

S H A R E

Technology • Connections • Results



Cheryl's Hot Flashes #22

Cheryl Watson
Watson & Walker, Inc.

August 28, 2009
Session 2109

www.watsonwalker.com

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

Agenda

- User Experiences
 - SMF Logger
 - CICS SMF Compression
 - SMF Type 30s
 - MIH for I/O Connections
- Workload Characteristics
- MTTR
- z/OS 1.11 Goodies in z/OS 1.10
- Software Pricing
- My Favorite Things
- Survey Answers

SMF Logger

- Great savings reported
- Some frustration; keep current with maintenance
- APAR OA27037
 - Provides z/OS 1.11 enhancements to z/OS 1.9 & 1.10
 - Most important – provides ARCHIVE, RELATIVEDATE, MAXDORM, DELETE
 - Answers most of concerns: no need for REXX program to specify the dates; allows deletion of records after they're retrieved; allows just deletion of records; uses MAXDORM to reduce the possibility of duplicate records.

SMF Logger (cont.)

- OA27037 (cont.)
 - The documentation about MAXDORM has confused some people.
 - Assume that you want to pull data from midnight to midnight and you use ARCHIVE to do that (e.g. RELATIVEDATE(BYDAY,1,1)).
 - Without this APAR, an SMF record might contain some records from yesterday and some from today. If you dump by date, you would dump all the records from the block today, and then dump those same records tomorrow. Because this APAR causes SMF to close a block each MAXDORM interval, there is little chance to offload duplicate records when using ARCHIVE. Most MAXDORM values are 30 minutes.
- Still some problems – track OA29903
 - *Unexpected Message IFA832I For Invalid Parameters With the Archive Option When Using IFASMF DL Utility (z/OS 1.9-1.11, OPEN)*
 - If you specify the following:
 - OUTDD(OUTRMF,TYPE(70:79))
 - OUTDD(OUTREST,TYPE(0:69,80:255))
 - Then you get an error. For now, you can specify:
 - OUTDD(OUTALL,TYPE(0:255))
 - OUTDD(OUTRMF,TYPE(70:79))

CICS SMF Compression

- CICS/TS 3.2 and 4.1
 - Option in the MCT to compress CICS SMF records (default in 4.1, but you can override it – DON'T!).
 - This only applies to type 110, subtype 1 (majority of CICS SMF records).
 - Significantly reduces SMF size using standard z/OS compression.
 - Some examples: 60% reduction in tapes; frequency of SMF dumps reduced from every 5-8 minutes to every 39-61 minutes; frequency of SMF dumps reduced from every 11 minutes to every 32 minutes.
 - IBM results show that type 110 records are reduced from 32K to 6K (80% reduction in size).
- My wish. . .
 - All SMF records could be compressed.

SMF Type 30 Records

- Problem: can produce multiple 32K records
 - The type 30, subtypes 2 & 3 (interval end), 4 (step end), and 5 (job end) can contain EXCP sections, with one 30-byte section for every DDname used by the job, plus one empty section for any alternate devices for certain SMS allocations.
 - For DB2, and similar subsystems, all of these subtypes result in multiple, continued 32K SMF records.
 - IBM has provided several options to reduce the size and processing time.
 - SMFPRMxx parameter DDCONS(YES/NO)
 - The default processing is to consolidate the EXCP sections (combining multiple sections for the same DDnames).
 - This consolidation also causes high CPU times when the records are being written. This is especially true when intervals are synchronized. So IBM provided DDCONS(NO) to reduce CPU time, but it can increase record sizes.
 - YES is the default, but NO is the recommendation.

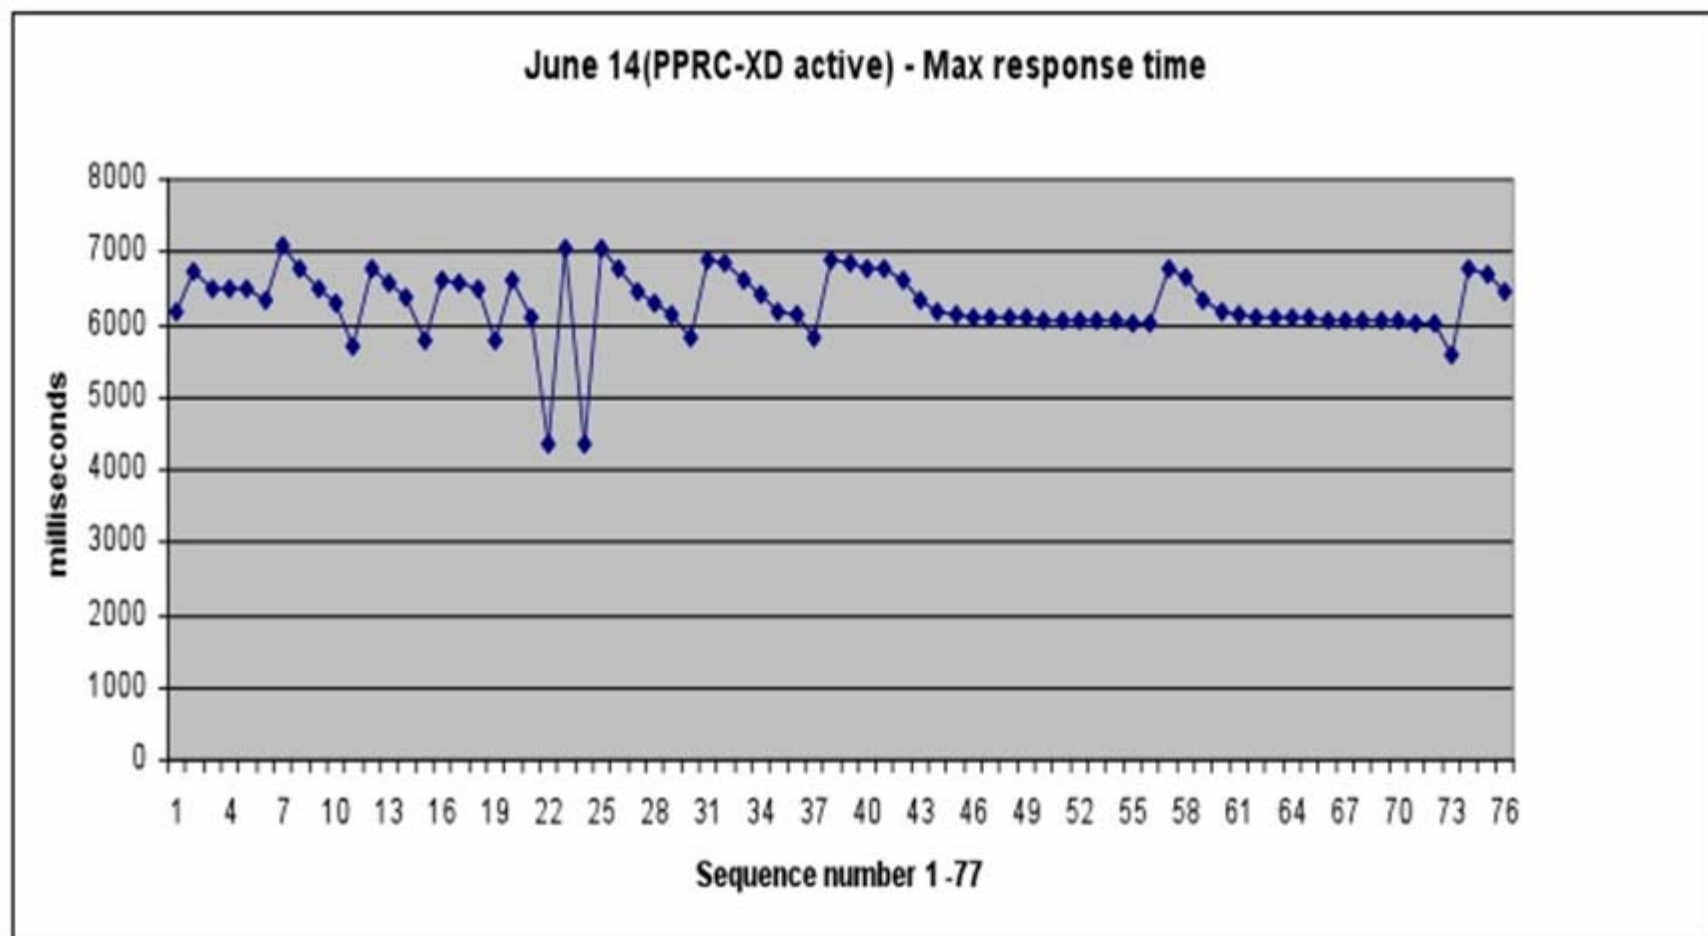
SMF Type 30 Records

- Problem: Multiple 32K Records (cont.)
 - SMFPRMxx parameter DETAIL/NODETAIL – Option that is confusing to many people.
 - For JES – NODETAIL has no effect – EXCP sections for batch jobs are written to subtypes 4 & 5, and also to subtypes 2 & 3 if INTERVAL is specified.
 - For TSO – NODETAIL applies only to the type 32 record, not the type 30 record.
 - For STC – NODETAIL might apply
 - If NOINTERVAL, then NODETAIL doesn't apply – EXCP sections written to subtype 4 & 5.
 - If INTERVAL, and default of DETAIL is used – EXCP sections written to subtypes 2, 3, 4 & 5.
 - If INTERVAL and NODETAIL – EXCP sections written to ONLY subtypes 2 & 3 (none in subtypes 4 & 5).
 - SMFPRMxx parameter EMPTYEXCPSEC(SUPPRESS/NOSUPPRESS)
 - Suppress removes any EXCP section that has zero values. The majority of these can be caused by the implementation of SMS Dynamic Volume Count function. Default is NOSUPPRESS.
 - See Hot Topics article in the August 2009 issue.

MIH Functions in I/O Communications

- Problem
 - Customer was seeing some poor response times in online transactions. For example a transaction that was normally complete in milliseconds might take 2 or 6 seconds. This backed up other transactions.
 - Customer did an analysis of SMF type 42 records, subtype 6, and plotted the maximum response times by time of day. Many of the maximum response times fell into the 2 second or 6 second time frame.
 - Thanks to Don Chesarek, IBM, who identified and is documenting this situation.
 - There are some internal MIH-like functions (FICON & FCP in adapters) that will send out an I/O, and if the device is busy and a response does not come back, will automatically redrive the request after a wait for 2 (FICON) or 6 (FCP) seconds before trying it again. If it makes it on the second try, it just looks like it took a long time.
 - The problem occurred about 1000 times a day at this site.
 - Maybe you should plot the maximum response times?

Max dataset response time from SMF 42-6 records (Intermittent errors(1,200/day) on PPRC paths)



Workload Characteristics

- zPCR is provided at no charge to customers from the WSC – [PRS1381](#) (or search for ‘zPCR Download’). Latest is 18Aug2009.
- 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.

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 delta 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 [2136](#) (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.
 - Contact John to send him your SMF types 23, 70, 71, & 113s.

- **MTTR – Mean Time To Recovery**
 - MTTR is the time to take down a system and re-IPL it. IBM is working on a multi-phase plan to reduce MTTR.
 - MTTR can take 15 minutes to 1.5 hours depending on . . .
 - I participated in Redbook residency in May. Book expected to be published in the next two months. Look for *SG24-7816 – System z Mean Time to Recovery Best Practices*.
 - Frank Kyne, Project leader, gave SHARE session 2235 this week on MTTR. The presentation has LOTS of good suggestions. The Redbook will have even more.
 - Two items from his presentation:
 - IPCS command: `IPCS VERBX BLSAIPST MAIN` formats the IPST control block. This shows the steps and their time during IPL. The Redbook will explain most of the steps.
 - Redbook *SG24-7328 – z/OS Planned Outage Avoidance* is terrific! It provides suggestions on avoiding IPLs.
 - An ITSO residency is an absolutely wonderful experience!

z/OS 1.11 SMF Goodies in 1.10

- Many z/OS 1.11 features are available as APARs in z/OS 1.10 (and z/OS 1.9)
 - OA24809/OA24816/PK64715 (z/OS 1.8-1.10) – Metro Mirror support for FlashCopy (Preserve Mirror)
 - OA25426 (z/OS 1.10) – API to Perform HMC Functions
 - OA25559 (z/OS 1.9-1.10) – Support for solid state devices
 - OA26084/OA29017 (z/OS 1.9-1.10) – zHPF Multitrack Support
 - OA26802 (z/OS 1.9-1.10) – Enhancements to SYSREXX
 - OA26832 (z/OS 1.10) – Expanded Service Unit Values in the SMF Type 30 Record.
 - OA26993/OA26995 (z/OS 1.10) – Allows AutoIPL with Sysplex Failure Management (SFM) Policies
 - OA27393 (z/OS 1.8-1.10) – DS8000 Drive-Level Encryption

z/OS 1.11 SMF Goodies in 1.10

- Many z/OS 1.11 features are available as APARs in z/OS 1.10 (and z/OS 1.9)
 - OA27495 (z/OS 1.10) – zAAP on zIIP Support
 - OA28323 (z/OS 1.9-1.10) – New STP Messages
 - OA29019 (z/OS 1.9-1.10) – Method to Track RTLS Usage
 - PK71213 (z/OS 1.9-1.10) – FTP Support for Named Pipes
 - PK85856/PK85899 (DB2 V8-V9, z/OS 1.10) – DFSORT Support for Additional zIIP Offload by DB2 Utilities
- See Bob Rogers' presentation 2228 (*Sysprog Goody Bag*)
 - He describes some of the little heralded features of z/OS 1.11

Software Pricing

- Sub-capacity Pricing
 - Can provide significant savings in IBM software, especially as new functions move more work to zIIPs and zAAPs.
 - Everybody should be running the Sub-Capacity Reporting Tool (SCRT) so that they understand what they are currently using.
 - New functions come out all of the time, but the savings usually aren't seen unless you're using Sub-capacity pricing (e.g. SyncSort & DFSORT ability to use zIIPs; XML support for specialty engines).
- ALERT! Most of the larger software vendors (like CA & BMC) provide sub-capacity pricing similar to IPLA, but you might not know about it.
- This is the reason that the addition of zIIPs and zAAPs can save some sites so much money. The software savings can usually pay for the specialty engines.

Software Pricing

- Chargeback on z10
 - Variability issues on z10.
 - Use ranges or bands for pricing.
- Lost Data
 - CPU time fields in SMF type 30 are only precise to one-hundredth of a second. On faster machines, the CPU time fields are zero.
 - Solution is to use service units for charge back & capacity planning.
 - But on newer machines, it's possible for the service unit fields to wrap.
 - Therefore, get APAR OA26832, which increases the length of the service unit fields. This is new in z/OS 1.11 or available via APAR on z/OS 1.9 & 1.10. You'll need to change any program that collects this data (check with your vendor).

Software Pricing

- The elephant in the room
 - A software vendor has just come out with a product that can move up to 50% of your work to specialty engines.
 - The question? Will IBM allow that?
 - First, look at the consequences:
 - Software charges would be dramatically reduced for all vendors who use sub-capacity pricing (this includes IBM, CA, BMC, and more). They won't be happy!
 - If IBM were to allow this, then every vendor would do the same thing, and try to move 100% of all work off the GPs.
 - It's VERY easy to do, but other vendors haven't done it because they've followed IBM's intentions that the specialty engines are used for a certain type of work. Software from ISVs that exploit specialty engines follow IBM's guidelines.
 - If IBM were to allow this, then somebody like me can come up with the same solution and sell it for \$50.
 - While you're still limited to having at least one GP for every zIIP or zAAP, if you have enough work for ten zIIPs and zAAPs, you would still need at least ten GPs.

Software Pricing

- Why I Don't Believe it Will Fly
 1. The reasons from the previous page.
 2. IBM's most recent statement:

To be clear, if customers use Specialty Processors/Engines to process workloads beyond the IBM authorizations, then customers must pay the hardware, software and maintenance pricing associated with full CP capacity.

 - See the full letter from Mark Anzani, VP & Chief Technology Officer, IBM System z:
 - <http://blogs.datadirect.com/media/Communication%20Letter%20DataDirect.pdf>
 - <http://blogs.datadirect.com/2009/08/violation-vs-exploitation-ibm-clarifies-authorized-use-of-ibm-specialty-enginesprocessors.html>
 3. I think that it's too big a risk of increased software charges from IBM and other vendors, and it might mar a good working relationship with IBM.

Some of My Favorite Things

- Favorite blogger
 - Willie Favero - <http://it.toolbox.com/people/wfavero/>.
 - On IBM Swat team for z/OS; DB2 sysprog.
 - Primarily DB2 blog, but also has many z/OS items – always useful.
 - You'll need to sign up for IT Toolbox (no charge); can subscribe with RSS feed.
- Neat Publications
 - Redbook SG24-6366-01 – *Introduction to the New Mainframe: z/OS Basics.*

Some of My Favorite Things

- Articles

- *50 Years of Big Iron*. This is an article from eWeek that shows photos of the "mainframes" for the past 50 years. My first experience with IBM machines was in 1962, and these bring back memories.
<http://www.eweek.com/c/a/IT-Infrastructure/The-IBM-Mainframe-50-Years-of-Big-Iron-Innovation-583073/>.
- *Enhance Your I/O Performance With IBM System z*. This article by Harry Yudenfriend in a Hot Topics "spotlight" shows the System z enhancements that have improved the performance of I/O processing. Some of the technologies are optional, so you have a choice in which improvements you implement.
http://publibz.boulder.ibm.com/zoslib/pdf/io_performance_enhancements.pdf.

Some of My Favorite Things

- New z/OS 1.11 Information Center
 - New documents are in html format, enabling more hits with search engines like Google
 - Go to - <http://publib.boulder.ibm.com/infocenter/zos/v1r11/index.jsp>
 - Prior to migration, you should review Changes section of relevant manuals (for example, increases in the SMF type 30 service unit fields aren't mentioned in the migration manuals)

Some of My Favorite Things

- Conferences on WSC
 - WSC link - <http://www-1.ibm.com/support/techdocs/atmastr.nsf>
 - Look at left side – select Auxiliary Materials
 - Then select Conferences, display by date
 - Get presentations for conferences
 - This is just starting, but I hope it will grow
- Plea to IBM Presenters
 - I really depend on your material
 - The SHARE proceedings are soon closed to non-members
 - Please post your sessions on the WSC site!

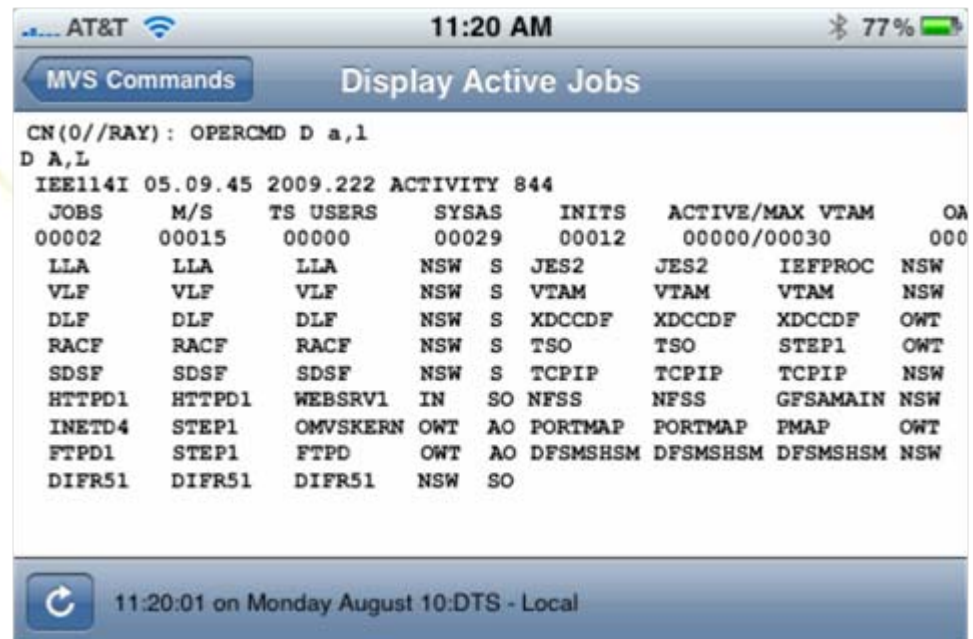
Some of My Favorite Things

- z/OS on iPhone
 - While at the STE, I found the DTS booth where they were showcasing their new iPhone app.
 - Being an iPhone enthusiast, this excites me! What's the future? RMF, Omegamon, SYSVIEW, Mainview apps?



Some of My Favorite Things

- z/OS on iPhone (cont.)



Survey Answers – Hardware

- Current Server Type (now or within next 12 months)
 - z800, z890? (8%)
 - Z900? (3%)
 - z990 (12%)?
 - z9-BC (26%)? z9-EC (28%)? (Marketing will be withdrawn soon)
 - Z10-BC (16%)? z10-EC (21%)
 - Using HiperDispatch? (7%)
- Using zAAP Processors (20%)?
- Using zIIP Processors (37%)?
- Using IFLs? (28%)
- Have Used On/Off Capacity on Demand (7%)?
- Using Sub-Capacity Pricing (31%)?

Survey Answers – Software

- Operating System (primary systems)
 - z/OS 1.7 (12%)? (EOS was September 30, 2008)
 - z/OS 1.8 (21%)? (EOS will be September 30, 2009)
 - z/OS 1.9 (43%)?
 - z/OS 1.10 (17%)?
- Using WebSphere on z/OS (19%)
- Have used zPCR (16%)?
- Using or planning to use HyperPAV (23%)?
- Using or planning to use zLinux (33%)?
- Using z/OS 1.9 SMF Logger (3%)?

IBM Survey

- Iris Rivera (IBM) is asking customers to fill out a “Consumability” survey. This is different than the yellow survey sheet that she’s been distributing this week.
- Please log on to

<https://www.ibm.com/survey/oid/wsb.dll/studies/consumabilitywebform.htm?product=2&icode=SHARE&brand=8&crit=18dc8dcfa3a35c@55&group=stg>

and complete the survey. It will help us all.

See You in Seattle!



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