPRINT!
 - PTF appears to relink entire LE component
 - See [PK75057](#) (3Dec2008) and [PK73828](#) (29Oct2008)

Thanks to **Ed Webb** of SAS Institute Inc

High DB2 zIIP CPU Times

- APAR [PK75214](#)
 - *The IIP CPU Time Being Reported is Higher Than Actual for RRSAF Threads (2Feb2009)*
 - Three zIIP accounting fields (QWACCLS1_ZIIP, QWACCLS2_ZIIP, and QWACZIIP_ELIGIBLE) could show higher CPU times than the capacity of the zIIPs.
 - This can happen on DB2 V8 for RRSAF threads after a parallel query is run on a thread that is reused.

- **MTTR – Mean Time To Recovery**
 - MTTR is the time to take down a system and re-IPL it.
 - This could take 15 minutes to 1.5 hours depending on . . .
 - This window needs to be shortened, so many groups in IBM are working on methods to reduce their startup/shutdown times (CICS, IMS, etc.).
 - A Redbook residency will be held in May to run benchmarks and publish results.
 - I'm very excited to join the residency in Poughkeepsie.
 - Please help!
 - If you're willing to provide either system logs or job logs, please send me an email.
 - If you've done some tuning and found things that reduce either the shutdown or restart, please let me know also.
 - How long does it take you to shutdown/restart? 30 minutes? 60 minutes? Over 1 hour?

Workload Characteristics

- zPCR is provided at no charge to customers from the WSC – [PRS1381](#) (or search for ‘zPCR Download’). Latest is 1Feb2009.
- 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. There is a new DI- Mix to see ‘worst case’ scenario. (**Kathy Walsh**, session [2500](#))

Workload Characteristics (cont.)

- What you can do:
 - Apply APAR [OA22414](#) (z/OS 1.8-1.9, 25Jun2008) which writes workload characteristics (e.g. I/O density) in the SMF Type 23 (Statistics) record.
 - Apply APAR [OA27161](#) (z/OS 1.8-1.10, 4Feb2009) which adds some fields.
 - Turn on type 23 records.
 - Contact **John Burg** (jpburg@us.ibm.com) and provide him some data.

Workload Characteristics (cont.)

- Second Stage – Hardware Instrumentation Service
 - Also called z10 CPU Measurement Facility
 - Very, very, neat!
 - Hardware monitor to collect counters and sampling.
 - Only available on z10 with APARs
 - See **John Burg**'s session [2536](#) (WSC Short Stories & Tall Tales) for more information.
 - Documentation is in MVS System Commands (!)
 - Two modes – collect counters (negligible overhead) and sampling (can be high overhead).
 - Caution when sampling – current default is sampling turned on producing 800,000 samples per minute (start with 320 instead and work up)! Don't set to over 130,000 for the z10-BCs.
 - Creates new SMF Type 113 record with statistics.
 - Output is a little cryptic (also see **Ed Jaffe**'s session [2839](#)).

z/OS 1.10 SMF Goodies

- Type 42
 - Subtypes 20 and 21 show when PDS and PDSE members are deleted, and who did it.
 - New z/OS 1.10 subtypes 24 and 25 show when PDS and PDSE members are changed or renamed.
 - Cool!
- Type 19
 - Contains DASD free space information for every online volume at IPL and SMF Switch.
 - Causes reserve on each volume.
 - z/OS 1.10 reduced the time during the reserve, so it may not be as bad now. Be careful of daytime SMF switches.
 - z/OS 1.10 added about 20 new fields, including the capacity of the device.

This SHARE

- GETMAIN Change in z/OS 1.10
 - In order to improve DB2 performance, the logic for obtaining storage with GETMAIN or STORAGE OBTAIN is changed.
 - Some ISVs and applications threw up on z/OS 1.10 because of this.
 - APAR [OA27291](#) (z/OS 1.10, 4Feb2009) was issued to change default behavior to previous releases, but you can override in DIAGxx member (VSM USEZOSV1R9RULES(YES/NO)).
 - If large DB2 site and high CPU, then use the z/OS 1.10 rules.
- My RedBooks
 - One of the neatest additions to IBM's Redbooks is a project called 'My Redbooks', which is provided by alphaWorks Software. You can find a link to it on the Web site at www.redbooks.ibm.com, or with a direct link to <http://services.alphaworks.ibm.com/MyR/>.

SMF Report of Subtypes

- EWCP Requirement
 - Requests that IBM change SMF dump program to produce listing by sub-types instead of just record types.
 - Type 42, for example, has 30 subtypes.
 - Type 99, subtype 6 is small volume and very useful, but other types are fairly complex and produce huge volumes.
 - Please logon to SHARE Requirements section and indicate an interest in the EWCP requirements, and add your vote if you think it's worthwhile.
- New “Free Tools”
 - New section of our website.
 - First is SAS program to create report of SMF subtypes.
 - If you use it, please send me the output of the report.

- The ITSO (International Technical Support Organization) is celebrating its 40-year anniversary. You can read about their history here:
<http://www.redbooks.ibm.com/redbooks.nsf/pages/history?Open>
.
- Interesting Redbooks statistics:
 - Every month over 500,000 IBM Redbooks are downloaded.
 - Over 100,000 Redpapers are downloaded per month.
 - Over 320 Redbooks materials were published in 2007 - over 400 by year end 2008.
 - Over 4,800 IBM Redbooks have been published.
 - Over 1,100 residents participate in the residency workshops each year.
 - In 2007, there were 287 workshops with almost 6,000 students.
 - Over 80,000 technical professionals subscribe to the IBM Redbooks Weekly Newsletter.

z/OS 1.10 Info Center



- New z/OS 1.10 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/v1r10/topic/com.ibm.zos.r10/zosr10home.html>
 - Send feedback with 'Contact z/OS' link at bottom of page
 - I ***LOVE*** it!

Politics (Soapbox Time)

- Chargeback on z10
 - Variability issues on z10.
 - Use ranges or bands for pricing.
- zLinux
 - Several installations know that zLinux will save money, but say that company politics prevents them from exploring.

See You in Denver!



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