# Hot Tips from Cheryl & Frank

*Cheryl Watson and Frank Kyne*

*Watson & Walker, Inc.*

[technical@watsonwalker.com](mailto:technical@watsonwalker.com)

*www.watsonwalker.com*

*August 7, 2014*

*Session 16251*
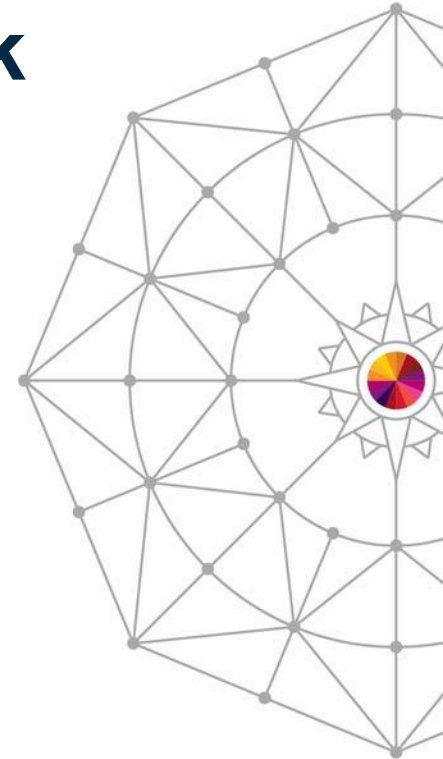
# Abstract

Cheryl Watson, best known for her newsletter on z/OS Performance (*Cheryl Watson's Tuning Letter*) will provide some of her favorite tips from her 23 years of writing. Our new associate, Frank Kyne, a long time author and editor of IBM's ITSO Redbooks, as well as a renowned teacher of the ITSO workshops, will provide some of his favorite tips from his 16 years of writing and teaching. Do you need z/OS courses or consulting? Come see and hear the people who teach these courses, and do the consulting. Everyone is guaranteed to take away something of value from this session. In Anaheim, we gave away copies of our Tuning Letter 2013 No. 2, and we'll have more copies for attendees in Pittsburgh. Our new, free SMF Reference Summary will be available, and a drawing will provide one lucky attendee with a DVD of the last 23 years of Tuning Letters. What's not to like?

Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval

# Agenda

- Who we are
- Hot tips from Cheryl
- Frank's lukewarm tips
- Drawing for DVD with with 23 years of *Tuning Letters*

# Who we are

- Watson & Walker, Inc. – A SMALL Business!
    - Tom Walker, CEO and Publisher
    - Cheryl Watson (Walker), President, teacher, author, consultant
    - Frank Kyne, Editor, teacher, author, consultant
    - Linda May, Business Manager
    - Consultants: Scott Barry of SBBWorks and others
    - Recognized ISV and IBM Partner in IBM PartnerWorld
    - Loving what we do every day!

# Watson & Walker, Inc. – A SMALL Business!
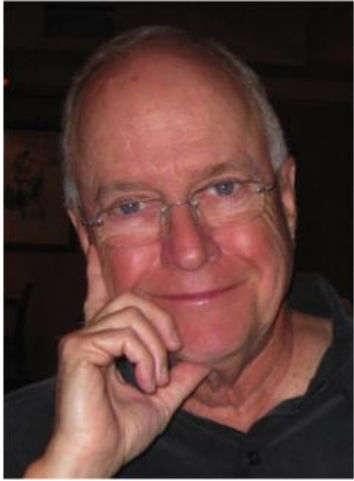
Tom Walker

Cheryl Watson

Frank Kyne

Linda May

# What we do

- *Cheryl Watson's Tuning Letter*
  - Began in 1991 as a 20-page monthly print newsletter.
  - Now sent via PDF with one subscription per physical site (special prices for multiple sites). Also includes yearly DVD with all issues, presentations, and CPU Charts since 1991.
  - Now published quarterly, but still providing over 240 pages a year with advice on z/OS. (Over 5,000 pages!) We should note that halfway through 2014, we already have 276 pages!
  - Unique in that we have no advertising and we give our own advice, rather than simply stating facts.
  - We aren't asking for your email address, but if you'd *like* to keep track of what we're doing, please sign up for our mailing list. Go to our website, select Cheryl's List archives to see the previous mailings; then sign up if you're interested. Cheryl's List is sent out about once or twice a month.

# What we do

- *Cheryl Watson's Tuning Letter – the proposed topics for the next issue are:*
  - zBNA update
  - RLS-Managed catalogs
  - Final installment of Capping article
  - Relative performance between Java and COBOL
  - Warning Track Interrupts
  - What's New
  - SHARE Trip Report
  - News (APARs, Techdocs, Redbooks, Blogs, etc)

# What we do

- *Cheryl Watson's CPU Chart*
  - Most comprehensive CPU Chart containing 1,077 processor configurations
  - Included as part of regular subscription of Tuning Letter; or available to ISVs for multiple sites as a stand-alone
  - Also used in our BoxScore product
  - For every processor, we provide 42 metrics, either from IBM published data or our own calculations
  - Many contracts are set up using our CPU Chart as the source for MIPS values

# What we do

- *Cheryl Watson's CPU Chart*
  - Example:

## Cheryl Watson's z/OS 1.13 Multi-Image CPU Chart

| Model | # of CPs | IBM PCI | Avg RNI MIPS | Avg RNI MIPS/ CP | MP | Low RNI MIPS | Low-Avg RNI MIPS | Avg-Hi RNI MIPS | High RNI MIPS | SU/Sec | UP SU/Sec | Common Name | Proc Grp | S/W MSUs | MIPS/ S/W MSU | H/W MSUs | MIPS/ H/W MSU | STIDP | STSI Model | ALS | GA |
|-------|------|------|--------|-------|------|--------|--------|--------|--------|-----------|------------|-------|------|-----|-----|-----|-----|------|-----|-----|-------|
| 2094-401 | 1 | 193 | 193.1 | 193.1 | 1.00 | 193.1 | 193.1 | 193.1 | 193.1 | 10107.3910 | 10107.3910 | z9-EC | 60 | 28 | 6.9 | 36 | 5.3 | 2094 | 401 | 7 | 05/06 |
| 2094-402 | 2 | 375 | 374.6 | 187.3 | 0.97 | 381.7 | 378.1 | 369.7 | 364.9 | 9708.7379 | 10107.3910 | z9-EC | 80 | 55 | 6.8 | 70 | 5.4 | 2094 | 402 | 7 | 05/06 |
| 2094-403 | 3 | 554 | 554.1 | 184.7 | 0.96 | 566.0 | 560.0 | 542.3 | 531.0 | 9433.9623 | 10107.3910 | z9-EC | 80 | 80 | 6.9 | 102 | 5.4 | 2094 | 403 | 7 | 05/06 |
| 2094-404 | 4 | 732 | 731.8 | 182.9 | 0.95 | 746.1 | 738.9 | 713.0 | 695.1 | 9195.4023 | 10107.3910 | z9-EC | IMLC | 103 | 7.1 | 132 | 5.5 | 2094 | 404 | 7 | 05/06 |
| 2094-405 | 5 | 897 | 897.0 | 179.4 | 0.93 | 922.0 | 909.4 | 874.7 | 853.4 | 8958.5666 | 10107.3910 | z9-EC | IMLC | 126 | 7.1 | 161 | 5.6 | 2094 | 405 | 7 | 05/06 |
| 2094-406 | 6 | 1056 | 1056.0 | 176.0 | 0.91 | 1093.9 | 1074.6 | 1028.0 | 1001.4 | 8762.3220 | 10107.3910 | z9-EC | IMLC | 147 | 7.2 | 189 | 5.6 | 2094 | 406 | 7 | 05/06 |
| 2094-407 | 7 | 1209 | 1209.0 | 172.7 | 0.89 | 1261.8 | 1234.8 | 1173.4 | 1139.9 | 8565.3105 | 10107.3910 | z9-EC | IMLC | 166 | 7.3 | 216 | 5.6 | 2094 | 407 | 7 | 05/06 |
| 2094-408 | 8 | 1356 | 1356.1 | 169.5 | 0.88 | 1425.9 | 1390.1 | 1311.3 | 1269.3 | 8368.2008 | 10107.3910 | z9-EC | IMLC | 185 | 7.3 | 241 | 5.6 | 2094 | 408 | 7 | 05/06 |

| Model | # of CPs | Low RNI MIPS/ CP | Low-Avg RNI MIPS/CP | Avg RNI MIPS/CP | Avg-Hi RNI MIPS/CP | High RNI MIPS/CP | % Diff - Avg to Low | % Diff - Avg to High | Speed of SP | % Diff - Avg to SP | MSU in PS Mode | % Diff in PS Mode | z/OS 1.11 Avg MIPS | % Diff in 1.11 vs 1.13 | SRM Sec / Real Sec | z/VM High RNI MIPS | z/VM MIPS/ CP | z/VM % Diff From AVG RNI | Linux Low RNI MIPS | Linux MIPS/ CP | Linux % Diff From AVG RNI |
|-------|------|-------|--------|--------|--------|--------|-----|-----|-------|------|------|------|--------|-----|----------|--------|--------|------|--------|--------|------|
| 2094-401 | 1 | 193.1 | 193.1 | 193.1 | 193.1 | 193.1 | 0% | 0% | 559.8 | 190% | | | 190.3 | 1% | 234.9623 | 190.3 | 190.3 | -1% | 190.3 | 190.3 | -1% |
| 2094-402 | 2 | 190.9 | 189.1 | 187.3 | 184.8 | 182.5 | 2% | -3% | 559.8 | 199% | | | 375.1 | 0% | 234.9623 | 363.9 | 181.9 | -3% | 380.7 | 190.3 | 2% |
| 2094-403 | 3 | 188.7 | 186.7 | 184.7 | 180.8 | 177.0 | 2% | -4% | 559.8 | 203% | | | 554.2 | 0% | 234.9623 | 531.8 | 177.3 | -4% | 571.0 | 190.3 | 3% |
| 2094-404 | 4 | 186.5 | 184.7 | 182.9 | 178.2 | 173.8 | 2% | -5% | 559.8 | 206% | | | 733.3 | 0% | 234.9623 | 699.7 | 174.9 | -4% | 750.1 | 187.5 | 3% |
| 2094-405 | 5 | 184.4 | 181.9 | 179.4 | 174.9 | 170.7 | 3% | -5% | 559.8 | 212% | | | 895.7 | 0% | 234.9623 | 856.5 | 171.3 | -5% | 929.3 | 185.9 | 4% |
| 2094-406 | 6 | 182.3 | 179.1 | 176.0 | 171.3 | 166.9 | 4% | -5% | 559.8 | 218% | | | 1058.0 | 0% | 234.9623 | 1007.6 | 167.9 | -5% | 1108.4 | 184.7 | 5% |
| 2094-407 | 7 | 180.3 | 176.4 | 172.7 | 167.6 | 162.8 | 4% | -6% | 559.8 | 224% | | | 1209.2 | 0% | 234.9623 | 1147.6 | 163.9 | -5% | 1281.9 | 183.1 | 6% |
| 2094-408 | 8 | 178.2 | 173.8 | 169.5 | 163.9 | 158.7 | 5% | -6% | 559.8 | 230% | | | 1354.7 | 0% | 234.9623 | 1281.9 | 160.2 | -5% | 1449.9 | 181.2 | 7% |

# What we do

- *Cheryl Watson's Advisor Series* of software products
  - BoxScore
    - SAS-based product to compare the speed and capacity of a processor after a hardware change, software change, application tuning effort, or configuration change.
    - Shows how individual job steps and transactions behaved.
    - Used by both outsourcers and those outsourced to determine normalization factor for billing.
    - Used to verify IBM's zPCR estimates.
  - GoalTender
    - Provides Workload Manager (WLM) analysis to help tune your WLM policies.
    - Provides Cheryl's warnings and reports.

# What We Do

- Education and Training
  - First class – Last Sunday was *z/OS Parmlib: The Heart of z/OS – Updates in z/OS 1.13 and 2.1*.
  - We expect to repeat this in Seattle on the Sunday before SHARE.
  - Next public classes (tentative names) in Dallas:
    - December 8-10 – Exploiting z/OS To Save Money
    - December 11-12 – How To Save Money With IBM License Charge Options
  - Policy – we NEVER cancel a scheduled class (unlike some other training companies).
  - Benefit – classes are taught by the actual class developers.
  - We also provide on-site and tailored classes.

# What We Do

- Consulting
  - Parallel Sysplex issues
  - Workload Manager (WLM) issues
  - Performance analysis
  - Capacity planning
  - Outsourcing contract review and negotiation
  - Software pricing

# What We Do

- Our website - www.watsonwalker.com
  - Contains *Tuning Letter* subscription information, prices, order forms, sample issues, and table of contents for 23 years.
  - Downloadable presentations (SHARE presentations since 1998), published articles, and some free software.
  - New References page – currently contains our SMF Reference Summary (currently 75 record types and 600 subtypes); this is a work in progress as we expect more ISVs to provide information.
  - SMF 101 contains articles from TL that are directed to new sysprogs, and are free to the public.
  - Subscribe to free *Cheryl's List* and get emails (monthly or so). You can also see the archives with summary of each *Tuning Letter* and other information we think is important.
  - BoxScore and GoalTender information, brochures, pricing, and order forms.

Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval

A couple of photos taken by Cheryl while leaving Sarasota.
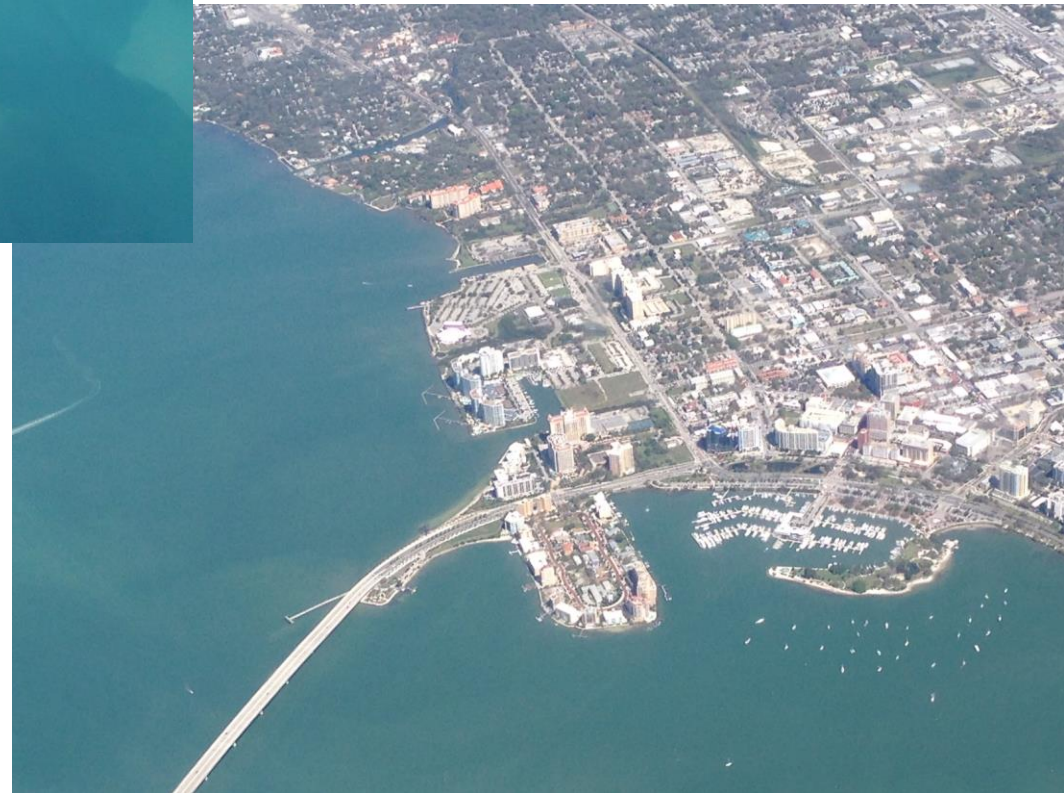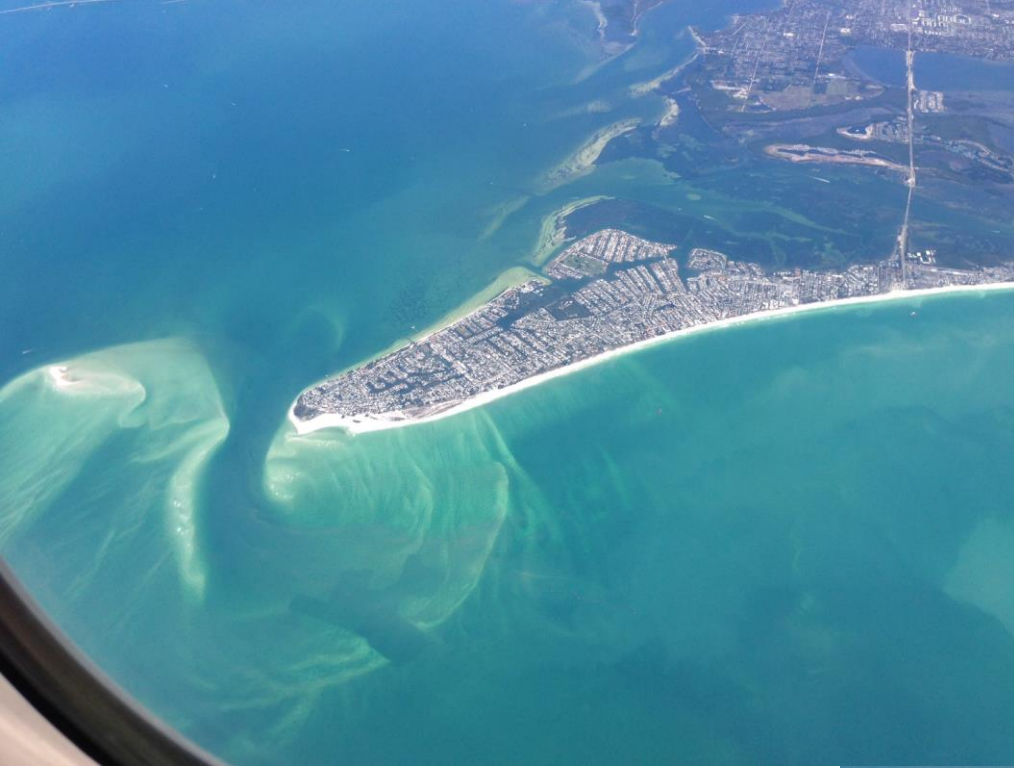
# Hot Tips From Cheryl

- Compiler enhancements in PL/I and COBOL
  - The C compiler has traditionally provided the ability to specify what CPC architecture the program will run on, letting you get the maximum benefit from architectural enhancements in CPC design.
  - But PL/I and COBOL traditionally did not provide the ability to specify what architecture level the resulting program would run on.
  - IBM Enterprise COBOL for z/OS 5.1 adds this ability, with the base level being for a 2084.
  - IBM Enterprise PL/I for z/OS 4.1 added this ability to PL/I

# Hot Tips From Cheryl

- Compiler Enhancements (cont.)
  - Some customers have experienced significant (>10%) CPU savings by re-compiling their COBOL and PL/I programs with an ARCH level that matches their installed CPC.  But:
    - <u>You cannot just recompile all your programs and expect identical results.</u>
    - Remember to allow for fallback situations such as if your DR site CPC is at an older architecture level than the production site.
  - For more information about migration considerations:
    - Refer to **Tom Ross**' COBOL sessions at this SHARE – sessions 16149, 16150, 16151, and 16152
    - **Chau Dickson's** PL/I session 16091

# Hot Tips from Cheryl

- Money Savings Tips
  - Reserve new capacity
    - After an upgrade, hold part of the new capacity back, especially if you want to retain consistent chargeback. IBM's LSPRs are based on running online work at 90% busy, batch work at over 90% busy, and mixed work at 99% busy.
      - If you don't, people may develop different techniques of working, and end up increasing the CPU usage.
    - Releasing all the capacity immediately means that jobs will cost less until new workload comes along and then they'll go up. For each 10% CPC busy, the job CPU time will increase between 3% (Low-RNI) and 5% (High-RNI). ***Don't set normalization factor at this time!***
    - Latent demand could increase CPU use and you'll need more capacity even sooner.

# Hot Tips from Cheryl

- Money Savings Tips
  - Re-evaluate pricing yearly
    - More savings are found by re-evaluating software pricing than almost anything else.
    - Change LPAR configurations to minimize software costs. See Kathy Walsh's session 15703 this week.
    - See Al Sherkow's sessions (Boston SHARE) on pricing. Investigate his LCS software.
    - Investigate Capacity Provisioning Manager (part of z/OS).
    - Understand that chargeback to customer will never be 100% accurate or repeatable.
    - Learn early and well that jobs run during peak time will cost more than those during off-peak.
    - See our session 15836 on "Why Is the CPU Time so Variable?" and session 15845 "Software Asset Management" by **Scott Chapman** and **Nancy Pearlman.**

# Hot Tips from Cheryl

- Money Saving Tips
  - **Jeff Magdall** said early this week that z/OS 2.2 will probably not support z9s; and the next release or two will probably not support z196.
  - Don't stay on old machines – maintenance and new purchase/lease can often cost less than retaining non-current machines.
  - Consider adding zIIPs (not zAAPs) instead of adding more CPs. (More applications now move part of workload to specialty engines.) Try out ZAAPONZIIP.
  - Don't be afraid to negotiate with all vendors. The more expensive the product, the better chance you have of reducing the price.

# Optimize zIIPs to minimize HW and SW costs

The following are just *some* examples of products that use zIIP processors:
    RMF Monitor III (in z/OS 2.1)
    DB2, starting with DB2 V8 and more exploitation with each new version.
    Comms Server IPSec encryption and decryption
    z/OS XML Services
    z/OS XRC SDMs
    z/OS CIM Server
    OMEGAMON
    zAAP on zIIP
    Connect Direct V5.1.1
    IBM InfoSphere Guardium S-TAP for IMS on z/OS
    IBM DB2 Utilities Suite for z/OS
    eJES from Phoenix Software
    PKWARE PKZip
    CA Products (2008 paper listed 8 products sure there are many more now)
    BMC Products
    Many others  - check with your vendors and check RMF and SMF

# Optimize zIIPs to minimize HW and SW costs

Most products automatically use zIIP if available.

- Meaning that moving to a new release could (to your surprise) move work from CPs to zIIPs

Some products only use zIIP if you explicitly turn that on.

Some products only mark work as zIIP-eligible if a zIIP is available in that LPAR. This makes it difficult to predict how much of your work would move to a zIIP before it is actually installed.

- And don't forget to set PROJECTCPU to YES in IEAOPTxx to help you see how much CP time is spent on code that could have run on a zIIP.

Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval

# CF Response times

- The z/OS Heuristic algorithm aims to minimize z/OS CPU cost of each request. Generally speaking, asynchronous CF requests use more z/OS CPU time than synchronous requests, so aim for mostly synchronous requests.
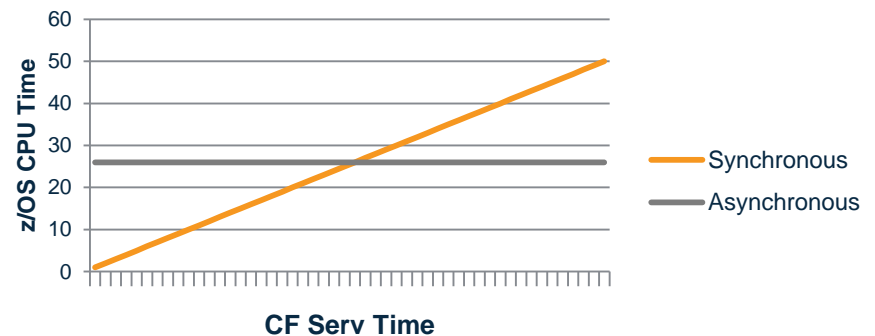


- Additionally, synchronous requests are less likely to result in L1/L2 cache getting overwritten.

# CF Response times

- How do you minimize CF service times:
  - Use fastest link technology available for your configuration
    - When using 12X Infiniband, try to use HCA3 adapters and run them in IFB3 mode.
  - Use dedicated CF engines.
  - If dedicated CF engines are not possible, enable Coupling Thin Interrupts on CF if CF is on zEC12/zBC12 with Driver 15 and ensure that production CF has an appropriate LPAR weight.

Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval

# CFs get no love

- How many people are in your storage management group?  And how many tools/products do they have to monitor and manage disk performance?

- How large is your CF performance management team?  And how many tools/products do they have?

Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval

# CFs get no love

- If you are doing data sharing, you are probably doing 2-3x as many CF requests as disk I/Os.

- If you are doing DB2 data sharing, most of the CF requests happen DURING the transaction – while most DB2 disk I/Os happen BEFORE or AFTER the transaction, so DB2 can be more sensitive to CF performance than to disk performance.

- A 10 mics response time increase might increase CF service times by 200%, but might only have a 1% impact on disk I/O times.

- Make sure that you are applying your scarce performance skills where they will deliver the most value.
  - Fantastic tools are available – but someone needs to set them up and monitor them.

Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval

# RMF PM

- RMF PM is a Java client that lets you design and run RMF reports in real time, taking data from multiplex sysplexes.
  - It is a very powerful capability, however it requires that you install code on your PC, which is not an option for some people.

- If you like RMF PM (or like the sound of it), you will LOVE the latest version of the z/OSMF Resource Monitoring function.
  - It has most of the capabilities of RMF PM
  - It only requires a browser on your PC
  - The graphics look sharper
  - It is easier/more intuitive to set up
  - You can use it to monitor other platforms (if you are into that type of thing…)
  - The latest service level gives you the ability to save the data into a CSV file for further processing
- Get handouts for **Peter Muench's** z/OSMF Resource Monitoring hands-on lab earlier this week (Session 15673)
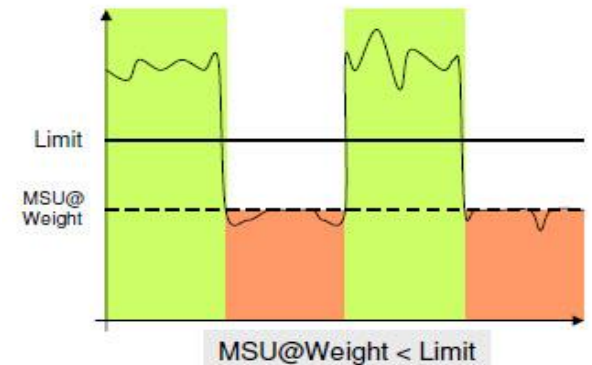
# What I can teach you about software pricing

- Do you LIKE paying more money than necessary to vendors for software? What about your management?

- If not, why on earth would you not use sub-capacity pricing?

- And why would you not use Defined Capacities and/or Capacity Groups (nearly 50% of respondents don't!)?
  - Nearly everyone has low priority work that can be rescheduled or can run for longer during prime hours.



Defined capacity limit enforcement
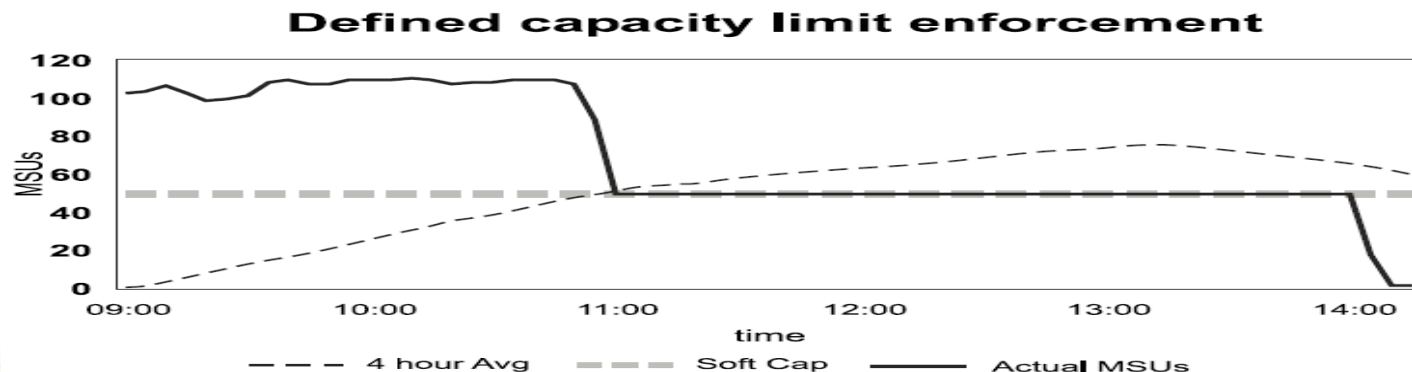
# What I can teach you about software pricing

- There are a multitude of LPAR capping options that can be used to limit the capacity of an LPAR – see **Horst Sinram's** excellent *Capping, Capping, Capping* Session 15719 and **Kathy Walsh's** *LPAR Performance* Session 15703 for more information.

- One scenario to be careful of is if the capacity guaranteed by the LPAR's weight is a lot less than the defined capacity.

- zEC12/zBC12, when used with z/OS 2.1, support negative phantom weights that result in LPAR being capped at its defined capacity rather than the lower LPAR weight.

- This month's Tuning Letter contains a 40+ page article on capping.

- **Robert Vaupel** and **Guenter Vater** have an excellent White Paper about the various System z capping capabilities -
*ftp://public.dhe.ibm.com/eserver/zseries/zos/wlm/**WLMCap**.pdf*

**What my daughter can teach you about software pricing**

- Every time my daughter comes home from the mall, laden down with shopping bags, she patiently explains to me (yet again!) how much she has saved me by emptying the shelves during a sale.

- What does this have to do with software pricing? Apparently (according to my daughter) more can be less!

- The ability to aggregate multiple CPCs and use PSLC pricing can result in significant savings. In order to be able to qualify for aggregation, the same sysplex must be using >50% of the used MVS MIPS (CP + *specialty engines*) on each CPC that you want to aggregate.

# What my daughter can teach you about software pricing

- What is your natural response to an edict to reduce SW costs? To REDUCE the amount of work being run.

- However, depending on the relative CPU consumption of your different sysplexes, INCREASING the amount of work running in the largest sysplex might SAVE you money if it makes the difference between being able to aggregate a CPC or not.

  – Move work from one sysplex to another

  – Move work between shifts

  – Move work from other platforms (web serving, for example)

  Test
  Prod

- If you do, or will, have mobile devices accessing your z/OS systems, start getting up to speed on Mobile Workload Pricing now – it is NOT simple.
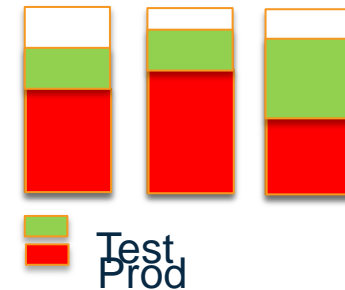
# zEC12/zBC12 and large pages

- On z10 and later, Fixed 1MB pages go in an OPTIONAL area of memory called LFAREA.

  - You only get this if you define it in IEASYSxx.

  - You control how much storage you want to dedicate to this.

  - Some products, like DB2, will use 1MB fixed pages by default IF you specify an LFAREA parameter.

  - You can get information about how LFAREA is being used with the D VS,LFAREA command and the RMF Mon III STORM report (although the corresponding RMF Data Portal report is much better)

Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval

# zEC12/zBC12 and large pages

- On zEC12 and later (any CPC that is SCM-capable) z/OS will *automatically* assign roughly 1/8 of total memory as the preferred location for Pageable 1MB pages.
  - You have nearly no control over this – you either get it (by default) or you can specify PAGESCM=NONE and you get no storage put aside for pageable large pages, but also no ability to use SCM for paging in that LPAR.
  - You get a PLAREA <u>EVEN IF THERE IS NO SCM INSTALLED ON THE CPC.</u>
  - There is no D VS,PLAREA command, and the support in RMF is limited, making it challenging to know how large your PLAREA is, and how it is being used.
    - RMF reports use of large pages, but doesn't report on utilization of PLAREA or LFAREA.

# zEC12/zBC12 and large pages

- Both PLAREA and LFAREA can contain 4KB pages if no 4KB frames are available.
- But the process of coalescing 4KB pages back into 1MB pages is not trivial (all frames must be contiguous and the 1MB frame must start on a 1MB boundary), so it is important to try to manage your storage so that 4KB pages do not need to be allocated in either of these areas.
- We have been contacted by a number of customers that have had issues with PLAREA, mainly related to the lack of metrics and lack of documentation.
  - When moving to a zEC12/zBC12, plan for LFAREA and PLAREA and ensure that the LPAR has sufficient memory.
  - There are also a number of APARs related to this:
    - OA41968, OA44207, OA44109, OA44183
  - Watch out for a possible DOC APAR from IBM (a follow-up to OA34024).
- There is also an article on this topic in the latest Tuning Letter.
- Elpida has a session on large pages on Friday morning (Session 16071).
- See White Paper 102464 for information about the benefit of large pages for DB2 buffer pools.

Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval

# More Coupling Facility Goodies

## Improved CF recovery and reduplexing times

z/OS 2.1 delivers new function to speed up recovery from failed CF, and reduced re-duplex time when CF comes back online – **CFLCRMGMT**.  Savings of >50% expected.

Also gives you control over the sequence in which structures are recovered, via new RECPRTY keyword in CFRM policy.

This is shipped DISABLED by IBM.  Your whole sysplex should be running z/OS 2.1 before you enable it (using the SETXCF FUNCTIONS command).

See Riaz's WSC Hot Topics slides for more info.

# More Coupling Facility Goodies

## Coupling Thin Interrupts.

New capability shipped with Driver 15.

Applies to either end of a CF link.

Each end can enable Coupling Thin Interrupts independently – you could have z10 CF and z/OS on zEC12 and z/OS could enable them and get the benefits from them.  Or CF could be on zEC12 and z/OS on z196 and CF could enable.

# More Coupling Facility Goodies

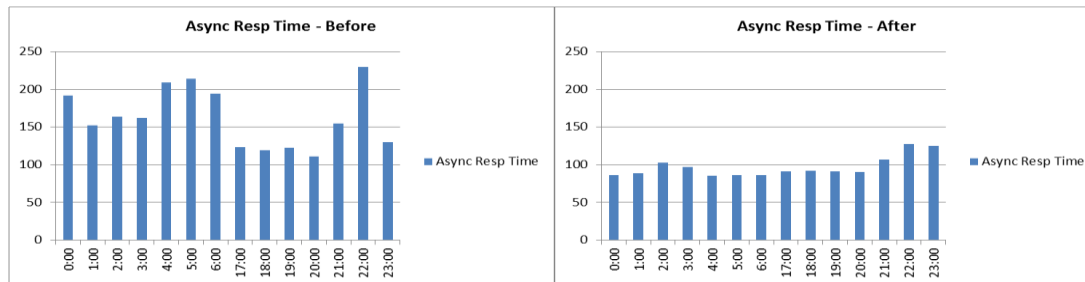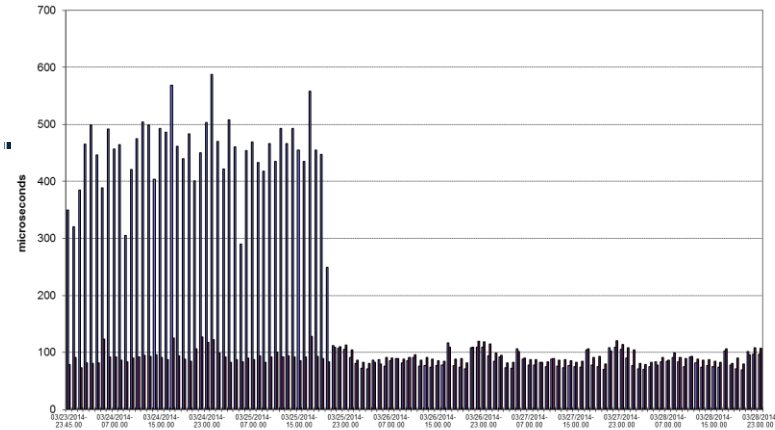Average Asynchronous Service Time by Coupling Facility

## Coupling Thin Interrupts.

Performance benefits are impressive.

Implementation is trivial – anyone with CFs with shared engines should be looking at this. And if your z/OS is already on EC12/BC12, you are probably already using this!

See WP102400 Coupling Thin Interrupts by **Barbara Weiler** and my session 15602 from this SHARE.

Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval

# More Coupling Facility Goodies

## MQ use of Storage Class Memory

CF Level 19 (with some MCLs), z/OS 1.13 and later with OA40747 support the use of Flash (SCM) as an overflow area for MQ structures.

Support delivered around the end of June.

Works with all supported levels of MQ.

Interesting thing is that this capability and the MQ Shared Message Data Set (SMDS) function appear to address the same issue – an MQ structure filling – but in different ways.

See Chapter 14 in draft Redbook SG24-8218 for a discussion of the relationship between these functions and how they can be used together.

Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval

# Catalog processing

- We heard a number of people complaining about catalog performance issues this week.  IBM's strategic answer is to use RLS-managed catalogs.  But there are some things you can do to help tide you over…..

- VLF provides the ability to cache catalog records more intelligently than In Storage Cache (ISC).

- However, the default when you define a new catalog is to use ISC.
  - If you want a catalog to benefit from VLF, you need to explicitly identify the catalog in the COFVLFxx member of Parmlib.

- You can EASILY get a summary of which catalogs are using ISC and which are using VLF by issuing the F CATALOG,ALLOCATED command.
  - Try to make sure that *at least* your busy catalogs are

# Catalog processing

- At the last SHARE, **Neal Bohling** (session 15089) *Using RLS with Your Catalogs – a How-To* provided some very interesting benchmark results comparing VLF-managed catalogs to RLS-managed catalogs:
  - RLS-managed catalogs showed significant elapsed time and CPU time savings for direct processing and decreased elapsed time but small increase in CPU time for sequential catalog processing
- However, the measurements did *not* include Enhanced Catalog Sharing.  ITSO ran similar experiment and found that <u>enabling ECS for a catalog delivered roughly half the elapsed time improvement that RLS provided</u> – but CPU savings were much smaller than RLS.
- You should consider enabling ECS for your catalogs that are not shared outside the sysplex until you are ready to move to RLS-managed catalogs.

# TECHDOCS status

- Some people have noticed that still-valuable material (FLASHes, White Papers, Presentations, etc) were being deleted from the TECHDOCS web site (http://www-03.ibm.com/support/techdocs/atsmastr.nsf/Web/TechDocs)

- TECHDOCS provide invaluable information to both IBMers and customers – if you have not gone there, do so NOW.

Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval

# TECHDOCS status

- Good news - Cheryl raised her concerns about the loss of this invaluable information.  IBM Management are aware of this issue and it is being escalated within IBM now and they hope to have a resolution soon.

- Ehhh news – if you select ANY value other than "Any time" in the "Include docs updated" search field in Techdocs, you will get 0 hits:

# TECHDOCS status

- With "Any time":

Techdocs Library >

## Searching All of the Techdocs Library

**Techdocs Library**
**Flashes**
**Presentations & tools**
**Technotes & tips**
**FAQs**
**White papers**
**Solution scenario profiles**
**Customer support plans**
**Sizings**
**Auxiliary Material**
**Search Techdocs**
- Most recent Techdocs updates
- Advanced search
**Techdocs feedback**

**Related links**
- Redbook publications
- IBM Software Support Handbook

The results of your search are below. You may also refine & resubmit the search.

**Search:** All of the Techdocs Library ▼    ☐ Allow word variants
**for:** z/os                                    ☐ "Fuzzy" search
**Hits:** 1000 ▼  **Order by:** most recent ▼   🔍 **Search**
Include docs updated: any time ▼    Help for Search

Documents 1 through 629 (of 629 found)

| Document ID | Document Title | Last Update | Doc Author or Owner |
|---|---|---|---|
| WP102439 | IBM WebSphere Liberty z/OS Connect | 08/06/2014 | Don Bagwell |
| PRS1778 | Mainframe System Education -- Wildfire Workshops | 08/05/2014 | Judith Vadnais |
| PRS4229 | Data Extraction Program (CP3KEXTR) for zPCR and zBNA | 08/05/2014 | Valerie Spencer |
| PRS1381 | Getting Started with zPCR (IBM's Processor Capacity Reference) | 08/04/2014 | John Campbell |
| TD103286 | z/OS Best Practices: Large Stand-Alone Dump Handling - Version 4 | 08/04/2014 | Riaz Ahmad |
| FQ127122 | Attached is the replacement to "High Performance FICON (zHPF) FAQs March 15, 2013 | 07/31/2014 | Jasmine N Burgess |
| WP101490 | WebSphere z/OS Optimized Local Adapters | 07/31/2014 | Don Bagwell |
| WP102468 | System zHPF: Commonly Asked Questions | 07/29/2014 | Jasmine N Burgess |
| WP102110 | WebSphere Liberty Profile for z/OS | 07/29/2014 | Don Bagwell |
| WP102467 | Performance Considerations on Capturing Workload with IBM InfoSphere Optim Workload Replay for DB2 for z/OS | 07/28/2014 | Akiko Hoshikawa |
| FQ131485 | IBM has announced SMC-R in z/OS v2.1 and RoCE feature support in zEnterprise EC12 and zBC12. SMC-R and RoCE provide high throughput with low overhead communications. This | 07/23/2014 | Linda Harrison |

# TECHDOCS status

- With "last 6 months":

# All Watson & Walker Sessions

- Thank you for coming!
- If you liked this session, please check out others from Cheryl and Frank:

    - 15836 – Tue, 3 pm – Cheryl & Frank – *Why Is the CPU Time for a Job so Variable?*
    - 15602 – Wed, 1:30 pm – Frank – *The Skinny on Coupling Facility Interrupts*
    - 15567 – Fri, 10 am – Cheryl & Frank – *Exploiting z/OS Tales From the MVS Survey*

- If you like SMF data, please see our new *SMF Reference Summary* at www.watsonwalker.com/references.html

# See you in Seattle!





Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval