

Cheryl and Frank's zRoadshow

Cheryl Watson (cheryl@watsonwalker.com)

Frank Kyne (frank@watsonwalker.com)

August 14, 2015

Session 17433



#SHAREorg



SHARE is an independent volunteer-run information technology association that provides **education, professional networking and industry influence.**

Copyright (c) 2015 by SHARE Inc.  Except where otherwise noted, this work is licensed under <http://creativecommons.org/licenses/by-nc-sa/3.0/>



Welcome



- Hi, thank you for coming.
- Who are we and what we do.
 - Three-person company; started in 1987
 - Frank Kyne joined us in 2014 from ITSO
 - Quarterly subscription-based newsletter –
Cheryl Watson's Tuning Letter
Cheryl Watson's System z CPU Chart
 - Classes on z/OS new features, WLM, performance, software pricing, chargeback, Parallel Sysplex, and high availability
 - Consulting on all the above topics
 - Software products – GoalTender and BoxScore

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

Welcome

- Education
 - z/OS Software Pricing Strategies
 - October 12-14, Orlando, FL
 - November 11-13, Hamburg, Germany
 - Exploiting New Features of z/OS To Minimize Costs
 - November 18-20, Hamburg, Germany
 - SYS1.PARMLIB: Enhancements and New Features in z/OS 1.13, 2.1, and 2.2
 - November 16, Hamburg, Germany

- All classes are available on-site

New Pricing Announcement

- IBM System z Software Pricing
 - <http://www-03.ibm.com/systems/z/resources/swprice/index.html>
 - Go to References, then Announcements to get Announcement Letters for CMP, zCAP, MWP
- Price Increases for 1/1/2016
 - Announcement Number [316-028](#) (US, July 28, 2015) or [516-001](#) (EMEA, August 11, 2015)
 - Average increase is about 4% in MLC software for z/OS, CICS, DB2, IMS, MQ, WAS, Tools, z/TPF, Compilers
 - SDSF V2 has 20% increase, even if version number has remained the same. See SHARE session 17432, SDSF Product Update for z/OS 2.2

SHARE Requirements



- **Industry Influence Through SHARE Requirements**
 - The SHARE Requirements system is an online advocacy program for SHARE members to influence IBM products and services. SHARE members can submit, discuss, and vote on requirements. All members are encouraged to participate - this is a great opportunity to make your voice heard!
- **What is a Requirement?**
 - A requirement is a statement suggested by SHARE members, describing a business need related to an IBM product, service, policy, or strategy of IBM, and must include a business justification.
 - Requirements are ranked and voted on by SHARE membership, and are submitted as an organization of many customers to IBM. IBM then responds to the Requirement back through SHARE.

SHARE Requirements

- From John Eells z/OS 2.2 Session 17438:

Availability

- **JES3 DSI Change**

- Not Dynamic System Interchange; that “other” DSI: Data Set Integrity
- In releases before z/OS V2.2, PPTNSDI must be set in IEFSDPPT (and not overridden by specifying DSI in SCHEDxx)
 - Default PPT entry for IATINTK remains:
 - C9C1 E3C9 D5E3 D240 ED10 (byte 8 bit 5 is PPTNSDI)
- This causes JES3 to use S99NORES (“don’t ENQ”) for its allocations
- z/OS V2.2 designed to support specifying DSI for JES3 in SCHEDxx
 - Default PPT still contains PPTNSDI for JES3 for now

- **Better Subsystem Interface (SSI) Initialization Processing:**

- SSCVT entry no longer intended to be built when initialization routines (INIRTNs) are not found
- Support for a new command to delete a subsystem:
 - SETSSI DELETE,SUBNAME=ssss,FORCE
 - (There are some restrictions!)

Top 50

Top 39er!

- **Dynamic Exit support for O/C/EOV**

- Support for the Tape Installation Exits: Volume Mount, File Start, File Validate, File End and Label Anomaly

IBM

© 2012, 2015 IBM Corporation 39

SHARE Requirements

- Status
 - IBM has been VERY supportive of our requirements and works to incorporate our top requirements into the development cycle (or even when code is opened for another reason).
 - As one example, the MVSE project created a Top 39 list in 2013, and z/OS 2.1 satisfied about 25% of the requirements
 - Then we created a Top 50 list in 2014; z/OS 2.2 satisfied about 30% of the requirements
 - The weight of multiple companies voting on a requirement and ranking them helps IBM prioritize these requests

SHARE Requirements



- Current Status
 - Starting in 2012, IBM began moving requirements from FITS to RFE (Request For Enhancement)
 - Since April 2014, no new requirements have been submitted through the SHARE system
 - That is changing!!!!
 - There is a team working with IBM and SHARE to create a new requirements process
 - If you're a project that wants to participate in (or learn about) this project, please contact me at Cheryl@watsonwalker.com
 - If you're a SHARE member who would like to participate, please go to current SHARE system and sign up for your favorite projects (Look for 'Getting Involved' or 'zNextGen' and click on Advocacy & Requirements)

SHARE Requirements

- Speaking of requirements
 - The MVSS project is looking on some feedback for prioritizing some of their requirements relating to 9 PDSE Small Enhancements
 - Please take this survey:
 - <https://www.surveymonkey.com/r/DVC7BFS>

CICS PBs Redux

- From our 1998, No. 1 Tuning Letter:
 - WLM Goal mode has very little overhead. The WLM policy adjustment runs once every 10 seconds, and WLM resource adjustment runs once every 2 seconds. These are very low frequencies and any overhead associated with them should not be noticeable. However, in goal mode, WLM state sampling runs once every quarter second. Still, this is not enough to cause very noticeable overhead in environments such as TSO and Batch. However, this could contribute to quite a bit of overhead in environments such as CICS.
 - In CICS environments, **if a response time goal is specified**, each transaction is represented by a 'performance block' or PB. Every time a new transaction is started it is associated with a PB, and every time there is a change in the state of a CICS transaction the new state is reflected in the PB for that transaction. Then every quarter second WLM samples the state of each transaction to gain insight as to where (i.e. in what regions) the transactions are spending their time processing. This state sampling is also used to provide to performance monitors information as to where (i.e. what state) transactions are spending their time processing.
 - In large CICS environments the number of transactions sampled during a sampling interval could be in the thousands. The number of PBs created and sampled is directly related to the CICS MAXTASKS parameter. It is for this reason that it is strongly recommended that MAXTASKS not be set higher than it need be. The larger the number, the higher the potential is for WLM sampling overhead. The lower the number, the higher the potential is for delaying your CICS transactions.

CICS PBs Redux

- From our 1998, No. 1 Tuning Letter:

High CPU When Sampling CICS

APAR **OW32140** addresses potential severe and dramatic increases in CPU usage during WLM sampling of CICS transactions. Some customers have experienced up to **40% to 50% of a single engine** being devoted to this type of WLM sampling. This overhead usually happens only in very large CICS environments.

- From our 2012, No. 3 Tuning Letter:

Chuck Hopf had a very large CICS configuration (350 test regions and 40 to 50 production regions) and had a total MXT of 30,000 (sum of MXT for all regions). Reducing the MXT to 8,000 reduced the WLM CPU usage from 40-50% of a single CPU to 16% of a single CPU.

CICS PBs Redux

- From our 2015, No. 2 Tuning Letter:

WLM APAR and CPU Time

One of our regular readers and contributors, **Todd Havekost** from USAA, kindly let us know about a recent WLM APAR. Todd noticed that their WLM CPU consumption had increased significantly (in the region of 50 to 85%) since migrating to z/OS 2.1. It turned out that a particularly expensive hardware instruction had been added to a routine in WLM that processes performance blocks (PBs) 4 times a second. You might wonder how a single instruction could make such a difference – the answer is that it depends on how many times it is executed.

In Todd's case, the system had nearly 40,000 PBs, so that instruction was being executed at least 160,000 times every second. How did he have so many PBs? This particular system was a test system with 9 DB2 subsystems, 80 CICS regions, and VSAM RLS.

CICS PBs Redux

- From our 2015, No. 2 Tuning Letter:
 - IBM opened APAR [OA48161](#) (Increase in WLM CPU Consumption in z/OS 2.1) to address this situation. At the time of writing, it is still open.
 - **RECOMMENDATION.** We recommend that you add yourself to the interested parties list for the APAR so that you will be informed when the PTF is available.
 - From the APAR:

PROBLEM DESCRIPTION: After migrating to z/OS 2.1, the CPU utilization of the WLM address space may increase on a system with tens of thousands of PB control blocks. In z/OS 2.1 the size of the Performance block (PB) control block was increased by 256 bytes. During WLM PB sampling in module IRASASRV, the PB control block will be copied to private storage of the sampling module using a MVCSK instruction(s). With the increased size of the PB it was necessary to use two MVCSK instructions instead of one because MVCSK is only able to copy up to 256 Bytes. The additional MVCSK can lead to a higher CPU utilization by the WLM address space in z/OS 2.1.

CICS PBs Redux

- From our 2015, No. 2 Tuning Letter:

CICS TS V5 and Changed Default

Regarding PBs, you need to watch for CICS TS V5. Prior to CICS TS V5.1, the minimum, default, and maximum values of MXT were 1, **5**, and 999. Effective in V5.1, the minimum is now 10, the default is now **500 (!!!)**, and the maximum is 2000. If you've been using the default MXT for CICS V4 and move to CICS V5, you might see a similar increase in the WLM CPU time. This is especially true if you have gobs of CICS regions in your test and development systems

CICS PBs Redux

- Lessons Learned:
 - Specify CICS MAXTASK to correspond to the average (or even maximum) number of concurrent transactions, and don't use the default
 - Monitor WLM usage
 - Reduce WLM CPU time on test and development systems by using velocity goals instead of response time goals, unless you need to meet your service level objectives (SLOs)
 - Have a weekly tea time (or meet up at the bar) so that CICS and z/OS sysprogs/performance analysts can learn to communicate!

Preparing for Audits by ISVs

- Lately, several ISVs have been increasing their number and frequency of audits
- Because many ISV contracts are based on decades-old license agreements, the terminology does not reflect the current configuration
- Most problems occur due to disagreement in the wording of the contract as it relates to sub-capacity pricing
- Most of the items in the contract checklist on the next page are NOT currently specified in contracts, but instead left to the interpretation of the reader


Preparing for Audits by ISVs

- Contract items:
 - If license is in MIPS instead of MSUs, how are the MIPS calculated?
 - How is sub-capacity pricing handled (i.e. LPAR-based or CPC-based)?
 - How is the usage of the product determined?
 - What happens if maintenance is dropped?
 - Is maximum usage based on a yearly average or single month?
 - Is usage based on total consumption or peak usage or average usage?
 - How are peak seasons handled?
 - How is OOCoD handled?
 - How are processor upgrades handled?

Preparing for Audits by ISVs

- Contract items (cont.):
 - Are penalties included on either the part of the vendor or the customer?
 - How far back can audits go?
 - What happens if the product is taken over by a different vendor?
 - If the product produces a file or database, it is defined or proprietary (that is, can the data be exported to a different format or product)?
 - How quickly will the product be updated after a new processor, operating system, or subsystem (e.g. CICS) is released?
 - What is the length of the initial contract (our advice – never agree to more than two years)?
 - How much can the price or maintenance increase each year?
 - What is the cancellation policy
 - What is the recourse in case of a disagreement (arbitration, etc.)?

Preparing for Audits by ISVs



- Contract items (cont.):
 - If SCRT is used, HOW is it used?
 - Simple example: CPC is 188 MSUs; max of z/OS is 181 (2 LPARs); COBOL is only 57; sum of N5 is 220 (more than max of machine)
 - Don't let ISVs use N5 for multiple LPAR usage!

```

===== B5===== SCRT SUB-CAPACITY REPORT - IBM Corp =====
Run Date/Time          11 Aug 2015 - 11:28
Machine Type and Model 2827-701
Machine Rated Capacity (MSUs) 188
===== E5=====
PRODUCT SUMMARY INFORMATION
MLC Product Name      MLC Product ID  Tool MSUs  Customer  Customer Comments (255 chars ma
z/OS V2                5650-ZOS        181
DB2 10 for z/OS       5605-DB2        181
Enterprise COBOL for z/OS V5  5655-W32        57
IBM Enterprise Cobol for z/OS V4  5655-S71        181
===== N5=====
DETAIL LPAR DATA SECTION
                Highest      Hour Count  Date/Time      2nd High Hour Co  Date/T
PROD01          162          1 28 Jul 2015 - 11:00      140      1 28 Jul
PROD02           58          1 25 Jul 2015 - 08:00       57      2 24 Jul
Total of N5      220
CPC              181          1 28 Jul 2015 - 11:00      161      1 28 Jul
===== P5=====
PRODUCT MAX CONTRIBUTORS
                LPAR      LPAR
Product Name      Product ID      Highest      Date/Time      PROD01      PROD02
z/OS V2           5650-ZOS        181 28 Jul 2015 - 11:00      162      19
DB2 10 for z/OS  5605-DB2        181 28 Jul 2015 - 11:00      162      19
Enterprise COBOL for z/OS V5  5655-W32        57 24 Jul 2015 - 13:00       0       57
IBM Enterprise Cobol for z/OS V4  5655-S71        181 28 Jul 2015 - 11:00      162      19
    
```

Neat Ways to Learn

- New Era Software
 - <http://www.newera-info.com/>
 - New Era produces some really neat software tools: SAE, ICE, IPLCheck, TCE, and many more
 - Also provides huge contributions to the z/OS community
 - Webcasts
 - Series of Mainframe Security webcasts from Mark Wilson
 - Series of z/OS Crypto webcasts from Greg Boyd
 - RACF Updates from Mark Nelson
 - ISPF from Tom Conley
 - CA ACF2 & TSS from CA
 - Slide decks from: Marna Walle, John Eells, Steve Warren, Angelo Corridori, Greg Daynes, Thomas Cosenza, Glenn Anderson

Neat Ways To Learn

- The treasury of SHARE.org
 - Webcasts
 - Over 50 webcasts from SHARE presentations
 - SHARE Videos
 - Look under SHARE Content
 - Lots of regular blogs:
 - MindSHARE, Marna's Musings, MVS Program Blog

Neat Ways To Learn

- The treasury of SHARE
 - Blogs
 - E.g. LISTSTAT command - <http://www.share.org/p/bl/ar/blogaid=380>
 - White Pieces
 - Several short papers on SHARE research initiatives
 - E.g. Mobile App Development

SHARE in Orlando | August 9-14, 2015 | Register Today

Home About Members Only Events Content Center Blogs Our Partners zNextGen Volunteers Store

Blogs Home Search Favorite Blogs Articles in Favorites All Blogs Options More Options

mindSHARE

IDCAMS SHCDS LISTSTAT

By: Terri Menendez and Pooja Bhargava, Software Development, IBM

Frustrated that you have to close your files just to see the latest catalog stats? Wondering when that next opportunity will be? Well, stop wondering, because the time is NOW! Introducing the new IDCAMS SHCDS LISTSTAT command, where file statistics you could only dream about are now at your fingertips!

Whether your files are open or not, the new LISTSTAT command displays sysplex-wide catalog statistics for your VSAM record-level sharing (RLS) data sets. Interested? Here's how:

In z/OS V2R2 (or z/OS V2R13 and later with the PTF for APAR OA42435), issue the following AMS command:

```
SHCDS LISTSTAT ('<cluster name>')
```

The essential file statistics such as totals for records, inserts, deletes, updates, retrieves, CA/CA splits, extents information, free space, index levels, and high-used/allocated RBA are immediately returned to you.

-----Sample Output-----

```
TSO SHCDS LISTSTAT ('SYSPLEX.VSAM.KSD2')
```

```
LIST STATISTICS (LISTSTAT):
CLUSTER-----SYSPLEX.VSAM.KSD2
DATA-----SYSPLEX.VSAM.KSD2.DATA
TOTAL RECORDS----- 3000    CI SPLITS----- 86
RECORDS DELETED----- 0      CA SPLITS----- 2
RECORDS INSERTED----- 1498  EXCPS----- 6613
RECORDS UPDATED----- 0      EXTENTS----- 5
RECORDS RETRIEVED----- 0     FREE SPACE----- 1769472
PT 2  PPA  1147200  PT 2  PPA  1147200
```

Recent Stories

SHARE in Orlando: Opening Keynote Celebrates Mainframe Strength in Age of Disruption

Updated 11:50AM EDT, Wed Aug 12th, 2015

Tom Rosamilia, Senior Vice President of IBM Systems at IBM Corporation, kicked off the opening general session of SHARE in Orlando on a note of celebration — both for SHARE's 60th annivers...


Introducing SHARE White Pieces

Updated 4:40PM EDT, Sat Aug 8th, 2015

SHARE is excited to announce its new research project, Whitepieces. The objective of a Whitepiece is to share information about technology or a process to implement technology. This can be new te...

Want to REST or Get WET in

Is there a Sub(Capacity) in your future?



CPU MF

- We have been seeing an increasing drive from IBM for customers to consider sub-capacity models when looking at upgrades.
 - Sub-Capacity models provide more points of dispatch, more L1, L2, and L3 cache (than a 7xx with similar MIPS), ability to have more Vertical High CPs and zIIPs.
 - Many customers do not really need 1700 MIPS engines.
 - Use IBM's free zBNA tool to help you identify CPU-intensive batch programs and model impact of CPU change.
 - Customer experience (based on zEC12) has been very positive.
 - Use [zPCR](#) to see if a sub-capacity model would be a good fit for *your* configuration before making a final decision.

Large Memory

- One of the hot (and recurring) topics of the week was the value of large memory on z13.
- One of the changes on z13 compared to previous generations is that rather than striping an LPAR's memory across all books, PR/SM now attempts to place all an LPAR's memory in the same drawer as its CPs and zIIPs. If successful, this should result in speedier access to memory than was the case on previous generations. And one of the best ways of ensuring its success is to provide as much memory as you can afford AND to order memory 'shrewdly'.
- LUCKILY, IBM have special offers to encourage you to do just that. If you install 3x as much memory as the box that is being replaced, you will only pay 1/3 as much per GB as you paid for zEC12 memory.
 - IBM also have a special offer if you install 5x as much memory, presumably at an even bigger discount than if you go for the 3x option.
 - According to IBM, 75% of z13s shipped in 1Q availed of 3x memory offer.

Large Memory Sessions

- The following are some of the sessions wholly or partly dedicated to large memory this week:
 - [17878](#): *How to Leverage Large Memory on z* by **Elpida Tzortzatos**
 - [17634](#): *z/OS Processor Storage Management* by **Kathy Walsh**
 - [17321](#): *Memory Management in the TB Age* by **Scott Chapman**
 - [17434](#): *The New IBM z13 Part 1: Processor Design, Server Structure, z/Architecture Enhancements, and Operating System Support* by **Harv Emery**
 - [17748](#): *z13 Performance* by **Tarun Chopra**
- Other great reference:
 - IBM Redbook [SG24-8251](#), *z13 Technical Guide*, Section 2.4

Large Memory

- If you have DB2 V11, use the Buffer Pool Simulation function (*not* rolled back to DB2 V10) to help estimate the benefit that larger buffer pools would deliver in terms of the number of DB2 I/Os that would be eliminated.
 - It is estimated that each I/O costs between 20 and 70 mics of CPU time.
 - Additionally, the I/O intensiveness of your workload is one of the contributors to your Relative Nest Intensity.
 - I/O intensive workload contributes to High RNI which means less MIPS.
 - Depending on the model, a low RNI workload can get up to 34% more MIPS than a high RNI workload on the same CPC.
- For more info about using this capability in DB2 V11, see [DB2 V11 Managing Performance](#)

Large Memory

- Few more sources of information about exploiting large memory for DB2:
 - **Rob Catterall** blog [entry](#) about DB2 and really large buffer pools.
 - **Adrian Burke** [Session 17290](#): *From VIRTUALLY Constrained to REALly Overcommitted: a Study of DB2 Storage*
 - IBM DB2 SAP White Paper [IBM z Systems: Performance Report on Exploiting Large Memory for DB2 Buffer Pools with SAP](#)
 - For everything you wanted to know about RNIs, **John Burg's** [Session 17556](#): *2015 CPU MF Update*

Counters

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

Little Known DB2 Factoid

- Who would like to reduce general purpose CP utilization, possibly reduce your software bill, AND get more value from your zIIPs?
- Who has DB2 batch jobs that run queries?
- IF you enable DB2 query parallelism, when DB2 splits a query, the different parts of the query execute under enclave SRBs, *and that means that they can use zIIP MIPS.*
 - Plans must be bound with DEGREE(ANY) to be eligible for query parallelism.
 - Maximum degree of parallelism is controlled via the PARAMDEG parameter in DB2 ZPARM. Note that the default value changed in DB2 V10. Also, don't specify a higher degree of parallelism than you have zIIP engines.
- For more information, refer to this [blog entry](#) by **Rob Catterall**.
 - Also, download the NO CHARGE DB2 Data Studio which will help you identify good candidates for query parallelism.
- And see this [blog entry](#) from **Willy Favero** regarding the PARAMDEG change in DB2 V10.

Country Multiplex Pricing

Country Multiplex Pricing

- IBM previewed a new software pricing option called Country Multiplex Pricing (CMP) as part of the z13 announcement in January 2015, and published the [formal announcement](#) on July 28, 2015.
- CMP is intended to provide significantly more flexibility in terms of how you configure your systems, to remove the financial incentive to create nonsensical sysplexes, and to reduce the cost of growing existing z/OS systems.

Country Multiplex Pricing

- So what does CMP do for you?
- Do you have data centers in two distant, exotic, locations (Hawaii and Poughkeepsie!) meaning that you can't aggregate them?
- No problem – CMP removes the requirement for the CPCs in the pricing plex to be in the same sysplex. The only requirement is that they are located in the same country.

Country Multiplex Pricing

- What else does CMP do for you?
- Do you find that you are constantly struggling to meet the dreaded 50% threshold??
- No problem. CMP eliminates the 50% threshold!!*
- * However, see comment later about sysplex aggregation.

Country Multiplex Pricing

- What *else* does CMP do for you?
- Driven to distraction (and unplanned outages) because you have unrelated systems in the same sysplex, PURELY to meet the 50% threshold?
- Good news. Because CMP eliminates the 50% threshold, there is no financial incentive to place all those disparate systems in the same sysplex any more.



Country Multiplex Pricing

- What **else** does CMP do for you?
- Agonizing over how to provide the failure isolation that a production CF requires (do you go for a standalone CF, or an internal CF that uses system-managed duplexing)?
- CMP has you covered. There is no driver to have all your CPCs in the same sysplex any more, so you can have production-only CPCs and test-only CPCs (if that makes sense for your organization). You can then place your production CFs in CPCs containing only test z/OS systems, and test CFs in CPCs containing only production systems. This gives you the required failure isolation without the cost of duplexing or a standalone CF.

Country Multiplex Pricing

- Anything else?
- Afraid to implement shared queues or dynamic workload routing because that makes your work more ‘mobile’, with the potential to adversely affect your software bill?
- No more. When using CMP, your monthly software bill is created by adding up the R4HA for each LPAR for each interval *across ALL CPCs in the country*. So moving xx MSUs of work from one CPC to another CPC has zero impact on your software bill because the sum of the MSUs for that interval has not changed.

Country Multiplex Pricing

- (This is getting boring) what else does CMP do for me? I suppose it lets me move MIPS from one CPC to another (ha ha)?
- Do you have scenarios where one CPC has unused Defined Capacity, but LPARs on another CPC are running at their cap?
- Not any more. Because R4HAs (or softcaps) are summed across ALL LPARs, you can now decrease a Defined Capacity on one CPC, increase a Defined Capacity on another CPC by the same amount, with no impact on your bill.

Country Multiplex Pricing

Is this cool or what??!!



WOOHOO!!!

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

© Watson & Walker 2015

8/16/2015

Country Multiplex Pricing

- Still not happy?
- Single Version Charging normally gives you one year to migrate from one version of a product to another. If you do not complete the migration in one year, you revert to paying for two separate products.
- If you are using CMP, you get a new thing called Multiple Version Measurement (MVM). This not only removes the one year window to complete the migration, it doesn't have ANY migration window! So you can keep using both versions for as long as you want, and you only pay for the combined MSUs under one product (the newer version).

Country Multiplex Pricing

- You STILL want more?
- For large customers, the AWLC pricing curve stops decreasing at about 16000 MIPS, meaning that the bulk of your capacity is in the top tier (tier 9).
- If you are using CMP, there are 15 tiers. The CMLC pricing curve doesn't stop decreasing until about 115,000 MIPS, and the z/OS \$/MSU for the top tier is about 27% cheaper than the top AWLC tier.
- See an example of savings in our SHARE session 17653 – *A Consultant's View of New z/OS Pricing Options*

Country Multiplex Pricing

- What's the catch?? (Yes, there is always a catch)
- CMP is designed to reduce the cost of *growth*.
 - If you sign up for CMP and do not change anything else, your bill will NOT decrease.
 - When you sign up for CMP, your last 3 months MSUs and software bills are averaged and those values form your new base. Any changes (more MSUs or fewer MSUs) are applied as deltas to that base.
 - **Do EVERYTHING in your power to minimize peak R4HA in the 3 months before you switch to CMP.**
- While CMP does not require a primary sysplex, you DO need to have a qualified sysplex *before* you go into CMP in order to have the lowest base. Once you switch to CMP, you can dismantle that sysplex if that is the right thing to do from a technical perspective.
- If you are planning on a CPC upgrade, consider whether you should move to CMP BEFORE the upgrade or after.
- There are special restrictions for outsourcers (and, by extension, outsourced customers) – for more info, check the section titled “Service providers” in the announcement letter.

Mobile Workload Pricing

Mobile Workload Pricing

- What is Mobile Workload Pricing?
- New IBM Software Pricing Option, announced in May 2014.
- Headline is that it offers a 60% discount on MSUs consumed by transactions that originated on a mobile device.
- 60 percent Off! What else is there to say??



- Quite a bit....

Mobile Workload Pricing

- The IBM announcement stated that YOU are responsible for determining (from SMF records), how much CPU is consumed by MWP-eligible work.
- Many customers start delving into their SMF Type 110 records (CICS) to find a field that identifies txns that originated on a mobile device:
 - This can be challenging, if not impossible.
 - The required information may not exist in the record.
 - The program doing this processing must understand the identifying data for every transaction that could potentially originate on a mobile device.
 - It involves processing potentially huge amounts of SMF data.
 - In the best case, it might identify 80% of the CPU consumed by or as a result of those transactions.

Mobile Workload Pricing

- While the announcement letter did not state it, IBM have agreed to accept address space-level data for regions and subsystems that are dedicated to MWP-eligible work.
 - It might be easier to identify the source of the transaction prior to region selection (in the network, for example).
 - No need to maintain the program every time a transaction changes or there is a change to how MWP-eligibility is determined.
 - ALL the CPU used by the region/subsystem is captured.
 - The volume of data to be processed is trivial compared to processing transaction-level records.
- Similarly, IBM will accept Type 70 and 89 records from LPARs that are dedicated to MWP-eligible work.

Mobile Workload Pricing

- All the other prerequisites and rules remain the same.
- Still must sign agreement and associated addenda
- Have to document and prove to IBM that only MWP-eligible transactions can run in these dedicated regions.
- Initial setup effort may be a little higher, but the ongoing maintenance costs are negligible, and the difference in captured CPU time might be significant.

Mobile Workload Pricing

- Secret sauce to success – promise little, and deliver a little more than that.
- MWP reduces your *R4HA* by 60% of the CPU used by MWP-eligible work. It does NOT reduce your bill by 60%!
- If your peak R4HA happens at a time when there is zero MWP-eligible work, MWP will not reduce your bill by 1 cent.
- BUT, mobile workloads are growing at a huge rate. Not only are NEW transactions coming from mobile, users of traditional sources are switching to use mobile devices instead.
- You cannot afford to ignore this trend.
- If you do not have much mobile work today, take this opportunity to plan for how you will handle it when it does arrive.

Sysplex and High Availability

Sysplex and High Availability

- Mark Brooks from Sysplex Development in Poughkeepsie had an excellent session about using Parallel Sysplex to achieve high availability – session 17431.
 - This was actually part 1 of a 2-part session.
- Mark's presentation is really chock full of interesting facts and some painful insights into the promise of High Availability and the reality....
 - For example, 75-80% of unplanned outages are caused by process failures, rather than technology failures.
 - z/OS is estimated to have 2 ^{1.2 million} unique code paths. There have been 2 ⁸⁹ nanoseconds since the big bang. So what are the chances that IBM can test every possible code path in z/OS after every change they make?
 - Identify the largest source of application down time and address that. Then identify the next largest, and address that. And so on.
 - 'Planned' downtime is typically far larger than unplanned, so don't ignore those outages in your investigations.
- If you are responsible for availability in your site, you **HAVE** to pull down Mark's slides. And attend part 2 of this session at the next SHARE.

Powerpoint Karaoke meets High Availability

- Mark's description of the process weaknesses that he has encountered strangely coincided with one of the slides from the low-moment of my week, the Powerpoint Karaoke.

Transaction
abend

~~Moderation~~



DB2
abend

~~Precision~~



System
Down

~~Creativity~~



Sysplex
Down

~~Panic~~



Miscellaneous

Miscellaneous

- One of the new features on z13 is 16Gb FICON links.
 - 16Gb is new on z13, but has been available on switches for a number of years.
- When customers started rolling out 16Gb links on switches, they encountered a huge increase in error rates.
 - It transpires that 16Gb links are FAR less tolerant of less than optimal cabling than 8Gb or slower links.
- If you have plans to run the 16Gb links at 16Gb, include a comprehensive review, cleaning, and potentially replacement of cables as part of your implementation plan.

Miscellaneous

- How many times did you have to decide which of two SHARE sessions you wanted to see? Or three sessions? Or five?
- It is great that SHARE has such a wealth of fantastic sessions, but it is very frustrating to have to miss a session that is really important to you.
- The answer to at least part of your problem might be SHARE Live.
 - This gives you access to recordings of a subset of sessions, for up to 6 months after the conference.
 - So your colleagues that couldn't travel can enjoy some of the sessions.
 - And you can sit in on your 2nd choice session, secure in the knowledge that you can watch your 1st choice after you return home.



Miscellaneous




- SMT2 support is new on z13, and is currently only available for zIIPs and IFLs.
- A number of sessions at this SHARE touched on various aspects of the theory, but I didn't see any based on real user experiences.
- If you are using SMT2 with zIIPs, or plan to do so in the future, and would be willing to share your experiences with Cheryl and me, please send us an [email](#).

Astounding Technology

Technology and mankind...

- I attended some fantastic sessions this week, met many kind and friendly people, enjoyed the weather (when I was indoors), and learned a TON.....
- But the highlight of the week was this insight into how far we have come with technology and the promise that the future holds for us by one of IBM's most respected mainframe gurus...

Amazing technology

- Today we have phones that have more memory than 80 STRINGS of 3330s. You can watch TV, read books, have free video calls, book flights, control your home heating, and even speak to other humans, all on your  Phone.



- One z13 CPU processes as much work in 1 hour as my old 370/138 did in 294 DAYS!



- The wingspan of an A380 is twice the length of the Wright brother's first FLIGHT.



Amazing technology

- Husbands and fathers no longer have a monopoly on knowing everything



- You can use your house to power your car, and your car to run your house.



Google Translate

Google Inc. Tools

Everyone

This app is compatible with all of your devices.

Star Trek's Universal Translator is here. And it can read, listen, and speak

Amazing technology

- And what do people do with this fantastic technology?



- Watch cats

- And fight with people they don't know....



Thank you!

- Thank you for coming along and making this week so enjoyable. And we hope we'll see you again in San Antonio. Safe trip home...
- Please remember to complete an evaluation



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

© Watson & Walker 2015

