

# Cheryl's Hot Flashes #20

Cheryl Watson  
Watson & Walker, Inc.

August 15, 2008  
Session 2509

[www.watsonwalker.com](http://www.watsonwalker.com)

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



# Agenda

- Survey Questions
- User Experience
  - SMF & SMF Logger
  - Distributed DVIPA
  - HiperDispatch
  - z/OS 1.8 CPU Reductions
  - Hiper CICS APAR
- New Function APARs
- Hiper APARs
- Some Goodies!
- This SHARE
- Last SHARE

## Survey Questions – Hardware

- Current Server Type (now or within next 12 months)
  - z800, z900, z890 (34% at last SHARE)?
  - z990 (18%)?
  - z9-BC (19%)? z9-EC (32%)?
  - z10 (8%)
- Using zAAP Processors (19%)?
- Using zIIP Processors (23%)?
- Have Used On/Off Capacity on Demand (12%)?
- Using Variable WLC Pricing (18%)?

# Survey Questions – Software

- Operating System
  - z/OS 1.7 (75%)? (EOS is September 2008)
  - z/OS 1.8 (80%)?
  - z/OS 1.9 (21%)?
  - Older (22%)?
- Using WebSphere on z/OS (25%)
- Have used zPCR (11%)?
- Running WLM-managed initiators (50%)?
- Using or planning to use HyperPAV (19%)?
- Using or planning to use zLinux (60%)?
- Using z/OS 1.9 SMF Logger (0%)?

# SMF Logger (Last SHARE)

- Several Sessions on SMF Logger
  - Handles higher SMF rates with no overhead; looks great especially when you don't have contention on your CF (don't know about when there's contention, but speed of CFs would appear to handle additional load)
  - Can allow you to turn on more record types (especially DB2 records)
  - 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](mailto:cheryl@watsonwalker.com)) 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); also a copy of SMFPRMxx

# Last SHARE: SMF Analysis

- Volume – the great majority of LPARs analyzed produced less than 1MB per second of data. One LPAR produced about 2.2MB/second during a peak period, and a few produced about 1MB/second during a peak period.
- Number of Records – Roughly one-third of the LPARs produced less than a million records per day; another third produced between 1 million and 10 million records per day; and the other third produced over 10 million records per day, with many of these in the 30-65 million range.
- Record Size – For installations that don't collect DB2 data, the average SMF record size is about 1000 bytes, while those that collect a lot of DB2 data produce average record sizes of around 2000 bytes. I saw one installation with an average LRECL of 4,500 bytes. That was due to a large volume of records that can be 20K to 30K in size: SMF 30 (batch) and CICS 110 records.
- DB2 Type 102 – Most installations do not collect DB2 type 102 records. Those that do, typically see that the 102 records produce between 70% and 85% of the total number of SMF records.
- Intervals – The majority of installations collected RMF data at 15 minute intervals; many used an interval of 10 minutes; and a few used an interval of 30 minutes. Without seeing the SMFPRMxx members, I couldn't tell what the SMF interval was set to, but I suspect that it was 30 minutes, which is the IBM default.

# SMF Logger

- My biggest concern was the difficulty in preventing duplicate data, but there are three solutions now:
  - WSC White Paper [WP101271](#), provides everything you need (including a REXX program to control dates) to set up and run SMF Logger. If you downloaded the REXX program originally, pull it off again because there were some corrections. The paper was written by **Ozan Baran** and **Andrew Sica**, who are part of the Poughkeepsie Systems Test team.
  - Computer Associates' CA SMF Director (used to be CA-JARS/SMF) will also handle the scheduling of SMF dumps to prevent duplicates.
  - RYO

# DVIPA User Experiences

- Distributed Dynamic Virtual IP Addressing (DVIPA) became available with z/OS 1.9.
- VIPADISTRIBUTE is a parameter to specify the percentage of a zIIP processor to use for the TCP/IP work.
- Prior to VIPADISTRIBUTE, there is a problem with WLM and TCP/IP because WLM did not recognize the zIIP engine in the workload balancing.
- In z/OS 1.9 you can specify percentages that your DDF workload is expected to use of the zIIP processor and the general processor. This is specified for each node for each DB2 subsystem in TCP/IP with the VIPADISTRIBUTE parameter.
- Redbooks:
  - z/OS 1.9 Implementation Redbook - SG24-7427-00
  - Redpaper - REDP4449 – DB2 9 for z/OS Data Sharing



# DVIPA User Experiences

- NationWide Experience
  - After VIPADISTRIBUTE, the zIIP processor usage went up by 30%, and customers were thrilled because their bills were decreased.
  - Use the IIPCP value in the RMF workload report to see the split in the DB2 DDF service class in order to set VIPADISTRIBUTE.
  - They were using the DVIPA Sysplex Distributor & BASEWLM options.
  - Thanks to Janet Howie.
- Another Experience
  - At another site, when both z9s and z10s were in the same pool, the routing sent more work to the z10s. They didn't use zIIPs.
- Your mileage may vary, but it's worth looking into when you go to z/OS 1.9.

# HiperDispatch – User Experiences

- Several Sessions on HiperDispatch for z10
  - Decreases CPU time for CPU-intensive work by keeping tasks near their cache (APARs available for z/OS 1.8 +)
  - Default is turned off, but you can turn it on dynamically
  - LSPRs were run with HiperDispatch turned on
  - WLM policy may need to be changed. See WSC White Paper [WP101229](#) by **Steve Grabarits**.
  - There may even be problems with HiperDispatch turned off (see APAR OA25825 for zIIP processors).
  - One site found that IRD didn't work the way it used to on the z9, even with HiperDispatch turned off. Has anyone else seen that?

## z/OS 1.8 CPU Reductions

- An installation moved from z/OS 1.7 to z/OS 1.9 and saw a dramatic reduction (75%) in the RMF address space.
- The RMF team responded with this explanation: For z/OS 1.8 and up the processing within GPMSERVE was changed to request only data from active devices instead of data for all configured devices. An “active” device is one that has I/O occurring during the RMF Monitor III interval (in this shop, it was 60 seconds). Since I/O does not occur to ALL configured devices in a given minute, but instead to a much smaller subset, this is a great reduction in data collection overhead that shows up as a significant CPU reduction in RMF which actually collects the device data.
- They also saw a significant decrease in \*MASTER\* (probably due to change in UIC processing in z/OS 1.8).

Thanks to **Jerry Urbaniak** from Acxiom

# HIPER CICS APAR

- From **Ken Williams**: “This hit us hard here and could be worth more than 250 MIPS in our busiest CICS regions.”
  - **OA24094** (z/OS 1.7-1.9, 21May2008) - *In a CICS Transaction Isolation Environment (ie. Using Subspaces) High Performance Overhead In IAXVP Doing CSP.* In a CICS transaction isolation environment (ie. subspaces are being utilized), very high performance overhead occurs in IAXVP during CSP processing due to the private storage getmains being done which drives page table initialization. The cause of the high performance impact is the PTLB (purge translation lookaside buffers) which occurs upon the CSP instruction. Note: This problem will not be encountered if NOT running in a CICS transaction isolation environment where subspaces are NOT utilized. [Note – Are there too many ‘NOT’s in that last line?]

## New Function APARs

- **OA18324** (z/OS 1.7-1.9, 30Apr2008) - *New RRS Health Checks - Storage Usage.*
- **OA18990/OA24385/OA24386/OA24388** (z/OS 1.7-1.9, 20Mar2008) - *RRS Task Hung Waiting for an ENF Signal from Logger with No External Symptoms Displayed*
- **OA19843** (z/OS 1.7-1.8, 25Apr2008) - *RM Unable to Restart with RRS - New Function to Allow RM to be Unregistered with RRS*
- **OA22112** (z/OS 1.8-1.9, 14Feb2008) - *New Function - XRC Secondary VOLSER and Device Number Filters*
- **OA22160** (z/OS 1.8-1.9, 30Jan2008) - *New Function. The WLM Service Macro IWMEQTME Returns No zAAP Service Times*
- **PK37032** (Communications Server for z/OS 1.8+, 31Jul2007) - *Support For FTP Batch Pipes.*

Thanks to **Jerry Urbaniak** of Acxiom

## New Function APARs

- **OA22481** (z/OS 1.7-1.9, 3May2008) – *New Function - Add Support for CEX2C with RSA 4096-Bit Encryption/Decryption and Signature Generation / Verification.*
- **OA22573** (z/OS 1.6-1.9, 20Mar2008) - *New Function APAR - Provide an Ability to Set a Tunable Time Value for Certain IOS Recovery I/O Requests.*
- **OA22899** (z/OS 1.7, 21Mar2008) – *New Function VSAM LSR Buffering. Relates to CICS 3.2 threadsafety.*
- **OA23174** (DFSMS 1.8-1.9, 6Mar2008) - *New Function - zIIP Enablement for XRC.*
- **OA19194** (z/OS 1.7-1.8, 14Sep2007) - *New Function - AMATERSE, Alias TRSMAN, Is a Supported Service Aid Function.*
- **OA21635** (z/OS 1.6+, 20Sep2007) - *New Function - Recalibrate SYNC/ASYNCR Thresholds.*

## New Function APARs

- [OA17177](#) (z/OS 1.6-1.9, 30May2008) - *New Function APAR - Coupling Over Infiniband Support.* New function APAR to provide IOS support for Infiniband CHPIDs in a Coupling Facility environment.
- [OA22161](#) (z/OS 1.7-1.9, 2Jun2008) - *New Function - Add Unique Self-Description Data for ECKD Virtual DASD.*
- [OA23828](#) (z/OS 1.8-1.9, 18Jun2008) - *New Function - XML System Services Now 100% zIIP Offloadable in Enclave SRB Mode.* This APAR does fix a performance related issue as well as provide new function. For additional information, also see White Paper [WP101277](#), DB2 9 and z/OS XML Services Synergy Update (Aug2007).
- [OA24658](#) (z/OS 1.8-1.9, 9Jul2008) - *New Function - D GRS,C,LATCH Enhancement.*



## HIPER APARs

- **OA24049** (DFSMS 1.8-1.9, 17Mar2008) - *Increased CPU Consumption for Jobs That Are Using PDSE.*
- **OA24096** (z/OS 1.8-1.9, 14May2008) - *Corrections to the WLM Group Capacity Algorithm.*
- **OA24890** (z/OS 1.7-1.9, OPEN) - *Initiators Being Started May Terminate during Job Selection Unexpectedly after **OA20474**..*
- **PK55505** (Websphere MQ V6, 13Mar2008) - *High CPU in the Mstr Task.* The customer reported the following: a nnnnMSTR address space was using high CPU for several days. It was using over 90% when typically it uses around 2% or 3%.
- **OA24852** (z/OS 1.8-1.9, OPEN) - *Duplicate Jobs May Execute Concurrently With DUPJOB=DELAY. (After **OA22474**)*

Thanks to **Jerry Urbaniak** of Acxiom



# HIPER APARs

- [OA23576](#) (z/OS 1.7-1.9, OPEN) - *Corrupted Alternate Couple Data Set.*
- [OA24260](#) (DFDSS 1.8, 19Jun2008) - *After Physical Data Set Copy, the Extent Hi-RBA Value is Invalid for an Extended-Attribute VSAM.*
- [OA24543](#) (DFSMSrmm 1.7 – 1.9, 16May2008) - *All But Last Data Set Deleted from Multifile Tape after S222 when Job Cancelled.*
- [OA24920](#) (z/OS 1.7+, 2Jul2008) - *Dedicated LPARs, Shared LPARs with AWM Off May Experience Performance Degradation.*
- [OA25310](#) (z/OS 1.7-1.9, OPEN) - *Catalog Control Block Overlay When VVDS > 120 Extents.*
- [OA22506](#) (z/OS 1.8-1.10, 4Jun2008) - *Discretionary Work Unilaterally Swapping at z/OS 1.8.*
- [OA25143](#) (z/OS 1.9, OPEN) - *Swapout Delayed 60 Seconds when Address Space Using Div is Being Swapped ==> Mark Non-Swappable to Prevent Swapping.*
- [OA25320](#) (z/OS 1.9, OPEN) - *SRB Scheduling Keeps Dispatching Priority High for Batch Job.* APAR [PK66986](#) (closed on 11Jul2008) addresses the SDSF portion of this issue.

# This SHARE



- SMP/E 3.3 & 3.4 need toleration APAR before z/OS 1.10 GA
  - See WSC Flash 10649 dated June 12, 2008, and toleration APAR IO07480. As the Flash indicates, "Failure to install these PTFs will cause RECEIVE commands to fail with message GIM20311E....". You could see 30,000+ error messages on ++HOLDS, and the SMP/E job will get a return code of 8. It will still work, but it messier.

## Some Goodies!

- Articles on WSC Techdocs
  - PRS3172 - *zSref - System z Time Considerations*, **Leon Daniels**.
  - WP101208 – Running IBM System z at High Utilization, Gary King
- 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](http://www.redbooks.ibm.com), or with a direct link to <http://services.alphaworks.ibm.com/MyR/>.

## Some Goodies!

- From Watson & Walker
  - SHARE “Hot Flashes” from spring 2008 – missed deadline for submission, so not on SHARE proceedings. (My apologies.) You can get it from our website at <http://www.watsonwalker.com/articles.html> (then click on Presentations).
  - z/OS 101 – New series of articles for z/OS zNextGen people. This will be updated every other month. See ‘Articles’.

## Last SHARE: zPCR Analysis

- zPCR is provided at no charge to customers from the WSC – [www.ibm.com/support/techdocs/atmastr.nsf/Webindex/PRS1381](http://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. There is a new DI-Mix to see ‘worst case’ scenario. (**Kathy Walsh**, session 2400)
- **APAR OA22414** (z/OS 1.8-1.9) which writes workload characteristics (e.g. I/O density) in the SMF Type 23 (Statistics) record. After you apply APAR, please contact **Gary King** ([garyk@us.ibm.com](mailto:garyk@us.ibm.com)) and provide him some data.

# Last SHARE: zPCR Analysis

- In a previous newsletter, 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

# Last SHARE: 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')

## Last SHARE: 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](http://www.cmg.org/proceedings) for these public documents
- Sign up for MeasureIT – another publication for members and non-members
- Previous issues contain multi-part series on SAS 9 Tips by **Phil Mason**
- Go to [www.cmg.org](http://www.cmg.org) and select 'Publications', then 'MeasureIT'



## Last 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!

# See You in Austin!



- **Email:** [technical@watsonwalker.com](mailto:technical@watsonwalker.com)
- **Web site:** [www.watsonwalker.com](http://www.watsonwalker.com)