Cheryl's Hot Flashes #6

Cheryl Watson
Session 2543, July 27, 2001

Watson & Walker, Inc.
publishers of 'Cheryl Watson's TUNING Letter'
http://www.watsonwalker.com

Cheryl's Hot Flashes

- Hardware Issues:
  - z900 Notes
  - DSNUTILB & Capacity Analysis
  - z900 Compression
- Software Issues:
  - z/OS or OS/390 R10 on z900?
  - How Much CPU O/H in R10?
  - Software Pricing
  - Mysteries
  - WLM Notes
  - Miscellaneous Notes
- Good Reading
- APARs

My Thanks

- Many people are kind enough to send me emails about their hardware and software experiences. Their main purpose is to share their experiences (often bad, but sometimes good) so that others won't suffer. My unending thanks to the following contributors to this session:
  - Jerry Urbaniak of Acxiom
  - Brian Currah of BDC Computer Services, Inc.
  - Dale Mattison
  - Derrick Haugan of MetLife
  - Marty Wertheim and Frank Durand of Bank of America
  - Janet Howie of Nationwide

My Thanks

- And more...
  - Jim Petersen of HomeSide Lending, Inc.
  - Jim Marshall of Office of Personnel Management in Washington, DC
  - Jack Hudnall of SBC Services
  - Ed Jurgenson from Humana
  - Nils Abrahamsson of Volvo
  - Greg Dyck, John Hutchinson, Walt Caprice, Frank Yaeger, and Marcy Nechemias of IBM, for their help on some of these items
  - IBM-Main participants
  - And others who prefer to remain anonymous

z900 Notes

- In my March SHARE session, I talked about some programs that didn't fare as well as expected on the z900. Most everything, like SAS, has been resolved.

- Most sites are finding that the z900 just flies, especially for online work. I've found two more instances of some types of work behaving worse than expected on the z900:
  - DSNUTILB in one site turned out to have the worst ten performing programs when moving to a z900.
  - Three sites reported that their slowest programs on the z900 (that is, less than expected) were old, legacy, COBOL programs.
  - These are very isolated cases, and we don't have any answers - please let me know if you experience similar conditions.

DSNUTILB & Capacity Analysis

- When moving from a 9672 to a z900, DSNUTILB was the worst behaving program. We don't know why. Anybody else see this?

BOXSCORE/BATCH 10:19 THURSDAY, JUNE 28, 2001

12

V1R4 (c) Watson & Walker, Inc.

Programs With Worst Outcome

BOXSCORE BATCH Analysis for 2064 Upgrade

Item Description	9672-XX7 DATA	2064-1C6 DATA	Delta	% Delta	Comments
Job/Step/Pgm With Worst Outcome:					
1: JOBABC/STEP101/**DSNUTILB**					
Average CPU/IO, MSec	60.0	133.7	73.7	122.8%	WWCB053-W Large variance from average
Observed MIPS/phys CPU	**166.3**	**74.7**	**-91.6**	**-55.1%**	
Rec count	3	6	3.0	100.0%	
Total EXCP count	116	224	108.0	93.1%	
Transactions	3	6	3.0	100.0%	
EXCPs per transaction	**38.7**	**37.3**	**-1.3**	**-3.4%**	
Average Users	3	6	3.0	100.0%	
Total CPU time, seconds	7.0	29.9	23.0	330.2%	
Coefficient of variation	6.4	4.7	-1.8	-27.4%	

DSNUTILB & Capacity Analysis

- When searching my email, I found a note from 12/99. When moving from a 9021 to a 9672, DSNUTILB was also the worst behaving program! But for a different reason (we think that additional memory let DSNUTILB take fewer I/Os).

BOXSCORE/BATCH 10:35 Friday, October 29, 1999

V1R2d (c) Watson & Walker, Inc. Programs With Worst Outcome

H5/G5 UPGRADE

Item Description	9/27-9/30	10/4-10/7	Delta	% Delta	Comments
Job/Step/Pgm With Worst Outcome:					
1: JOBDEF/STEP040/**DSNUTILB**					
Average CPU/IO, MSec	6.6	42.2	35.6	540.8%	WWCB053-W Large variance fr
Observed MIPS/phys CPU	**57.8**	**9.0**	**-48.8**	**-540.8%**	
Rec count	3	3	0.0	0.0%	
Total EXCP count	129	9	-120.0	-93.0%	
Transactions	3	3	0.0	0.0%	
EXCPs per transaction	**43.0**	**3.0**	**-40.0**	**-93.0%**	**WWCB085-W Large change in I**
Average Users	3	3	0.0	0.0%	
Total CPU time, seconds	0.8	0.4	-0.5	-55.3%	WWCB055-I Small CPU usage
Coefficient of variation	2.0	4.6	2.5	123.7%	

DSNUTILB & Capacity Analysis

- But here's what also happened in that same run:

BOXSCORE/BATCH 10:35 Friday, October 29, 1999

V1R2d (c) Watson & Walker, Inc.

Programs With Best Outcome
H5/G5 UPGRADE

Item Description	9/27-9/30	10/4-10/7	Delta	% Delta	Comments
Job/Step/Pgm With Best Outcome:					
1: JOBXYZ/STEP08/**SYNCSORT**					
Average CPU/IO, MSec variance..	37.8	2.2	-35.6	-94.2%	WWCB053-W Large
Observed MIPS/phys CPU	**57.8**	**988.6**	**930.8**	**94.2%**	
Rec count	4	4	0.0	0.0%	
Total EXCP count	37	4344	4307.0	11640.5%	
Transactions	4	4	0.0	0.0%	
EXCPs per transaction	**9.3**	**1086.0**	**1076.8**	**11640.5%**	**WWCB085-W Large change..**
Average Users	4	4	0.0	0.0%	
Total CPU time, seconds	1.4	9.6	8.2	586.4%	WWCB055-I Small CPU usage
Coefficient of variation	6.6	1.1	-5.4	-82.6%	

DSNUTILB & Capacity Analysis

- Many people, including IBM, use a CPU per I/O analysis to confirm the capacity of a new machine. Here are some things to watch out for when using that technique:
  - watch for changes in the number of I/Os
  - watch for the same programs exhibiting the same reactions
  - be careful to do the analysis at the program level instead of period level that has multiple programs in it
  - look at all programs, not just a handful of pre-selected programs
  - realize that an increase in storage may may a processor look faster or slower
  - find a way to identify programs using hardware data compression because they may change significantly
  - watch for application, especially tuning, changes

z900 Compression

- Hardware compression took a major performance hit starting with G4.

- From z900 Ivory Letter:
  - "Hardware data compression on the IBM z900 (Freeway) has been enhanced to provide significant performance gains over IBM 9672 G6 (turbo) processors. Depending upon the length of the records: a 3 to 4 times improvement has been demonstrated for compression and a 2 to 3 times improvement has been demonstrated for expansion. For both expansion and compression, the longer the record length, the higher the performance gain. With expansion, records with lower compression ratios have higher performance gains".

z900 Compression

- Marty Wertheim at Bank of America wanted to look into compression comparing the z900 with a 9672 - Used IDCAMS to copy 2000 cylinders of data into a sequential data set. The compression rate was 60% for this data set. Results (CPU times):

9672	Compressed	**107.7 seconds**
	Not compressed	12.4 Seconds
	Ratio	8.68
z900	Compressed	**34.4 seconds !!!!**
	Not compressed	10.1 Seconds
	Ratio	3.41

- Conclusion: If you've seen sequential data sets that could benefit from compression but the overhead was too high, it may be worth the time to look at these again. Wait until all the CPU's in a node are z900s before implementing in production.

z/OS or OS/390 R10?

- Do you remember seeing the IBM slides that show that you can follow one of two paths for software and hardware?
  - OS/390 Rx on a G5/G6 ...to... OS/390 R10 on a G5/G6 ...to... OS/390 R10 on a z900 (31-bit mode) ...to... OS/390 or R10 z/OS on a z900 (64-bit mode)
    - or
  - OS/390 Rx on a G5/G6 ...to... z/OS on a G5/G6 ...to... z/OS on a z900 (64-bit mode)
- Some speakers make it seem like there's no difference - BUT THERE IS!

z/OS or OS/390 R10?

- z/OS on a z900 can *only* run in 64-bit mode, but OS/390 R10 on a z900 can run in either 31-bit or 64-bit mode.

- If you go to R10 on the z900 *before* installing z/OS, you can test 64-bit with no risk since you can always back out. z/OS on a z900 has no backout position. IBM may be ready for 64-bit, but do you have ALL of the PTFs from your ISVs to support 64-bit?

- I've seen many sites have problems with this, so I always recommend installing R10 on a z900 before z/OS.

How Much CPU O/H in R10?

- Frequent question:
  - Q - What's the overhead for going from OS/390 Rx to R10 or z/OS 64-bit mode? How about new releases of CICS and IMS?
    - A - It depends on: how much paging you're doing; how much storage you have; which workloads are you running; what release of the subsystems are you running; do you have any storage constraints now?
    - You can check with your IBM account rep, and maybe get CP2000 run. Or you can use a new, free, tool that WSC has just released - SoftCap V2.2 (it was an internal tool).
  - Download it from: **www.ibm.com/support/techdocs** (should be available soon)

Software Pricing

- You're probably aware of the software pricing announcement from this week announcing the delay of ILM on z/OS and the temporary replacement to due sub-capacity pricing from reports.
- Look at the ILM Web site for details of the July 23 announcement (but there is no IBM announcement letter). Also sign up for ILM listserver. **www.ibm.com/servers/eserver/zseries/wlc_lm**
- Announcement also indicated that there is sub-capacity pricing for 'Value Unit' priced data management tools. What are these? They'll be priced at the capacity of the z/OS capacity.

Software Pricing

- On March 6, 2001, announcement 201-057 was released. Title is 'IBM Introduces Value Unit Based Charges for Data Management Tools'. This was introduced with little fanfare - more like "stealth marketing". Tivoli had a press release using similar terminology.
- There is a list of 50 to 70 products that used to be priced features of other products, but are now individual products in their own right. Examples are: DB2 Performance Monitor and IMS Performance Monitor. (Find the full list from the ILM Web site.)
- MSU pricing is gone for these products; new Value Unit Pricing is in use effective March 2001.

Software Pricing

- How to calculate Value Units:

	MSUs	Value Units/MSU
Base	1-3	1.00
Tier A	4-45	.45
Tier B	46-175	.36
Tier C	176-315	.27
Tier D	316+	.20

- Example: 200 MSUs (e.g. 2064-106, 1160 MIPS) =

(3 * 1.00) + (42 * .45) + (130 * .36) + (25 * .27) = 75.55 = 76 Value Units

Software Pricing

- What's the result?
  - I don't know. I did the pricing for one product - DB2 PM. Value Unit pricing is 15% less than MSU pricing for both the one-time-charge and the S/S (Service & Subscription, i.e. annual charge).
- Why did they do this?
  - Probably to be more competitive with ISV products, especially when sub-CEC capacity pricing starts to be exploited.
- What does this mean to you?
  - This will take more work to determine whether you can save with sub-CEC capacity; but you could end up saving money.

Mysteries

- Oracle V8I (OSDI)
  - After installation, the system capture ratio dropped to 50%. In a stress test with only Oracle running, the CPU busy was 100%, but the combined TCB/SRB time was 20-40%. Anybody else see this?
- JES3 CPU Drops
  - After installation of JES3 for OS/390 R10, user noticed a 35% decrease in JES3 CPU time. Then they found it in the CONSOLE address space. No firm answer from IBM, but it appears that some functions moved to CONSOLE. Just be aware of it.

WLM Notes

- Compat mode disappears in 1Q02 - yeah!
  - The EWCP sessions on WLM migration reminded me that there are still sites that need to prepare for conversion to goal mode by getting their IPS/ICS structures organized. An article in my newsletter in 1995 addressed those issues. I posted it on my Web site after last SHARE. Wait for a week, then see the article on positioning for goal mode at: **www.watsonwalker.com/articles.html**
  - Come on, folks! It typically takes less than a day or so to get a sample policy running in your test LPAR. Everyone should be at least at this point. Use my QuickStart Policy, RMF's migration tool, or WSC's sample policy to get started. It's easier than you think!

WLM Notes

- I/O Priority Management
  - When I/O priority management is turned on, the velocity calculation is calculated as:

$$\frac{\text{(CPU Using + I/O Using)}}{\text{(CPU Delay + CPU Using + Storage Delay + I/O Using + I/O Delay)}}$$

  - Where:
    - I/O Using = Connect time + disconnect time
    - I/O Delay = Pend time + IOSQ

WLM Notes

- I/O Priority Management
  - This made no sense to performance analysts who have always believed that disconnect time (seek time and RPS misses) is a delay.
  - Because of the calculation, there were several problems: velocities were quite variable; the poorer the DASD response times, the more that jobs looked like they were running better than expected and they'd get knocked down by WLM (getting the reverse effect of what's needed); the amount of disconnect time was disproportional to other delays and dominated the velocity calculation.
  - I always recommended that people NOT turn on I/O priority management. So did WSC. See my TUNING Letter, 2000 No. 5, pg 55.

WLM Notes

- I/O Priority Management
  - Now with dynamic CHPID management, you MUST use I/O priority management.
  - So new APAR, **OW47667**, changes the calculation to omit disconnect time from the I/O USING. Will close next week, R8+.
  - Apply the APAR, recalculate your velocities, and THEN you can turn on I/O Priority Management.
  - Related APAR:
    - RMF APAR **OW43954** - "High Disconnect Times May Occur For Devices That Present Large Device Active Times"; this causes Ex Vel to show as 100%. (RMF R6-R9, 8/8/00).

Miscellaneous Notes

- z/OS will soon be available on Fundamental Software's FLEX-ES! In my last newsletter, I discussed it but said that only R10 was supported. That means that developers and small sites can be running z/OS on either a PC or UNIX (like) machine. Whoopee!

- My favorite sessions this week were given by **Glenn Anderson** on how to set up and measure a Webserver on OS/390. Look for sessions 2599 & 2547 when SHARE publishes them.

Miscellaneous Notes

- **Jerrie Stewart's** session 2412 stated that z/OS 1.4 will run on z900, G5/G6, and Multiprise 3000. That means that these machines are safe for awhile, but I also means that there could be another architectural level set any time from March, 2003 on. (I.e. when the G5/G6 won't be able to run the newer operating systems.)

Good Reading

- Keynote Speaker - **Dave Barry**
  - To see some of his columns, see:

    www.miami.com/herald/special/features/barry/

- Online manuals:
  - **Bob DuCharme's** "The Operating Systems Handbook," a crash course in UNIX, OS/400, VMS, MVS, and VM/CMS is available for free download:

    www.snee.com/bob/opsys.html

- Old Computers:

  www.obsoletecomputermuseum.org

Good Reading

- Neat Redbooks at **www.ibm.com/redbooks**
  - **SG24-6129** - DB2 for z/OS and OS/390 Version 7 Performance Topics
  - **SG24-5351** - DB2 UDB for OS/390 V6 Performance Topics
  - **SG24-5952** - z/OS Intelligent Resource Director

- DFSORT Changed URLs:

  www.ibm.com/storage/dfsort

  www.ibm.com/storage/software/sort/srtmhome.htm

APARs

- OS/390 R10 and z/OS - base, RMF, WLM
  - **OW48782** - HIPER, A.S. not swapping in after system has been IPLed for 51 days, 5/11/01, r10+
  - **OW49509** - HIPER, Recursive page faults and dataloss, PE, 6/19/01, R10 in 64-bit
  - **OW49322** - HIPER, Preemption Improvement, 5/10/01, R10
  - **OW49617** - HIPER, z/architecture mode high paging logically swapped users being physically swapped out, 7/3/01, R10+ in 64-bit
  - **OW46476** - RMF new function to support new WLM features, SMF type 79/2 new fields documented in z/OS but not R10, 12/12/00, R10+

APARs

- OS/390 R10 and z/OS - base, RMF, WLM (cont.)
  - **OW50358** - HIPER, IEEVWAIT + x'348' waiting while holding SYSIEFSD, OPEN, R10+
  - **OW47887** - Increased paging due to high protective processor storage targer, 3/27/01, R10+ in 64-bit
  - **OW47959** - HIPER, ARC0387A is issued & recover does not process when recovery is from alternate tape in disaster mode, 6/22/01, DFSMS R10+
  - **OW49980** - HIPER, Random ESQA, ECSA, SQA, CSA, overlay with blanks, 7/17/01, R10+
  - **OW49480** - HIPER, Enqueue lockout SYSIEFSD.VARYDEV & SYSZIOS.SWITCH are held, lockout at IPL, 7/13/01, R10

APARs

- OS/390 R10 and z/OS - DFSMS, IOS
  - **OW49094** - Isn't HIPER, but should be for capacity planners & performance analysts!, SMF15NTR is zero with R10, SMF30BLK negative, 5/10/01, DFSMS for R10. The following fields are also bad: SMF15NER, SMF30DCT (device connect time). SMF30BLK negative value looks like millions and millions of EXCPs!
  - **OW49672** - HIPER, Slower throughput processing batch jobs and TSO ids making SVC 25 catalog locate requests, 6/6/01, DFSMS for R10; tuning GRS and MIM can help (see my TUNING Letter, 2001, No. 2, pg 9)
  - **OW49736** - HIPER, ABEND878 RC 10 in SMSVSAM A.S. Excessive SP 229 key 5 storage used after OW4222 applied, VSAM RLS, 7/17/01, DFSMS R10

APARs

- OS/390 R10 and z/OS - DFSMS, IOS
  - **OW45559** - New function - add vector table support for CICSVR, 11/02/00, DFSMS R10 - got wrong one - should be OW46559
  - **OW49546** - HIPER, ANTQMGR loop hangs concurrent copy jobs & memterm processing, 7/13/01, DFSMS 1.4+
  - **OW50098** - HIPER, Waits & slowdowns due to enq contention using extended format data sets, 6/28/01, DFSMS R10
  - **OW48995** - SP 231 Key 5 ECSA Storage grows using PDSEs, 5/17/01, DFSMS R10+
  - **OW44428, OW44429, OW47519** - New function for FICON Director, documentation for SMF 74/7 not in manuals, 12/15/00 & 3/1/01, R10+

APARs

- TCP/IP, IOS, DFSMS, GRS
  - **OW49151** - HIPER, GRS contention exit may leave enq request hung, GRS=STAR, 7/3/01, R6+
  - **OW50495** - ECS waiting for IOSAS, IPL takes too long, after OW48879, OPEN, 7/23/01
  - **PQ46754** - HIPER, TCP/IP traffic stops on a connection, 5/3/01, TCP/IP V3
  - **OW49640** - HIPER, Corrupt VSAM RLS KSDS, 7/18/01, DFSMS
  - **OW38548** - New function SRM support for Fibre channels, added new IEAOPTxx parm, CPMF, for FICON but not documented until z/OS; chg to SMF 73 records, 7/15/99, R3-R8
  - **OW48166** - HIPER, Using PDAS or TDMF to move catalog under ECS control may result in broken catalog, 7/12/01, DFSMS

APARs

- TCP/IP, IOS, DFSMS, GRS
  - **PQ44379** - New function - performance improvement for CISCO routers, 12/6/00, TCP/IP V3
  - **II11334** - Great TCP/IP info APAR
  - **PQ30143, PQ33474** - ABEND878-10 or ECSA exhausted due to a telnet client flooding telnet, high CPU, added new parm MAXVTAMSENDQ to prevent hot I/O, 11/04/99 & 12/30/99, TCP/IP V3 & OS/390 R5
- Base, WLM, RMF, JES2
  - **OW49125** - WLM - BPXAS inits not marked as privileged, some high priority work (like FTPs) could fall into default started tasks, 7/18/01, R6+

APARs

- Base, WLM, RMF, JES2
  - **OW49428** - ASCBTCBS & ASCBTCBE counts invalid after enclave TCB suspended for CMS lock, could cause CPU overhead in LPAR management or uncaptured, especially in DB2 sites, similar to OW44517 from last year, OPEN (target 10/3/01), R6+
  - **OW47328** - HIPER, JES2 New function, Fixes problem that can cause hangs in JES2, new function for allowing multi-JES levels in MAS, 5/15/01, JES2 R7+
  - **OW48934** - HIPER, Requests to CF lock structure hung, 7/18/01, R6+
  - **OW49165** - HIPER, WLM doesn't select jobs due to incorrect info from JES2, 5/10/01, JES2 R5+
  - **OW49716** - HIPER, Various RMF problems, several abends noted, 7/12/01, R6+

APARs

- Base, WLM, RMF, JES2
  - **OW49773** - HIPER, RMF cache data missing UCBDCE fields incorrect after UW77972 Shark D/T2105, 7/11/01, R3+
  - **OW44297** - HIPER, When two members IPL simultaneously, signalling connectivity is not achieved, 7/25/01, R6+
  - **OW48624** - Application performance impacted when a system is partitioned from parallel sysplex, 7/10/01, R8+
  - **OW50081** - HIPER, Enq contention after T MMS or D MSS command issued, OPEN, R8+
  - **OW50107** - HIPER, High CPU after $T JOBCLASS with MODE, OPEN, JES2 R5+

Questions

- **Email: cheryl@watsonwalker.com**
- **Web site: http://www.watsonwalker.com**