Cheryl's Hot Flashes #16

Cheryl Watson

August 18, 2006, Session 2509

Watson & Walker, Inc.

www.watsonwalker.com

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

Agenda

- Survey Questions

- MIPS, MSUs, Service Units

- ROTs (Rotting and Otherwise)

- User Experiences

- Interesting APARs

- Downloads

- 6-Month Update

- This SHARE

Survey Questions – Hardware

- Current Server Type (now or within next 12 months)
    - z800, z900, z890 (*80% at last SHARE*)?
    - z990 (*75%*)?
    - z9-EC (*30%*)?
    - z9-BC?
    - Older Hardware (*15%*)?

- Using zAAP Processors (*10, 15-20 at last two SHAREs*)?

- Using zIIP Processors (*15-20*)?

- Activated IRD CPU Management (*15, 15*)?

- Have Used On/Off Capacity on Demand (*12*)?

- Using Variable WLC Pricing *(25, 25-30%)?*

- Doing Heavy Cryptographic Work (*10, 10*)?

Survey Questions – Software

- Operating System
  - z/OS.e *(5, 6)*?
  - z/OS 1.4 or 1.5 (*90%, 90%*)?
  - z/OS 1.6 (*20, 75%*)?
  - z/OS 1.7 (20%)?
  - Earlier than z/OS 1.4 (*12, 3*)?
  - Note: End of Service for z/OS 1.4 & 1.5 is March 2007

- Using WebSphere on z/OS (*35, 45%*)?

- Using VSAM RLS for CICS (*12, 15*)?

- Using Transactional VSAM (*0, 1*)?

- Have used zPCR (*10*)?

- Debug dumps with IPCS?

MIPS, MSUs, Service Units

- Some common questions we're asked:

    - How do I convert from MSUs to MIPS?

    - How do I convert from Service Units to MIPS?

    - Should I use MIPS, MSUs, or Service Units?

    - How can I normalize CPU time for chargeback?

- Answer: IT DEPENDS!

MIPS

- Where do MIPS come from?
    - At some point during the announcement of a new line of processors, IBM makes a statement such as "equivalent to a 450 MIPS machine," or "1.5 times the capacity of the xxx machine" (where the xxx machine had been 450 MIPS)
    - IBM publishes LSPR (Large Systems Performance Reference) rates for different workloads that give the relative capacity of a machine compared to a base machine.
    www.ibm.com/servers/eserver/zseries/lspr
    - These LSPRs can be changed when the software changes (e.g. there are LSPRs for z/OS 1.4 and also for z/OS 1.6)
    - Analysts use the average rates and the IBM statements to produce MIPS tables[1]

[1] Most analysts use the "average", but we publish all workloads

MIPS

- Example:

  - The 2084-301 was announced as a 450 MIPS machine
  - The z/OS 1.4 LSPR for a 2084-314 says that the ratio between a 301 and 314 is:
    - Average = 10.29 = 4630 MIPS
    - Low I/O = 11.49 = 5170 MIPS
  - But the z/OS 1.6 LSPR for a 2084-314 in single-image mode produces the following:
    - Average = 10.98 = 4941 MIPS
    - Low I/O = 11.51 = 5193 MIPS

MIPS

- Example (continued):
    - And the z/OS 1.6 LSPR for a 2084-314 in multi-image mode lists the following:
        - Average = 10.30 = 4635 MIPS
        - Low I/O = 10.75 = 4838 MIPS
        - Also, the 301 in multi-image mode is now rated as 428 MIPS instead of 450 MIPS
    - So MIPS for the 314 range from 4630 to 5193 (563 MIPS variation or 12%)
    - For most sites, the z/OS 1.6 multi-image low I/O value is the closest to reality (4838 MIPS)

Service Units

- Where do CPU Service Units come from?
  - Although IBM doesn't specify, these appear to be derived from the average MIPS (because the ratios haven't changed much in years)
  - These are published by IBM at www.ibm.com/servers/eserver/zseries/srm
  - z/OS 1.6 single image shows an average of 47.4 service units per MIPS (max = 54.3, min = 46.7)
  - z/OS 1.6 multi-image shows an average of 49.8 service units per MIPS (max = 54.6, min=45.5)
  - z/OS 1.4 shows an average of 48.8 service units per MIPS (max = 54.3, min = 46.6)
  - Non-zSeries machines used to be about 52.5 service units per MIPS

Service Units

- ## How are CPU service units used?
    - SRM uses the service units per second value specified for the LPAR view, not the CEC view and multiplies it by the CPU time and the TCB/SRB service definition coefficient
    - So if you are running a 3-way LPAR on your 314 machine, the service units are from a 303, not a 314
    - In our example, a 314 has a value of 15640.2737 su/sec and a 303 has a value of 20075.2823. That's a 28% difference. Due to LPAR overhead, the work will actually experience closer to the lower value. That is, it will appear that work takes more resources on a smaller LPAR

Service Units

- Cautions:
  - Be sure to only use the CPU components of service units when using SMF data
  - Be sure to back out the service definition coefficient if you don't use a value of 1.0
  - I would convert all service units to use the CEC value

- My wish:
  - IBM should remove the service definition coefficients altogether and should use the CEC service unit rate for all internal calculations

MSUs

- ## Where do MSUs come from?

  - Originally, MSUs were used to replace model groups for software pricing and were calculated from service units. The calculation for a machine was:

    ((su_per_sec * no_of_cps) * 3600) / 1000000

  - Example: z900-101:

    - ((11585 * 1) * 3600) /1000000 = 41.7, MSUs are 41

  - Originally, there were between 5.8 and 6.0 MIPS per MSU; the z900 had more variation and ranged from 5.5 to 6.8 MIPS per MSU

- ## These are published by IBM at:

  www.ibm.com/servers/eserver/zseries/lspr

MSUs

- How are MSUs used?
    - Now they're controlled by marketing in order to provide reduced software prices
    - The original MSUs were called hardware MSUs
    - The z800s came out with hardware MSUs (ranging from 5.1 to 5.9 MIPS/MSU); followed by decreased pricing (software) MSUs (these were quite varied and ranged from 5.1 to 8.5 MIPS/MSU)
    - The z990s followed with hardware MSUs of about 6.5 MIPS/MSU and software MSUs of about 7 MIPS/MSU
    - The z9s range from 7.1 to 8.0, with an average of 7.6 MIPS/MSU (only software MSUs are published, no hardware MSUs are available)

MIPS, MSUs, Service Units

- Now, how about those questions?

- How do I convert from MSUs to MIPS?
  - You can't – the ratio varies from 5.1 to 8.5 MIPS per MSU – it's too inconsistent (but higher ratios produce better price/performance due to software pricing)

- How do I convert from Service Units to MIPS?
  - You're probably safe using 49-50 service units per *average* MIPS, unless you're running in an LPAR that uses less than the total number of CPs, where that number could be 50% off!
  - Remember that this is based on the *Average*, not the low I/O that you might be using for sizing!

MIPS, MSUs, Service Units

- Should I use MIPS, MSUs, or Service Units?

  - This REALLY depends!
  - Realize that there is **no answer** that really works – this isn't a science, it's an art
  - If you want to do chargeback or capacity planning by job, then you need to use service units; but if running on a small LPAR (i.e. using less than all CPs) then you should first convert to use the CEC service units (not in the RMF records)

MIPS, MSUs, Service Units

- Should I use MIPS, MSUs, or Service Units (cont.)

  - If you are doing any type of sizing of the machine for upgrades or capacity planning, then don't go there alone – use zPCR or one of IBM's service offerings; there are too many "gotchas"

  - Because most managers want to use MIPS, then use the MIPS from zPCR after you define your configuration; for zPCR, you can use 450 MIPS for a 301 as a base machine

- Our belief:

  - The only MIPS you should use should come from zPCR. www.ibm.com/support/techdocs/atsmastr.nsf/ WebIndex/PRS1381

MIPS, MSUs, Service Units

- More questions:

  - How can I normalize CPU time for chargeback?

    - As before, use service units, but change them to use the CEC service unit value:
      - *Modify reported CPU service units by dividing them by the LPAR su/sec and multiplying them by the CEC su/sec*

Rotting ROTs

- For years performance and capacity planners have used Rules of Thumb (ROT) for easy tuning and planning

- Are those old ROTs still valid?

- Do we still need ROTs?

- Let's examine some that are getting a bit shaggy, some new ones that can be set, and an EWCP project that's coming

Rotting ROTs – Paging

1. Paging

   - Old ROT – Keep page fault rate less than 20 (non-cache) or 15 (cache) pages per second

   - Now – Paging should be almost zero due to increased real storage (a blip of 10 pages per second should come as a surprise)

Rotting ROTs – DASD

2. DASD

- Old ROT – Keep average DASD response time less than 20 ms

- Now – Keep average DASD response time less than 3 ms due to improved caching, devices, ESCON, FICON, PAVs, etc.

Rotting ROTs – CPU Busy

3. CPU Busy (1)

- Old ROT – Keep average CPU busy less than 80% in order to provide consistent good response time to online users

- Now – Run machine at 100% busy as often as possible. WLM can ensure that online systems get good response times while adequately managing the other workloads. Remember that unused CPU time can never be regained.

Rotting ROTs – CPU Busy

4. CPU Busy (2)

- Old ROT – For trending or capacity planning, use average CPU Busy

- Now – Use total CPU Busy – CPs can be added or deleted dynamically. An average CPU busy of 33.3% on three CPs is the same as 50% busy on two CPs or 100% busy on one CP. But they all have a total CPU busy of 100%.

Rotting ROTs - Upgrades

5. Upgrades

- Old ROT – When adding resources (storage and CPU) during an upgrade, distribute all of the new resources among the LPARs

- Now – Only add resources as needed. If you add more memory than you need, you'll simply cause CPU overhead. If you add more CPU power than you initially need, you'll change the users' expectations and they'll never want you to take the resources away.

Rotting ROTs – SMF Interval

6. SMF Interval

- Old ROT – Use 15 minute intervals to get detailed data

- Now – Use 30 minute intervals to reduce SMF overhead. This will tend to hide some spikes that could be seen in 15 minute intervals, but there are probably more important things to work on and better uses for your DASD and nighttime processing.

Rotting ROTs – Java

7. Java

- Old ROT – Don't use Java on the mainframe because it's slow as molasses.

- Now – With latest SDK, and especially with zAAP engine, Java can be competitive with other platforms; much of WebSphere uses Java. Caution: Be VERY current with SDKs and maintenance! Watch the HIPER APARs.

ROTs – Some New Ones

- ## WLM ROTs (e.g.)
  - PIs over 1.0 for importance 1 and 2 work are exceptions and should be dealt with
  - PIs over 1.5 for low importance work are exceptions and should be dealt with
  - System and SYSSTC work shouldn't take over 15% of the system[1]
  - Velocities should vary by at least five, if not ten to be effective

      Thanks to Peter Enrico (see page 38)

1

ROTs – Some New Ones

- Other Areas:
    - Networking
    - TCP/IP
    - Coupling Facility measurements
    - IRD
    - zIIPs, zAAPs, ICFs
    - LPARs
    - Etc., etc., etc.

Rotting ROTs – Still Valid?

- Are ROTs still valid?
  - Yes, because the only way that you can manage a system is by exception, and you need to know what defines an exception
  - Start with an industry ROT and modify it
  - Best way to tune is to find the Uglies! Find the worst ten DASD devices, the ten largest CPU jobs, the peak interval of the day, the ten largest storage users, the ten heaviest CICS transactions, etc., etc.
  - Set ROTs in your monitors and let them tell you when there's a problem

Rotting ROTs – Still Valid?

- From Norman Hollander:
    - "The best monitor is a blank screen."
        - That says it all. You really don't need dials and graphs. All you really need to know is whether anything is exceeding your guideline.

- Exception Processing
    - The only way to go.
    - But how to find the ROTs?

Rotting ROTs – Future

- Finding ROTs
  - Start with your top ten lists (understand your worst DASD response times, CPU busy times, etc.)
  - Come to SHARE and obtain ROTs from speakers
  - Join the EWCP project – send an email with your name, company name, and email address to me at cheryl@watsonwalker.com and we'll add you to the list. We'll soon be starting a project to create an industry-accepted list of ROTs. You can provide input or simply watch the progress.

User Experiences – IMS

- From Janet Howie of Nationwide:

  - *"We migrated from z/OS 1.5 to z/OS 1.7 and received a 25% response time improvement in our IMS Version 9 system that runs 2.5 million transactions per day.*

    "We run many WLM-managed DB2 stored procedures under IMS in a DB2 data sharing environment, so I'm thinking most of the benefit is coming from WLM's new management algorithms in z/OS 1.7. IMS DC Monitor traces show at least a 30% reduction in the average time for DB2 "Normal Calls."

    "I don't think I have seen this much improvement in IMS response for over 20 years!"

Interesting APARs

- **OA16302** (DFSMS for z/OS 1.6+, 24May2006)
  - *PDSE to PDSE Member Copy Output is Truncated Under DFSMS 1K0 and DFSMS 1J0 (HDZ11K0 z/OS V1R7 & HDZ11J0 z/OS V1R6).* Doing a PDSE to PDSE member copy with IEBGENER results in truncation of the output member, even though the program executes without error. (IEBCOPY works okay)

- **OA12857/OA12822/OA12861** (DFSMS for z/OS 1.6+, 22May2006)
  - *PDSE Performance - Put Member Caching Statistics in SMF Type14/Type15 Records.* This new function updates the SMF Type 14 and Type 15 records to include a new section containing performance information for PDSE data sets. These statistics help you calculate the effectiveness of the PDSE cache for specific data sets. Unfortunately, if multiple PDSE data sets are concatenated in the same allocation, the counts reflect the totals for all data sets.

Thanks to Jerry Urbaniak of Acxiom

Interesting APARs - Red Alerts

- ## IMS/VSAM (10Aug2006)
  - IMS/VSAM interface problem affecting IMS users of VSAM Hiperspace buffers
  - https://www14.software.ibm.com/webapp/set2/sas/f/redAlerts/home.html

- ## zIIPs (28Jul2006)
  - After installing JBB77S9 (z/OS 1.6 and z/OS.e 1.6 zIIP Web Deliverable) or JBB772S (z/OS 1.7 and zOS.e 1.7 zIIP Web Deliverable) users may experience the problem described by OA16005. This can lead to a multi-system outage.

    https://www14.software.ibm.com/webapp/set2/sas/f/redAlerts/home.html

- ## Red Alerts:
  - You can browse Red Alerts at www.ibm.com/servers/eserver/support/zseries/. From the main page, click on the link "Red Alerts for System z." Or go directly to www14.software.ibm.com/webapp/set2/sas/f/redAlerts/20060331.html. You can also subscribe to a service that will send you emails related to Red Alert activity.

Interesting APARs

- **FLASH10483 - Changes to Daylight Saving Time (DST) in the USA - 2007**
  - The new USA Energy Policy Act of 2005 has changed the dates that Daylight Saving Time (DST) will be observed in the United States, starting in 2007. This document lists the new dates, and the areas that will be affected (some states and regions do not plan to honor the new dates). Depending on the method used for setting the time, customers will have to change their procedures or install new microcode updates so that the time change occurs correctly. (30May2006)

- WSC Hints & Tips
  - TD103094 - Calculating LPAR Image Capacity (26Apr2006)
  - TD103105 - Using Workload (WLM) and DB2 Stored Procedures (SPAS) Some clarifying Questions and Answers (14May2006)
  - TD103128 - SMF Buffer Constraint Relief in z/OS 1.6 (23May2006)

Interesting APARs

- APAR OA14409 finally fixes MCCAFCTH defaults

  - APAR changes the default value of MCCAFCTH=(400,600) to 0.2% (low) and 0.4% (ok) of total real storage for the LPAR.

    If you've already coded overrides to IBM's 400,600 default values, you may be able to remove them after this APAR.

  - Thanks to Brian Peterson on IBM-Main

Interesting APARs & Downloads

- APAR OA16909
    - z/OS R1.7 users should have this APAR applied if using compression with VSAM data sets. Without it, IDCAMS uses strictly software to compress/decompress causing batch and online tasks to use a lot more CPU and run much longer. This also corrects a problem where BUFND is calculated incorrectly, leading to poor performance.

- 6 ½ years of SAS Coding Tips:

    - support.sas.com/sassamples/archive.html?ETS=4756&PID=123632

    - Thanks to Jerry Urbaniak of Acxiom

6-Month Update - Downloads

- From Hot Flashes at last SHARE:
  - OMEGAMON® z/OS® Management Console
    - No charge monitoring tool to look at Health Checker, system status, and configuration data
    - Designed for those new to z/OS
    - Announcement letter **205-329** on 13Dec2005
    - See the description in HOT Topics, found in your SHARE bag
    - It's *VERY COOL!*

- *Feedback from This SHARE:*
  - Takes 1 to 3 days to install – it's a headache!
  - "Output isn't worth the effort"

This SHARE

- Interesting Sessions
  - 2500, *"zOS Performance 'Hot' Topics"* by Kathy Walsh – latest in WSC information and performance APARs (always my favorite session at SHARE).
  - 2546, *"z/OS WLM SYSTEM & SYSSTC Service Classes"* by Peter Enrico – Great recommendations for use of these service classes.
  - 2852, *"A z/OS System Programmer's Guide to Migrating to a New IBM System z9 EC or z9 BC Server"* by Greg Daynes – Super handout (as usual) for the sysprogs.
  - 2892, *"One Last Drink from the Firehose: Migrating from z/OS R4 to R7"* by Marne Walle – The last migration recommendations for moving from z/OS 1.4 – Your last chance.

This SHARE

- IBM's zFavorites
  - Easiest way to find this page is to enter 'zfavorites' into Google and hit "I'm feeling lucky"

- Neat operator command for tuning XCF message sizes
  - D XCF,CD,CLASS=ALL
    - It shows the actual size of messages coming through; helps you properly size the transport classes
    - Only problem is that you need to be there; so try to schedule the message periodically and review syslog
  - For a description, see Joan Kelley's session 2523, ""*Parallel Sysplex Tuning Update*"
  - Also explained in WSC White Paper WP100743, "*XCF Performance Considerations*" by Kathy Walsh, last updated 22Jul2006

See You in Tampa!

- **Email: technical@watsonwalker.com**

- **Web site: www.watsonwalker.com**