The Watson and Walker zRoadshow

Session 23894, March 15, 2019

Mario Bezzi, Frank Kyne

www.watsonwalker.com

technical@watsonwalker.com

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

Welcome!

- Who are we?
  - Cheryl Watson, CEO of Watson & Walker Inc. since 1986, working on IBM mainframes since 1965.
  - Brenda White, Mario Bezzi, Scott Barry, and Frank Kyne, plus Tom Walker (Cheryl's husband), Alan Murphy, and Graham Horne.
- We publish Cheryl Watson's Tuning Letter (since 1991), teach classes, perform performance and sysplex reviews, provide software pricing consulting, and have two software products: BoxScore and GoalTender.
- Our latest passion is our SCRTPro Service Offering, where we help clients get the maximum value from their z/OS systems.
- We have two websites:
  - Watsonwalker.com General web site, free access for everyone.
  - Watsonwalkerpublications.com Subscriber-only – all Tuning Letters for last 28 years.
- This session provides some tidbits of information that we hope you will find valuable.

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

Cheryl and Tom's Big Trip

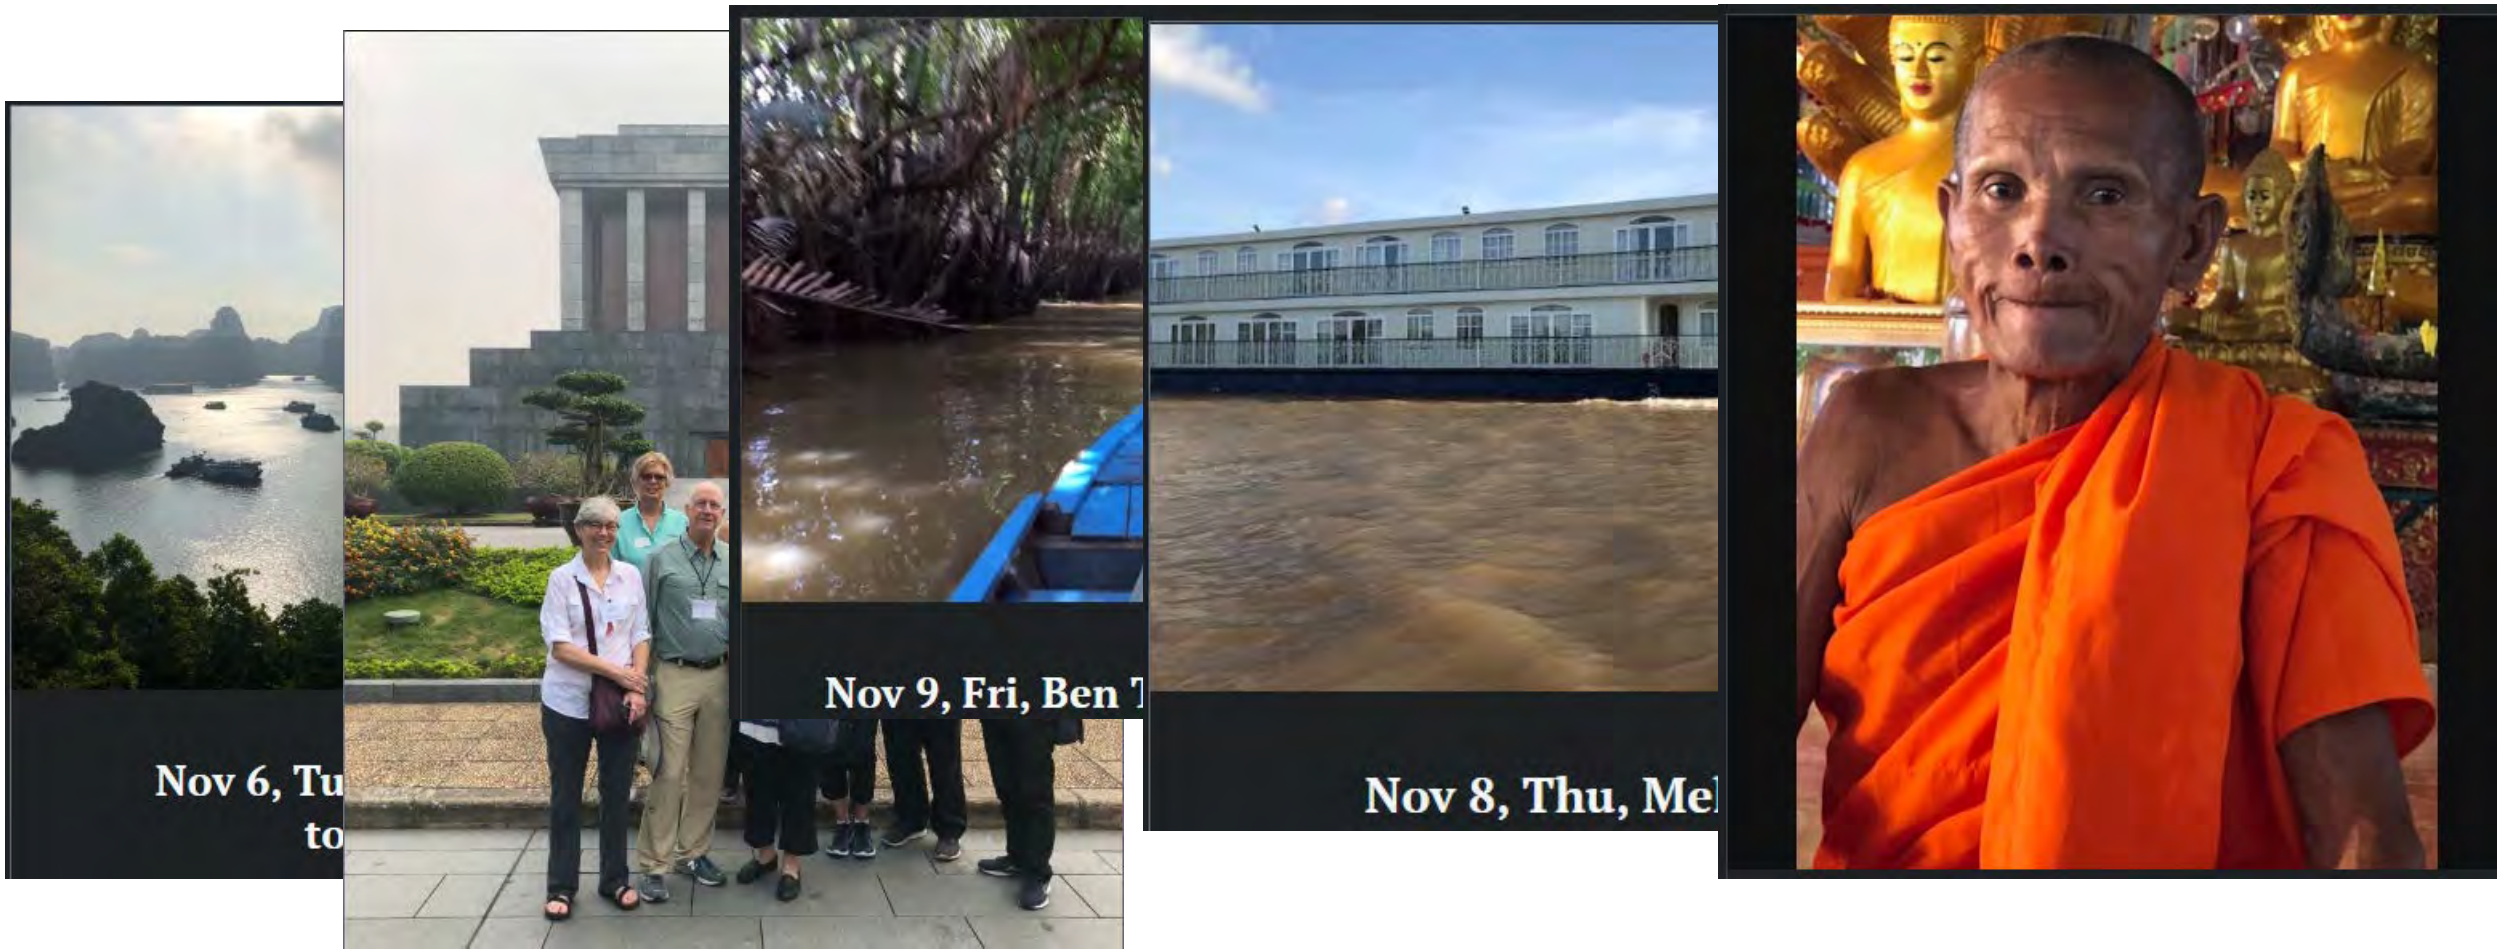

For more about life out in the real world, see http://tomandcherylstravels.blog

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

DFSORT and SMF Data

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

DFSORT and SMF Tips

- We spend a LOT of time working with SMF data.
- In some cases, we only want a subset of a particular set of SMF records – for example, just the type 30 records for job FATFRANK.
- However, the standard SMF utility, IFASMFDP, only supports selection at the type and subtype level.
- So I had a stroke of GENIUS and coded up a DFSORT job to extract out just the type 30 records that I needed. How difficult could that be? Even *I* could do that!
- Sigh, well, maybe not – "ICE218A 18 BYTE VARIABLE RECORD IS SHORTER THAN 37 BYTE MINIMUM FOR SORTIN FIELDS".
  - SMF Type 2 and 3 records are placed in every file created by IFASMFDP, and they are only 18 bytes long. This means that the fields that I want to filter on will always be beyond the length of the shortest records in the file.

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

DFSORT and SMF Tips

- Luckily for me (but unluckily for him!) I happened to be talking to **Sri Hari Kolusu** from DFSORT development about something else at the time, and he gave me a tip to add an INCLUDE control statement like:

```
INCLUDE COND=(01,02,BI,GE,330,AND,
  06,01,BI,EQ,X'1E',AND,
  23,02,BI,EQ,X'0005',AND,
  231,8,CH,EQ,C‘FATFRANK')
```

- This effectively tells Sort to ignore any record shorter than 330 bytes, thereby avoiding the ICE218A message.
- So now I can extract a few KB worth of type 30 records, rather than GBs of them. Yipee!

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

DFSORT Tips

- Given his unofficial new role as my DFSORT mentor, Sri Hari also kindly threw in this tip (no extra charge!):

```
//S1 EXEC PGM=SORT
//SYSOUT DD SYSOUT=*
//SYMNOUT DD SYSOUT=*
//SYMNAMES DD *
SMF30LEN,001,02,BI $ RDW WITH LENGTH
SMF30RTY,006,01,BI $ RECORD TYPE
SMF30STP,023,02,BI $ RECORD SUBTYPE
SMF30JBN,231,08,CH $ JOB OR SESSION NAME
SMF30PGM,239,08,CH $ PROGRAM NAME
S30MLEN,330 $ MINIMUM LRECL = 330
SMFRTYP,30 $ RECORD TYPE = 30
SMFSTYP,2 $ SUBTYPE = 2
JOBNAME,C'FRANK2' $ JOBNAME = FRANK2
/*
//SORTIN DD DISP=SHR,DSN=KYNEF.S0W1.RAWSMF
//SORTOUT DD DISP=(,CATLG),DSN=KYNEF.S0W1.SMF.EXTRACT.FRANK2,
// UNIT=SYSDA,SPACE=(CYL,(25,10),RLSE),
//SYSIN DD *
 OPTION COPY,VLSCMP,SPANINC=RC4
 INCLUDE COND=(SMF30LEN,GE,S30MLEN,AND,
 SMF30RTY,EQ,SMFRTYP,AND,
 SMF30STP,EQ,SMFSTYP,AND,
 SMF30JBN,EQ,JOBNAME)
/*
```

Rather than having to remember the offsets of my favorite fields in the various SMF records, we can use symbols.

Define them here (or in a shared PDS member)

And use them here…

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

DFSORT Tips

- Sri Hari also provided a sample ICETOOL job to read a file of SMF records and provide a count of how many of each SMF record type and subtype are in the file:

```
TYPE  SUBTYPE   COUNT
---   -------   -----------
30    1         17720
30    2         3984
30    3         59610
30    4         59838
30    5         18893
30    6         300
41    3         61
42    1         30
42    2         992
42    4         1
42    5         138
42    6         950961
42    9         3
42    21         195
42    24        2587269
42    25        11
42    27        92352
```

We hope to provide a sample program in the Free Tools part of our website soon that will list the start and end date and time, and source system, for each set of record type and subtype found in a data set.

- The sample is too long to repeat here, but it is in the latest Tuning Letter (2018 No. 4), in the User Experiences section

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

‘Optimizing’ Sorting of SMF Data

- If you read multiple generations of a GDG, and reference them using the GDG base name, the default sequence is that you get the most recent data set first, then the one before that, then the one before that, and so on.
- If the generations contain data that you want to use in oldest-to-newest order (SMF data, for example), you would traditionally have to sort them into the ‘right’ order before you could use the data.
- As an alternative, consider defining the GDG with the FIFO attribute.
  - Allocating the data set with DSN=MY.GDG now results in the data sets being allocated as 1, 2, 3, rather than 3, 2, 1 (the default).
  - You can also use the IDCAMS ALTER command to change this attribute for an existing GDG.
- See slide 11 in **Stephen Branch** and **Frank McCune’s** Session 23953 for more information. This idea was also described in a User Experience in *Tuning Letter 2015 No. 2*.

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

CPU Measurement Facility and SMF 113 Records

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

CPU MF

- We recently performed performance reviews for two customers, and tripped across some very interesting insights.
- One of the customers was running flat out nearly all the time, and wanted some advice on possible upgrade options.
- Because their CPC has just 2 general purpose CPs, and they have only 2 LPARs (1 production, 1 development), they didn't think that there was any value in enabling CPU MF and the SMF type 113 records – after all how much LPAR configuration tuning can you do with 2 CPs and 2 LPARs?

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

CPU MF

- They considered their CPC to be about 900 MIPS, based on IBM's LSPR numbers. So their search for alternatives was based on 900 MIPS + x% for growth.
- We managed to convince them to enable CPU MF and collect the type 113 records for both LPARs.
- According to LSPR, the range of MIPS for their processor runs from about 800 MIPS for a High RNI workload, to about 970 MIPS for a Low RNI workload.
- Using their SMF type 70 and 113 records to get an accurate MIPS number for THEIR workload, we discovered that the capacity they were getting from their CPC was actually less then 800 MIPS – 100 less than they thought, and at the bottom end of the 170 MIPS range of capacities for their CPC.
- Armed with this knowledge, rather than spending money on purchasing a hardware upgrade AND paying all the associated SW Upgrade charges, they launched a project to reduce their RNI.

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

CPU MF

- What lessons do we learn from this?
  - EVERYONE should have SMF 113 records enabled ALL the time for ALL LPARs.
    - The ‘overhead’ is so small that it is not measureable, and the volume of data is tiny compared to their total volume of SMF records.
    - They contain information that cannot be determined by any other means.
    - And make sure that you keep them!
    - We will not work with a customer to study upgrade options if they do not have the type 113 records enabled.
  - If you have a High RNI workload/configuration, try to reduce the RNI before you spend money on an upgrade.
  - There are LOTs of things that you can do to improve your RNI.
  - And LOTs of things to be gleaned from SMF (199 slides worth) – you just need the time and the tools and the curiosity.

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

CPU MF

- A few other tips:
  - If you have a more complex LPAR configuration, enable the SMF Type 99 subtype 14 records to help you understand logical CP placement.
    - If possible, we recommend collecting other SMF type 99 records as well. They are produced by WLM every 10 seconds even if you don't collect them, so the only additional overhead is the cost of writing them to SMF. They are very useful in understanding short term system utilization/behavior. By properly selecting specific subtypes you can keep the recording volume below 100MB/Day/System. For advices about subtype selection see Session 24243, *Introduction and Using High Frequency SMF 98 and 99 Measurements*, by **Peter Enrico**.
  - Remember that IBM's zPCR tool does NOT factor in the knock-on benefits of reduced CPU utilization if you move to a larger CPC.
    - And if you move to a SMALLER CPC, especially one with fewer CPs, don't forget to allow for the negative impact of higher utilization than you have today.

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

CPU MF

- The high cost of high CPU utilization:

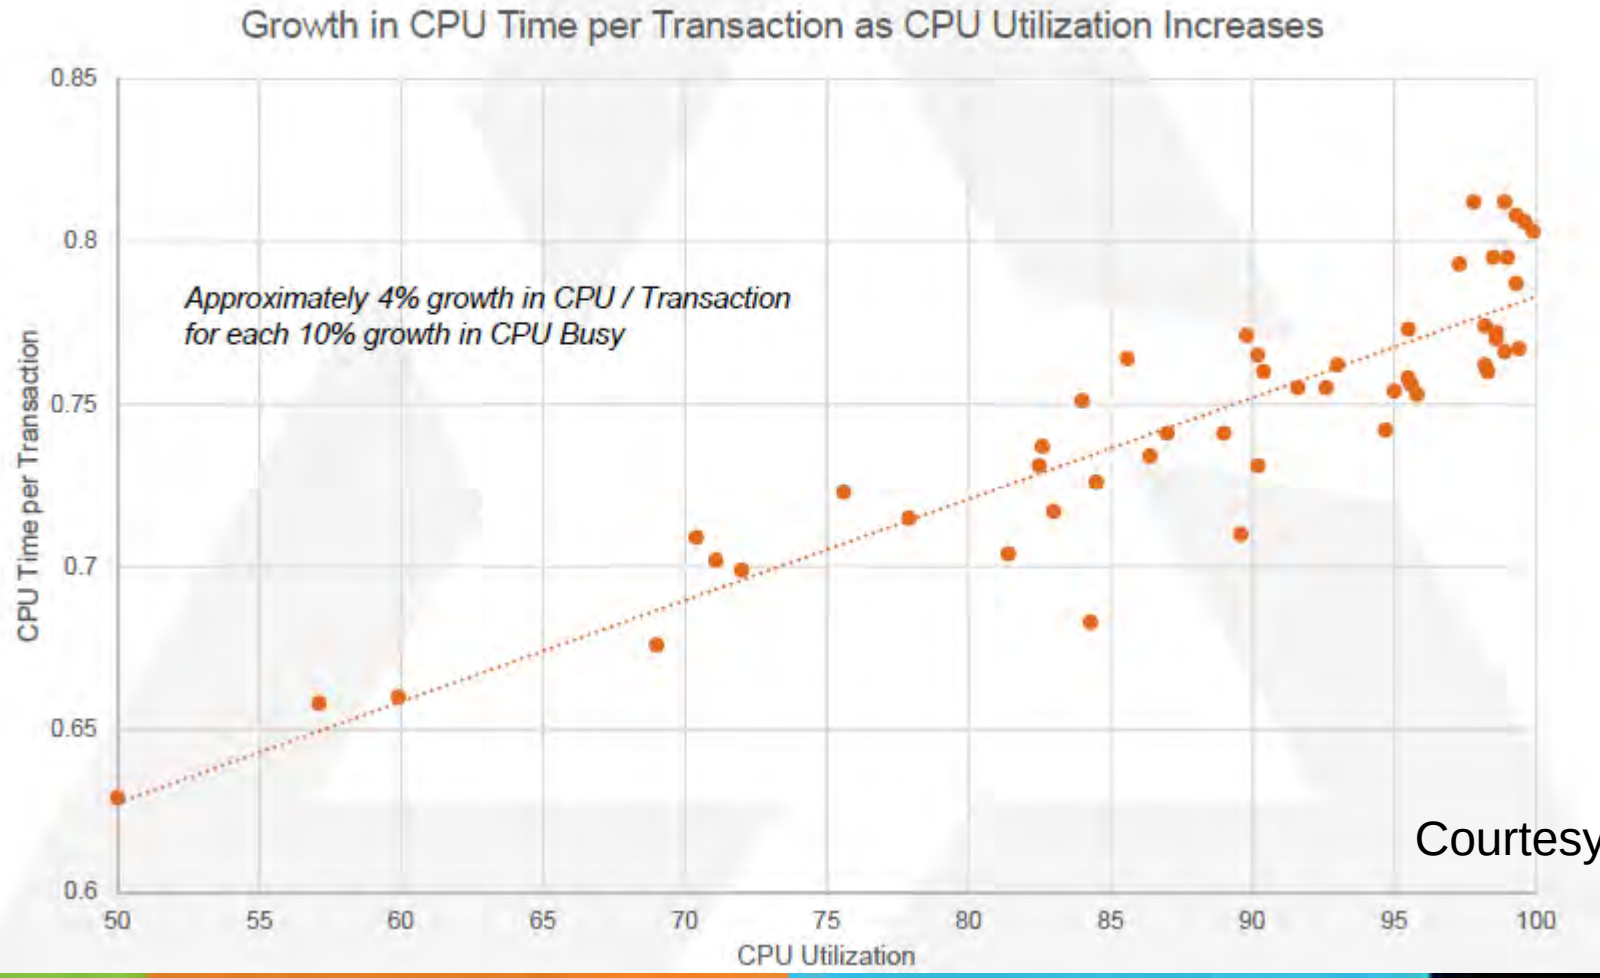

Courtesy IBM's Dave Hutton

Help!

- *None* of IBM's capacity planning tools currently attempt to predict how any change in the number of CPs in your CPC will affect your RNI.
- As described in our zRoadshow slides from SHARE in St. Louis, NASCO moved from a z13 709 to a z14 523. According to zPCR, the 523 should have delivered 1.7% more capacity. Instead, their peak MSUs dropped by 22% (yes, TWENTY-TWO percent!) after the upgrade. Do not underestimate the impact that changing the number of CPs can have on your performance.
- We are doing some work with IBM to see if they can provide some guidelines to help customers factor potential RNI changes into your upgrade plans.

- To help validate their theories, they are looking for SMF type 70 and type 113 records from customers that changed to CPCs with a significantly different number of CPs.
  - If you would like to contribute some data (and help yourself and the entire community in the process), please contact us at technical@watsonwalker.com.

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

SMF 98 High Frequency Throughput Statistics

Kathy Walsh – S24672

- Available since z/OS 2.2 + APAR
- WSC and IBM Pok lab both recommend enabling High Frequency Throughput Statistics (HFTS) SMF type 98 records.
  - The MVS Init & Tuning Reference currently recommends setting the HFTS interval to 20 seconds (HFTSINTVL in SMFPRMxx).
  - However, because the overhead is very low and the volume of SMF records is not very large, z/OS Development recommend 5 seconds, and hope to have the manuals updated later this year to reflect that.
- Currently exploited to collect Supervisor performance data (Subtype 1)
  - SMF 98 Subtype 1 records contain performance information about the workload and its significant jobs. It includes metrics such as utilization, concurrency, efficiency, contention, and queuing.
- HFTS data aids in diagnosing transient problems that are hard to spot in multi-minute intervals.

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

Reducing HSM Resource Consumption

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

HSM CPU

- HSM tends to sit in the background and quietly do its job of backing up your data and ensuring that your primary volumes always have space for new data set allocations.
- But HSM can move A LOT of data, and moving (and compressing) data costs CPU time.
  - In some customers, HSM can be one of the largest single consumers of CPU capacity.
- To help customers minimize the CPU cost of HSM, **Glenn Wilcock's** Session 24254, *Minimizing HSM CPU and Elapsed Times*, provides some excellent advice on ways to use standard HSM features to reduce HSM's impact on your R4HA.
  - We also worked with Glenn on an article on this topic for the most recent Tuning Letter (2018 No. 4), based on performance reviews with two customers.
- See the DFSMS Report Generator for help with formatting the HSM SMF records.

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

HSM CPU

- Summary of Glenn's advice:
  - Adding zEDC is the single best thing that you can do to reduce HSM CPU consumption (more anon).
  - Ensure that Fast Subsequent Migration is enabled and optimized.
  - Monitor/Reduce HSM Thrashing.
  - If you have a virtual tape subsystem, consider eliminating ML1, and just migrate directly to ML2.
  - If you have disk subsystems with multiple media types (Performance Flash, Capacity Flash, spinning disks), consider replacing Migration to Level 1 with Class Transitions – these use FlashCopy to perform the data movement, rather than z/OS CPU time.
  - Learn about Transparent Cloud Tiering – even if you don't want to use it today, the concept of moving data directly from disk subsystem to 'the cloud' (which could be your own TS7700) is the way of the future.
  - See Glenn's presentation for a PATCH command to help you identify HSM exceptions that waste CPU time.

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

zEDC

- z Enterprise Data Compression
  - Offloads compression and decompression work to PCIe cards (so that processing doesn't count towards your R4HA).
  - Generally achieves better compression ratios than traditional DFSMS compression.
- You need to purchase a minimum of 2 cards. IBM recommend a minimum of 4.
  - Also requires a chargeable (but cheap) feature on z/OS.
- Exploited by DFSMS for sequential data sets, SMF for its logstreams, HSM and DSS for migration and backup, zFS, Connect Direct, MQ, others.
- On z13, turning on zEDC for a data set that is currently *not* compressed results in a small increase to small decrease in CPU time.
- On z14, turning on zEDC for an *un*compressed data set nearly always results in less CPU time.
  - Obviously, on both z13 and z14, replacing traditional compression with zEDC will reduce CPU consumption.
- We are not aware of any customer that has zEDC that is not delighted with it.

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

zEDC

- IBM's free zBNA tool estimates benefits of zEDC, but ONLY for sequential data sets, and NOT for any HSM activity. This is because:
  - zBNA uses information from the SMF Type 14/15 records.
  - HSM uses DSS to migrate or backup sequential data sets, and DSS does not open the data set, so you don't get any Type 14/15 records.
- HSM DOES create its own SMF records (optionally), and they contain information about every data movement by HSM.
  - However, those records do *not* provide any information about whether the data sets are compressed or not.
  - If you don't know whether the input data set is already compressed, you can't estimate how much space or CPU time zEDC could save.
- To try to partially address this gap, we opened SHARE requirement MVSS 170, requesting that this information be added to the HSM SMF record. If this data would be valuable to you (regardless of zEDC), please review the requirement.

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

Thoughts on Pervasive Encryption

- Encrypting and decrypting data is very CPU-intensive. Much better on z14 than on z13, but still heavy on CPU.
- In general, data is decrypted when it is read into memory, and encrypted when written out.
- So, anything you can do to reduce the volume of data will reduce the CPU cost.
  - zBNA now provides the ability to estimate the cost of encryption on the current data set, and the cost if you compress the data using zEDC.
  - In one example in **John Burg's** zBNA lab (Session 24150) enabling encryption for a particular data set increased CPU time by 1.2 hours to *without* zEDC, and by just 13 minutes *with* zEDC.
  - Naturally YMMV, but make sure that you use this option in zBNA when evaluating the cost of encryption.

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

Thoughts on Pervasive Encryption

- PDSE V2 now supports encryption. To reduce the CPU cost, **Tom Reed's** Session 24327 recommends using Buffer Beyond Close and enabling PDSE member caching, both of which reduce the number of times the data must be read into memory, and therefore the cost of encryption.
- Various encryption measurements metrics are available via SMF 113 Crypto Counters, SMF 42.6 DFSMS data set statistics, SMF 82.31 Cryptographic usage statistics.
- New SMF70 LACCR field reports on system wide long-term average of CPU service (MSUs) consumed by DFSMS data set encryption. This is available on z/OS 2.1 and later via PTF (See APARs OA54404 and OA56097).

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

Pervasive Encryption Planning and Measurement

John Burg – S24533

zEDC Compression and *Then* Data Set Encryption Top Candidates

New

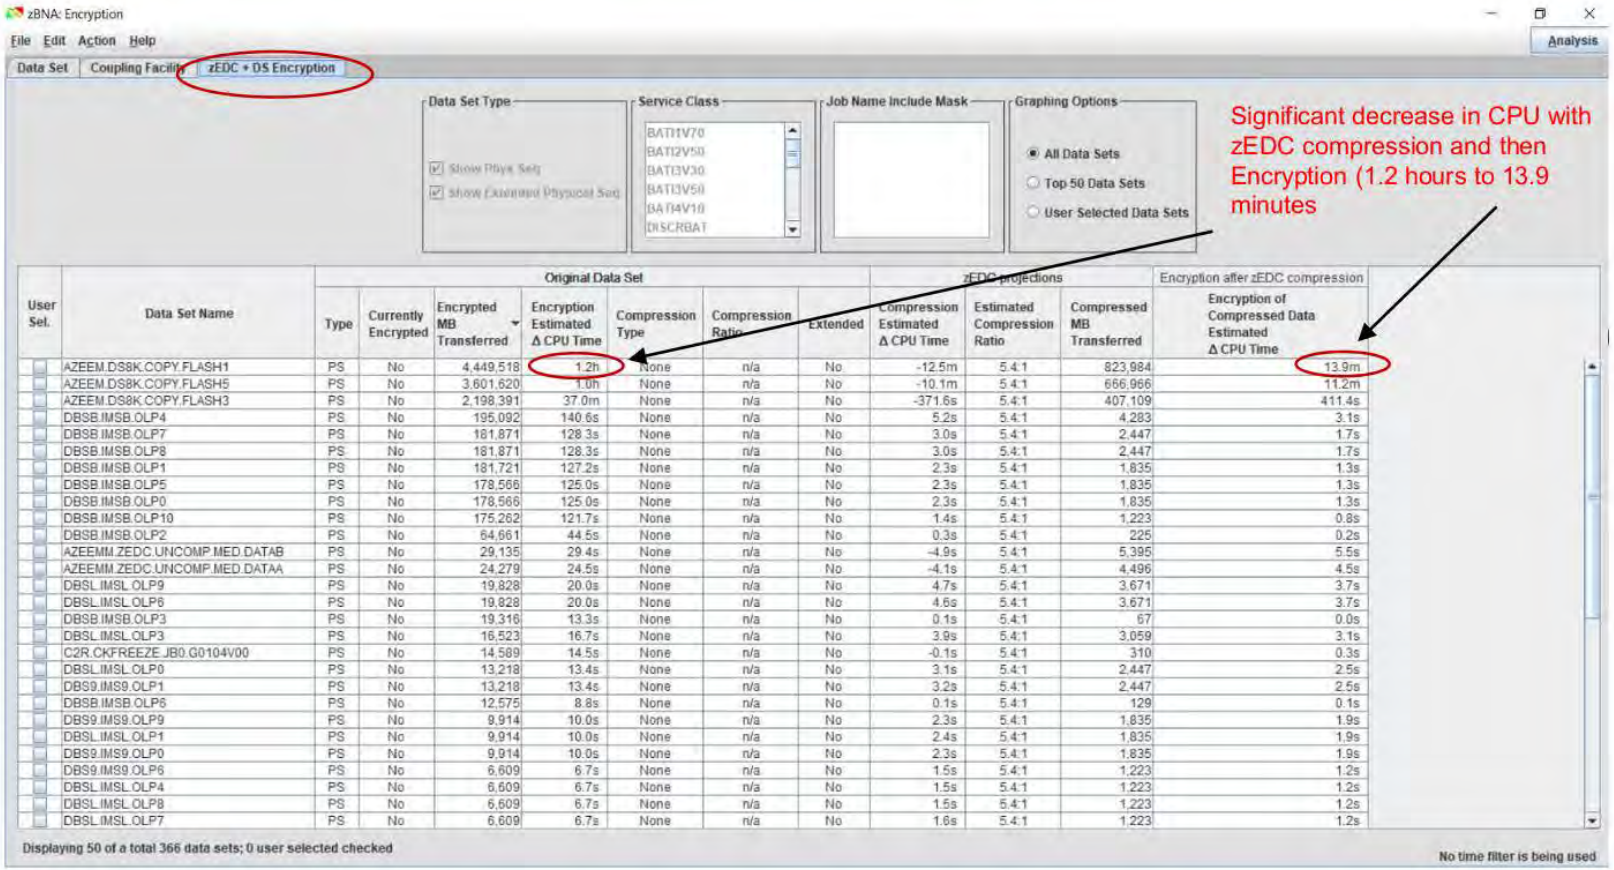

© 2019 IBM Corporation

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

Encrypted data sets and TS7700 Tapes

Kathy Walsh – S24672

- When writing encrypted data sets to tape, if compaction isn't disabled by the SMS dataclass, the TS7700 will attempt to compress them. Using the default FICON algorithm, they will expand by 12.5%

- Solution is to change your dataclass definitions to either disable compaction or to switch to the LZ4 or ZSTD algorithms introduced by TS7700 R4.1.2
  - LZ4 compression delivers moderate improvements in compression efficacy while consuming a small amount of the TS7700 server's processing power.
  - Zstandard compression delivers greater improvements in compression efficacy while consuming more of the TS7700 server's processing power.

- If can't disable compaction, then using LZ4 would be best for these workloads since it won't expand, has lower overhead, but will still compress those records which are not encrypted

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

Asynchronous CF Lock Duplexing

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

Asynchronous CF Lock Duplexing

- IBM announced Asynchronous CF Lock Duplexing in 2016.
  - Requires z13 or later CF, z/OS 2.2 or later, and Db2 V12.
- **Currently only supported by Db2 V12**, *and* only for its lock structure.
- The objective is to largely eliminate the CPU and elapsed time impact of traditional System-Managed Duplexing for Db2 lock structures.
  - Rather than z/OS sending requests to both CFs, and having to wait for the request to be completed in *both* CFs, with Asynchronous CF Lock Duplexing, XES only sends the request to the primary structure. Mirroring to the secondary is driven by the CF containing the primary structure and happens asynchronously, *after* the response to the original request has been sent back to z/OS.
- Today, most System-Managed Duplexed requests are sent to the CF asynchronously – this increases response times, and z/OS CPU time compared to a short synchronous request. Async CF Lock Duplexing should enable most requests to a local lock structure to be processed synchronously.
- Let's have a look at a real world example of the impact of Async Duplexing.

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

Asynchronous CF Lock Duplexing

This data is from Nedbank in South Africa. They have two z14s, internal CFs, both z14s are in the same data center (so no long distances involved), and they were originally using traditional System-Managed Duplexing.

Thanks to **Daniel Hamiel** for his help.

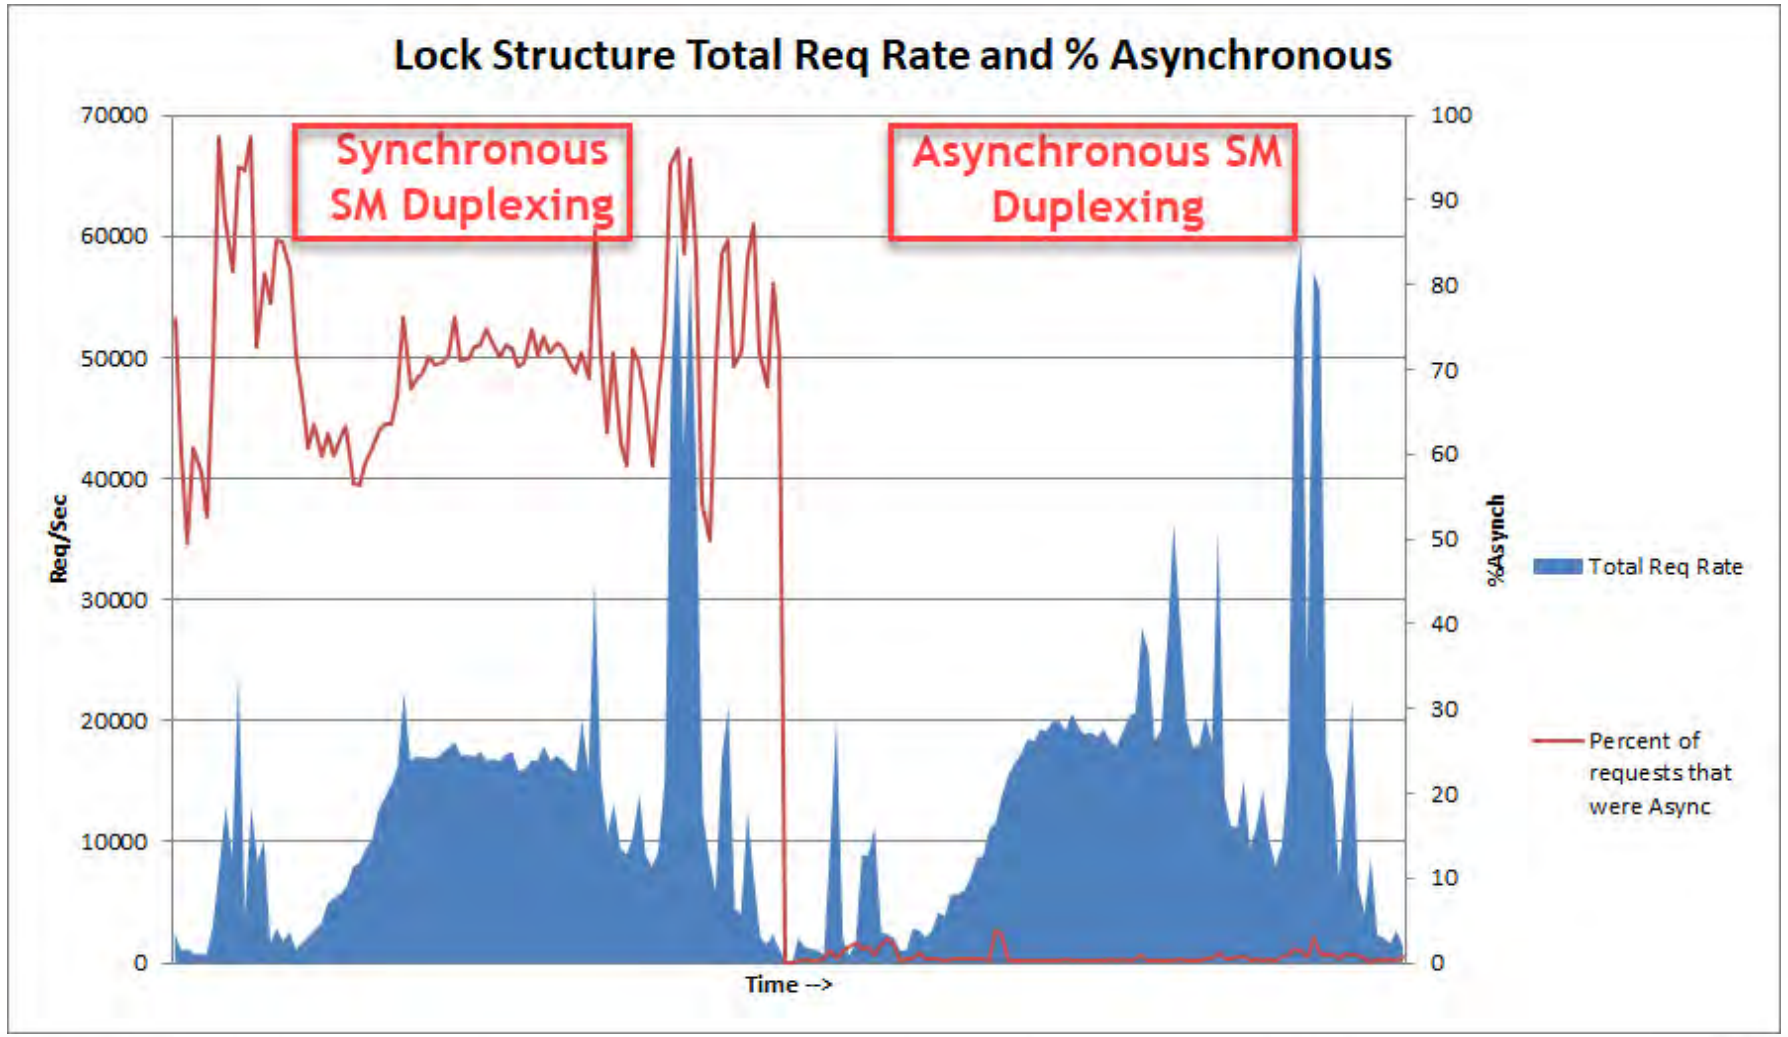

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

Asynchronous CF Lock Duplexing

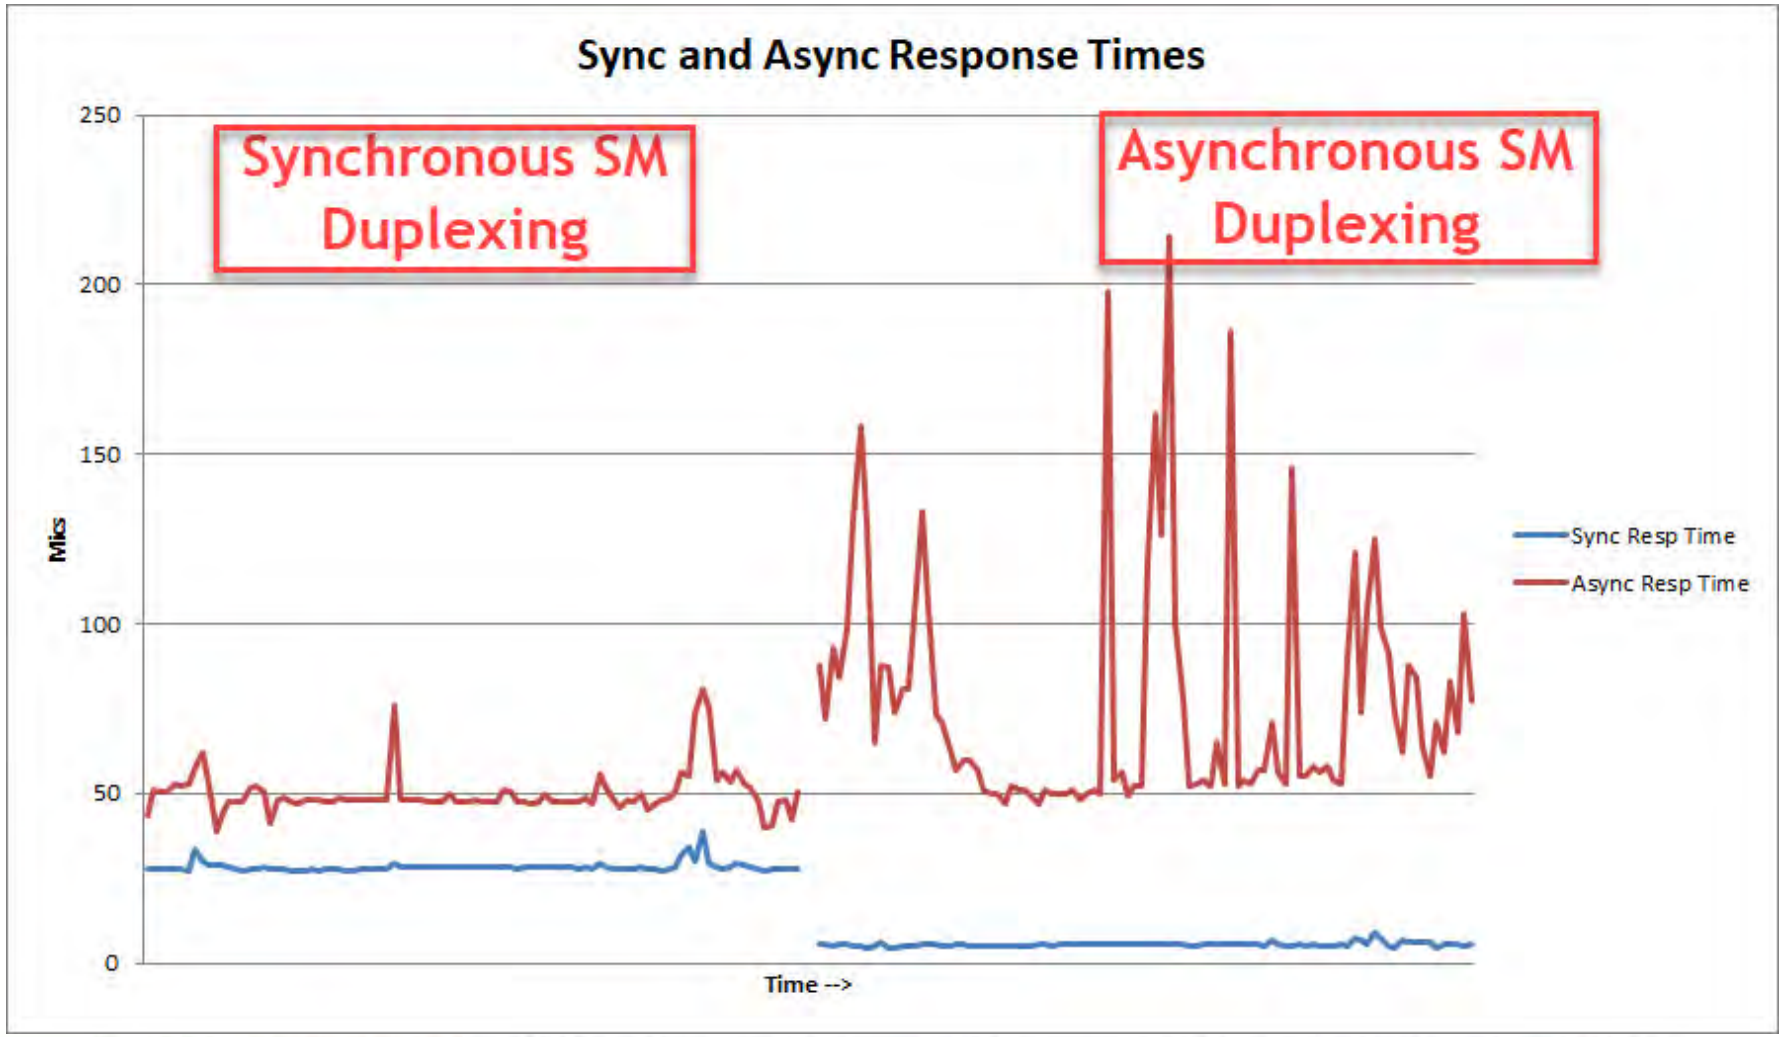

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

Asynchronous CF Lock Duplexing

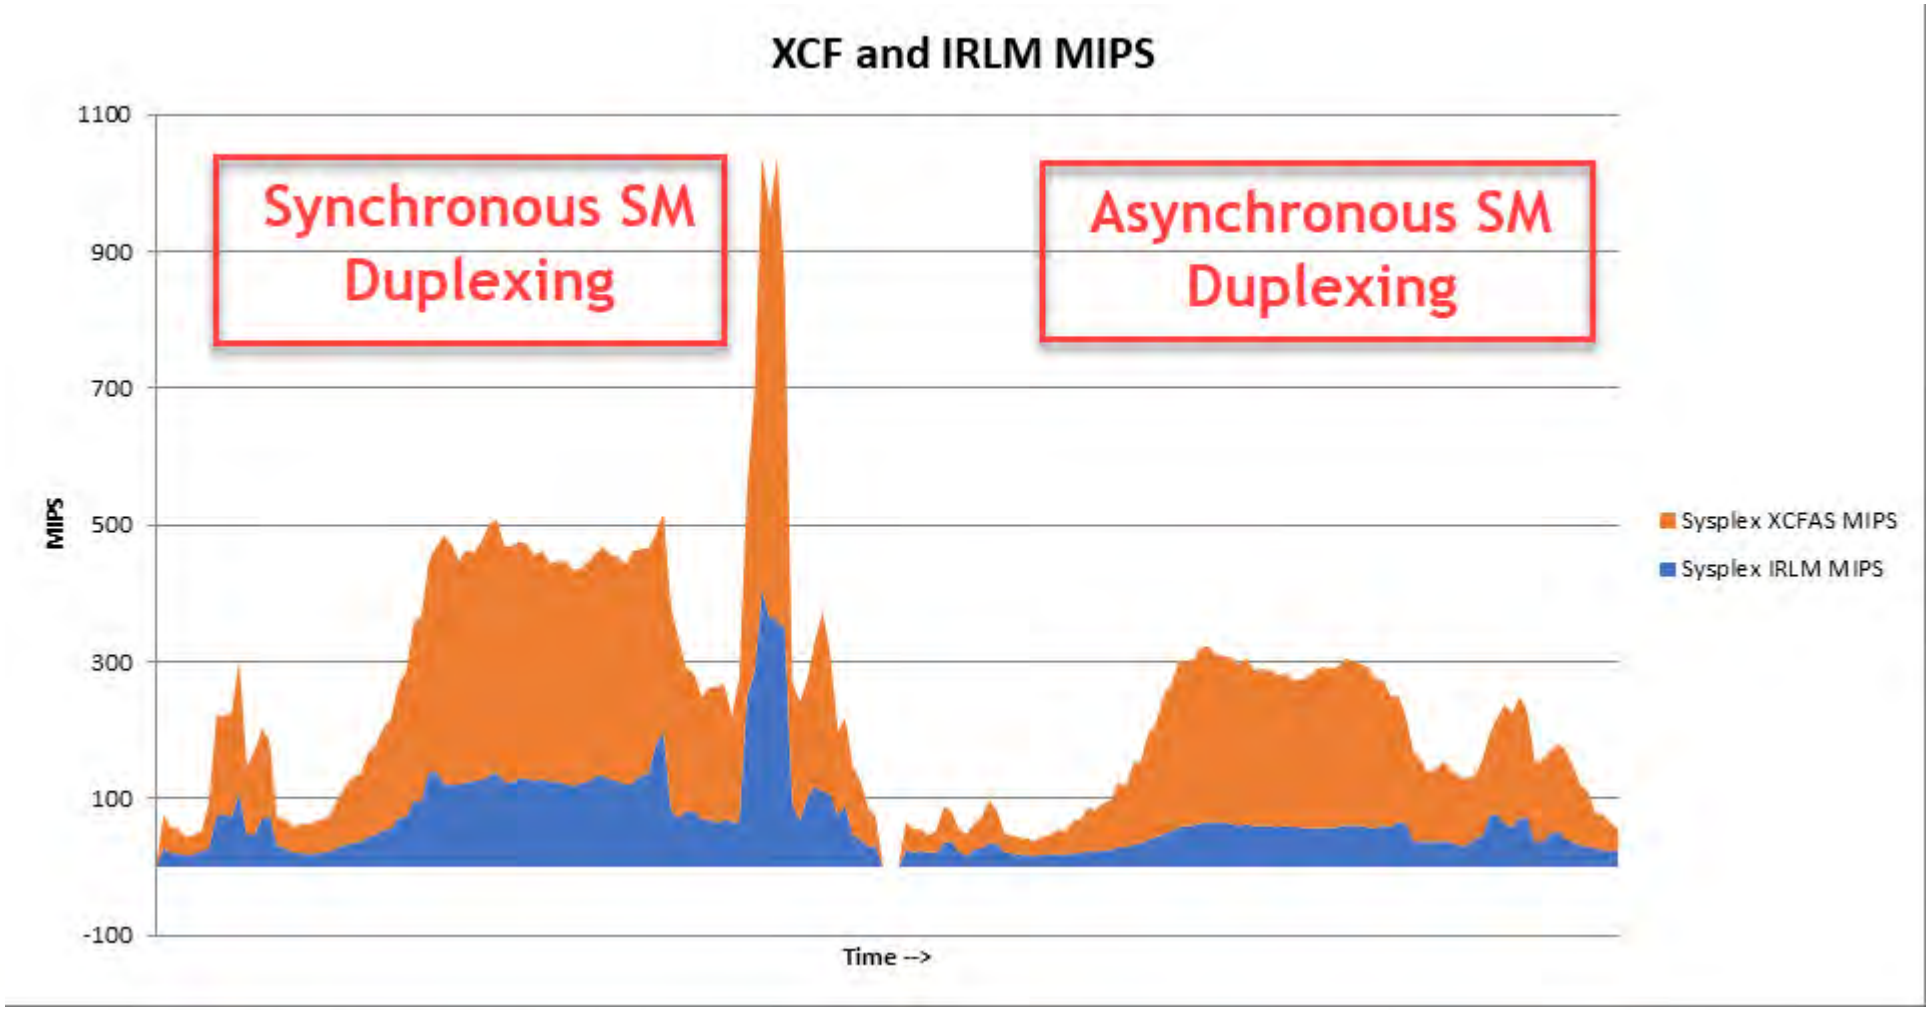

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

Asynchronous CF Lock Duplexing

- This is the first customer that we have worked with that has implemented Async CF Lock Duplexing.
  - Customer takeup has been slow due to requirement for Db2 V12.
- BUT, this customer was really delighted with the results.
- Based on these results:
  - We find it hard to see a valid reason why sites that are using System-Managed Duplexing for Db2 Lock structures today would *not* evaluate Asynchronous CF Lock Duplexing.
  - We believe that anyone that has a simplex, NOT failure-isolated, Db2 lock structure today should consider implementing Asynch CF Lock Duplexing.
- For more information, Tuning Letter subscribers should refer to Asynchronous CF Duplexing article in Tuning Letter 2018 No. 4.

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

Asynchronous CF Cross-Invalidate

- CF Level 23 (z14 GA2) introduced the option for updates to a cache structure to issue cross-invalidates asynchronously.
  - Prior to this, updates to a cache structure could not complete until a cross-invalidate signal is sent to every CPC that contains a local copy of the page being updated, AND acknowledgement is received back in the original CF.
    - If the target CPC is a long distance from the CF, this could *significantly* elongate these update requests.
  - Asynch XI sends the response to the requesting z/OS system as soon as the request is completed in the CF, without having to wait for the acknowledgement back from the target of the XI signal.
  - This should be especially attractive to customers with multi-site sysplexes.
- Requires z14, z/OS 2.2 or later with OA54688, and Db2 V12 with APAR PH05193 plus another, yet to be announced, Db2 APAR.
- We hope to have more news on this by the time of the next SHARE in Pittsburgh.

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

zHyperWrite

- zHyperWrite is intended to reduce the impact of synchronous disk mirroring for certain disk writes.
  - Rather than the write going to the primary subsystem, and the primary then mirroring the update to the secondary, simplex writes are sent to the primary and secondary devices in parallel.
  - The effective response time is the longer of the writes to the primary and secondary devices.
- zHyperWrite is supported by z/OS 2.1 & later, Db2 V10 and later, IMS V15, and System Logger. There is also a Statement of Direction for MQ support. A DS8870 with microcode level 7.4 or later is required, and HyperSwap must be enabled for the devices containing the data sets that use zHyperWrite.

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

zHyperWrite

- Configuration with two data centres, 20km apart, and use Db2 and IMS.
- Db2 subsystem DBP1 runs in the same data center as the primary disk subsystem.
- DBP2 runs in the same data center as the secondary disk subsystem.

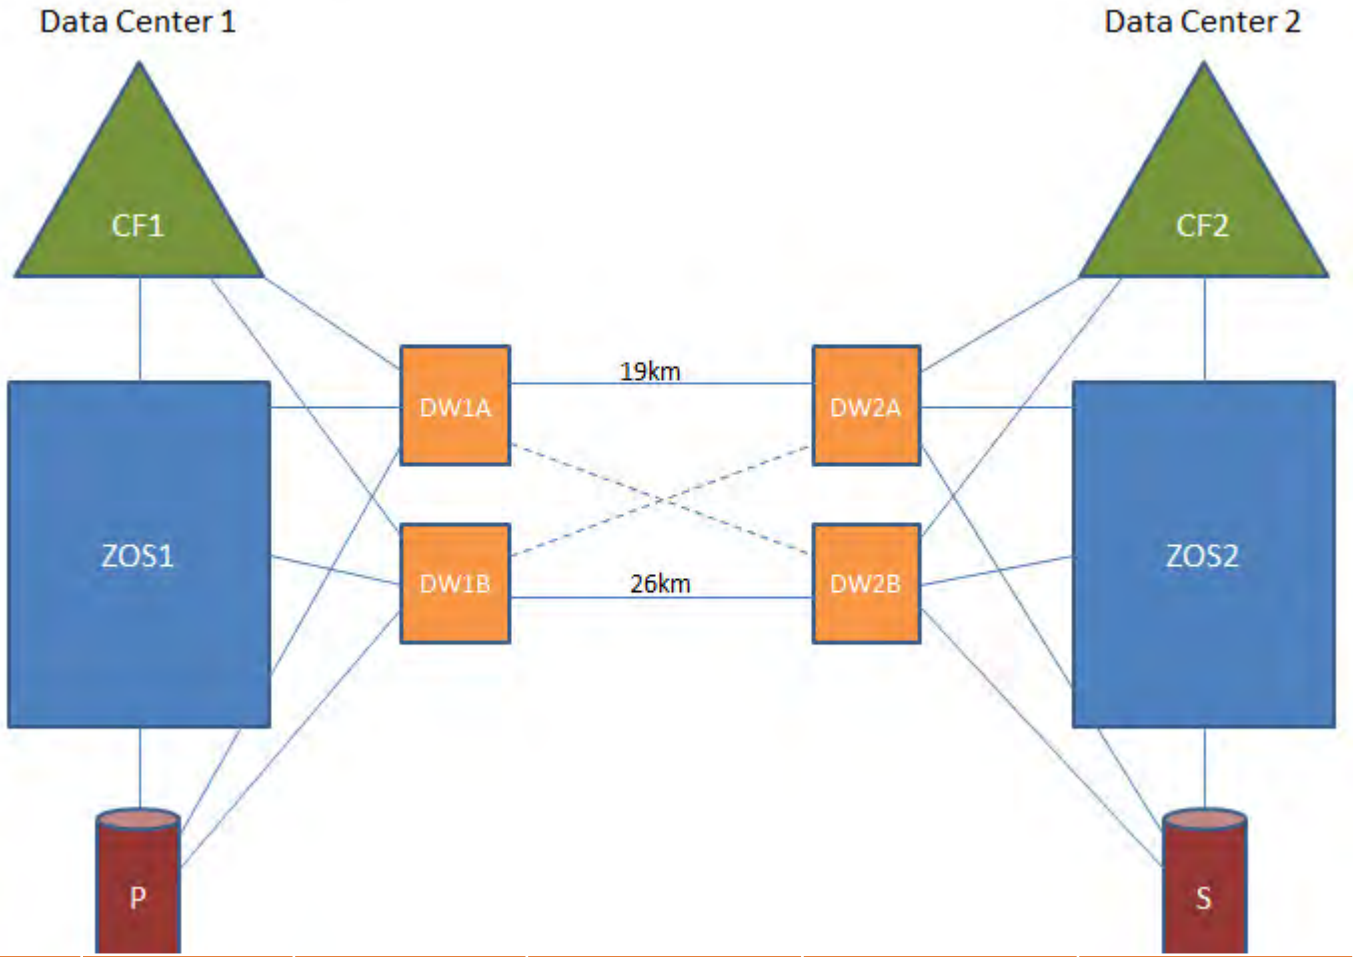

DB2 Subsystenm	Location	I/O Rate	Response time before zHyperWrite	Primary disk Response time after zHyperWrite	Secondary disk Response time after zHyperWrite
DBP1	DC 1	100-700/Sec	.6ms	.16ms	.4ms
DBP2	DC 2	100-150/Sec	.9ms	.4ms	.16ms

Copyright© by SHARE Association Except where otherwise noted, this work

Highlights From Phoenix

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

Recommendations This Week

- *z/OS Performance Hot Topics* – 24672 – **Kathy Walsh**
  - Always a must-see session.
  - Keeping SMF 113 records should be a requirement, NOT an option! This requires running CPU MF and HIS and turning on 113 records in SMFPRMxx.
  - New service definition coefficient “forward thinking” recommendations: CPU=1, SRB=1, MSO=0, IOC=0. “Trust me, you will thank me later”
  - Recommend changing SMFPRMxx parameter NOSMF30COUNT to SMF30COUNT.
    - This can help identify application changes.
    - Also provides the data to let you calculate the cycles per instruction (CPI) = (cpu time) / SMFCOUNT for that program.
      - This has been used to find possible SIIS (Store In Instruction Stream) culprits. (See Cheryl Watson’s Tuning Letter 2016 No. 1, pages 14-16)
  - 38% performance improvement in COBOL app with a lot of packed decimal on ZR1/z14 between COBOL 6.1 and 6.2, with just a recompile. See WSC White Paper WP102731.
  - **NOTE: COBOL V4 EoS is September, 2021. See announcement letter ENUS919-022**

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

Recommendations This Week

- Lists of VERY important APARs that EVERY site should know about:
  - MVS Program Opening - 24581 – **John Shebey**
  - MVS Performance Project Opening - 24672 – **Kathy Walsh**
  - What's New in JES2 - 24628 - **Tom Wasik**

- Important Enhancements:
  - What's New in DFSMS and MVSS Project Opening – 24345 – **Barbara McDonald**
  - Are You Leveraging the Latest DFSMShsm Enhancements – 24269 – **Glenn Wilcock**
  - What's New With DFSORT? – 24248 – **Sri Hari Kolusu**

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

OA53355: USERKEY COMMON MIGRATION

Kathy Walsh – S24672

Problem summary

The allocating, obtaining and changing common areas of virtual storage, such that the storage is in user key (8-15), will not be supported after z/OS V2R3. The currently documented methods for identifying user key common storage usage may be too disruptive.

In addition, there is no assistance provided to identify user key common storage usage created via the CHANGKEY service.

Additional ways to identify user key common storage usage are needed.

If you are running CICS Transaction Server for z/OS, ensure that you are running V5.2 or later version.

Temporary fix

This APAR provides additional methods for identifying user key common storage usage. The methods include:

- SMF Type 30 records were enhanced to identify jobs/steps that use user key common storage.

- Allow installations to set a singular SLIP trap to catch ALL user key common storage usage. Previously, different types of user key common storage usage require different SLIP traps. This addresses the restriction that multiple PER type SLIPs could not be active at the same time. Previously, there was not a SLIP trap for user key common storage created via the CHANGKEY operation.

- A new migration health check was created that will generate an exception message if user key common storage is used.

In addition, the severity of existing health check, VSM_ALLOWUSERKEYCSA, has been raised to MED.

https://www-01.ibm.com/support/docview.wss?uid=isg1OA53355

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

Some interesting z/OS Tuning tips

Peter Enrico and Scott Chapman – S24832

- Large memory configurations, being considered one of the key enablers for better performance, put increased pressure on System z virtual address translation mechanisms.

- For this reason IBM introduced 1MB and 2GB Large Pages, and have been encouraging customers to use them for years. If you didn't do it yet we urge you to consider them.

- In one of his charts Scott presented a zEC12 customer case where over 7% of the overall CPU capacity was being used for address translation, we also see similar cases with our customers.

- z14 helps here, as it introduced a new TLB engine which allows up to four concurrent translations with each translation step being ~2x faster compared to the previous design. But even on a z14, the cost of translating virtual addresses is higher when using 4KB pages. For this reason implementing large pages is important, especially for shops using large buffer pool configurations.

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

Some interesting z/OS Tuning tips

Peter Enrico and Scott Chapman – S24832

- Peter made another interesting tuning consideration about using high velocity goals:
- WLM finds it difficult to help periods having high velocity goals and getting close to their objective. This is because WLM only takes actions which it expects will make significant changes, and it is hard to significantly increase the velocity of such periods.
- In this case, one option you may consider is to disable WLM I/O Priority Management.
- When WLM I/O Priority Management is enabled, the Velocity calculation includes I/O Using and Wait samples:

$$\text{Velocity} = \frac{\text{CPU Using} + \text{I/O Using}}{\text{CPU Using} + \text{CPU Delay} + \text{I/O Using} + \text{I/O Delay}}$$

- I/O wait samples only include IOSQ and Pending times, which nowadays are usually very small. This means that I/O Using samples tend to inflate the Velocity, and on the other hand that the I/O Priority Queuing function provided by WLM can't provide much help.
- By disabling WLM I/O Priority Management, you take I/O Using and Wait samples out of the equation. This will have the effect of decreasing the Velocity of your workload, meaning that you should review your WLM goals. However, WLM has more control over the contributors to the Velocity, CPU Using and CPU Delay, which is good.

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

No Infiniband Coupling on Z servers beyond z14

Kenneth Stine – S24669

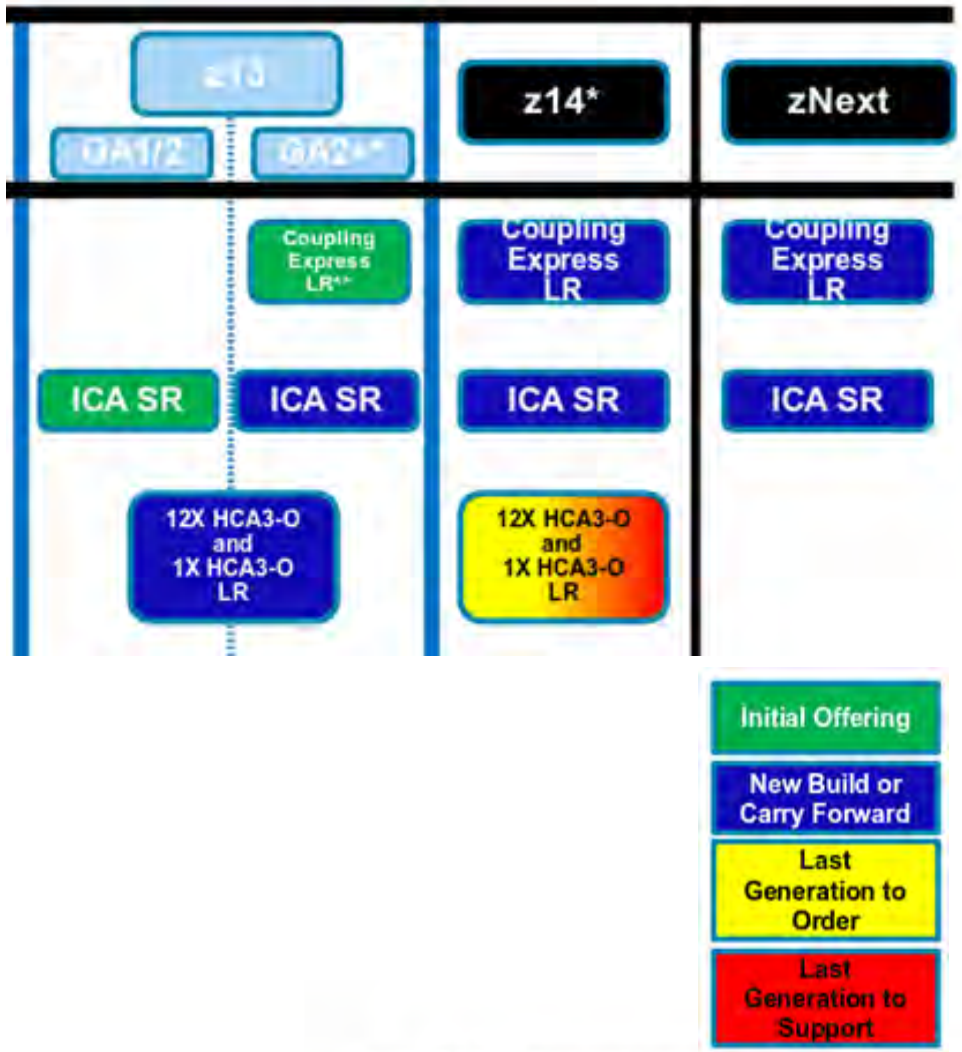

The timeframe for migration from Infiniband coupling links to current Coupling Technologies is really short!

If not yet there, you need to acquire and plan the installation of new ICA SR short-distance / CE LR on your z13s and z14s in order for them to be able to participate in a Parallel Sysplex with zNext.

Last day to place MES order for z13 is June 29 2019 *

http://www-01.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/1/897/ENUS918-071/index.html&lang=en&request_locale=en

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

zIIPs Beyond the basics..

Kathy Walsh – S24534

About capacity planning:

- zIIPs are NOT expected to run at high utilization rates. You simply want to add more zIIP capacity when the amount of Crossover to GPs makes it financially sound.
- Adding zIIPs to a configuration changes its MP ratio and reduces GP capacity. Use zPCR to estimate the effect, but make sure to set realistic zIIP loading values as the default is 100%.

About offloading:

- Marking a UoW zIIP eligible doesn't immediately move it from one processor type to another. Moving a TCB from one engine to another is expensive as it reduces L1 and L2 caches efficiency of both.
  - For this reason, z/OS implements a 'lazy switch' technique: When marked zIIP-eligible, a UoW keeps running on the GP until it gets interrupted or its time slice ends. In the same way when it is marked as no longer being zIIP-eligible, it stays on a zIIP until it gets interrupted or its time slice ends.
- DB2 uses a similar "lazy" technique for DDF. Instead of making every DDF thread 60% zIIP eligible, it marks 100% zIIP-eligible for 3 out of 5 DDF threads.

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

zIIPs Beyond the basics..

Kathy Walsh – S24534

About Crossover to GPs

- The IIPHONORPRIORITY parameter in IEAOPTxx determines, at the system level, if GPs can help process zIIP-eligible work to alleviate delays.
- DB2 v11 and later only exploit zIIPs for certain Db2 tasks if IIPHONORPRIORITY is set to YES.
- It is now possible to effectively set the IIPHONORPRIORITY to NO for individual WLM service classes. This setting overrides the global one.
- In a HiperDispatch configuration, deciding to unpark a logical processor is an action which takes a few seconds. On the other hand, the "need help" algorithm which may trigger a crossover to GP runs every few milliseconds. This means that bursty zIIP workloads may result in a significant amount of zIIP-eligible workload running on GPs. Same can happen in configurations with few zIIPs.
- A Discretionary zIIP eligible workload won't get help from CPs, even if IIPHONORPRIORITY is set to YES.

Various

- WLM Resource Group limits may now optionally include zIIP consumption.
- In any case, when a WLM Resource Group is capped, *all* the associated workload gets capped, regardless of whether it is zIIP-eligible or not.
- SRM doesn't provide Blocked Workload support for zIIP workloads

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

JES2

- Refer to Session 24628, *z/OS V2.3 JES2 Product Update and Latest Status*, by **Tom Wasik**.
- z/OS 2.4 is the last release that will support the z11 checkpoint format. Z22 mode was introduced by z/OS 2.2. If you have not already migrated to z22 format, recommend that you do so. No known problems with this migration.
- JES2 in z/OS 2.4 adds the ability to encrypt the JES2 spool.
  - As part of that support, zEDC compression support is also being added.
  - You can use the compression support without having to encrypt the spool.
  - Requires all systems in the MAS to be running z/OS 2.4 or later.
- z/OS 2.4 starts a loooong journey to replace JES2 exits with an XML-format policy.
- New Checkpoint Health Check.
- z/OS 2.5 will be the last release to support JES3.

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

Important APARs

- The following CA (Broadcom) APARs address potential performance issues in CA-View 14.0:
  - ST07584 PERFORMANCE ISSUE WHEN SORTING REPORTS BY DATE IN JOB MODE
  - ST06106 SARBCH DDNAME INEFFICIENT BLOCKSIZE
  - SO00832 IMPROVE AFP INDEX PERFORMANCE
  - SO00445 IMPROVE AFP INDEX PERFORMANCE
  - SO00181 SLOW RESPONSE ACCESSING INDEX FOR REPORT
  - RO97820 ENHANCE INDEXING I/O PERFORMANCE
  - RO97347 INCREASED CPU TIME ACCESSING REPORT INDEX

  - The first two APARs are still in test.

- Thanks to **Ed Blazejewski** in Broadcom for his help with this information.

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

Important APARs

- The following IBM APARs should be investigated:
  - OA55320 – Stops HSM/DSS trying to compress a data set that as already been compressed by zEDC.
  - OA55190 - Potential performance improvement for large clients using zEDC compression for DB2 & MQ log offload & restore.
  - OA56229 – JES2 Excessive CPU in HASPLIM subtask
  - OA55846 – JES2 delays when processing large number of commands
  - OA55591 – JES2 High CPU Usage After Starting Multiple Jobs With SCHEDULE DELAY=YES
  - OA54973 – Changes to PDSE Monitor to Reduce CPU Consumption
  - OA51549/OA51551 zSecure APARs. Make sure to explicitly specify PRECONSOLIDATE in zSecure 2.2.0 or 2.2.1.

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

What’s New in IBM Documentation

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

What's New in IBM Documentation

- Last SHARE
  - User community complained strongly after web pages and manuals were deleted and moved without backup or recovery due to IBM-wide marketing initiative.
- Since Then
  - IBM sent EIGHT people from the Poughkeepsie-based Content Development team to this SHARE, we hope that you took this opportunity to let them know what you need to do your job.
    - See our blog post for a list of all their sessions.
  - Most of the missing Redbooks have been restored.
  - Many of the missing links have been redirected (e.g. WLM, RMF, z/OS zFavorites)
  - Older z/OS releases have been restored
  - Out of support z/OS release PDFs now available in a single zipped indexed file (back to z/OS 1.13); individual PDFs back to z/OS 1.10 (WOW!) Also true for z/VM.
  - Before z/OS 2.3, PDFs were updated quarterly (up to every two years). Now they're updated as soon as the PTF is released. To register to be notified of PDF updates, go to: https://www-01.ibm.com/servers/resourcelink/svc00100.nsf/pages/zOSV2R3Library.
  - KC has been significantly enhanced – many small changes that make it much more efficient to use – see *Tuning Letter 2018 No. 3* article Knowledge Center – Your New Old Friend for more information.

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

What's New in IBM Documentation

- Improvements in Knowledge Center

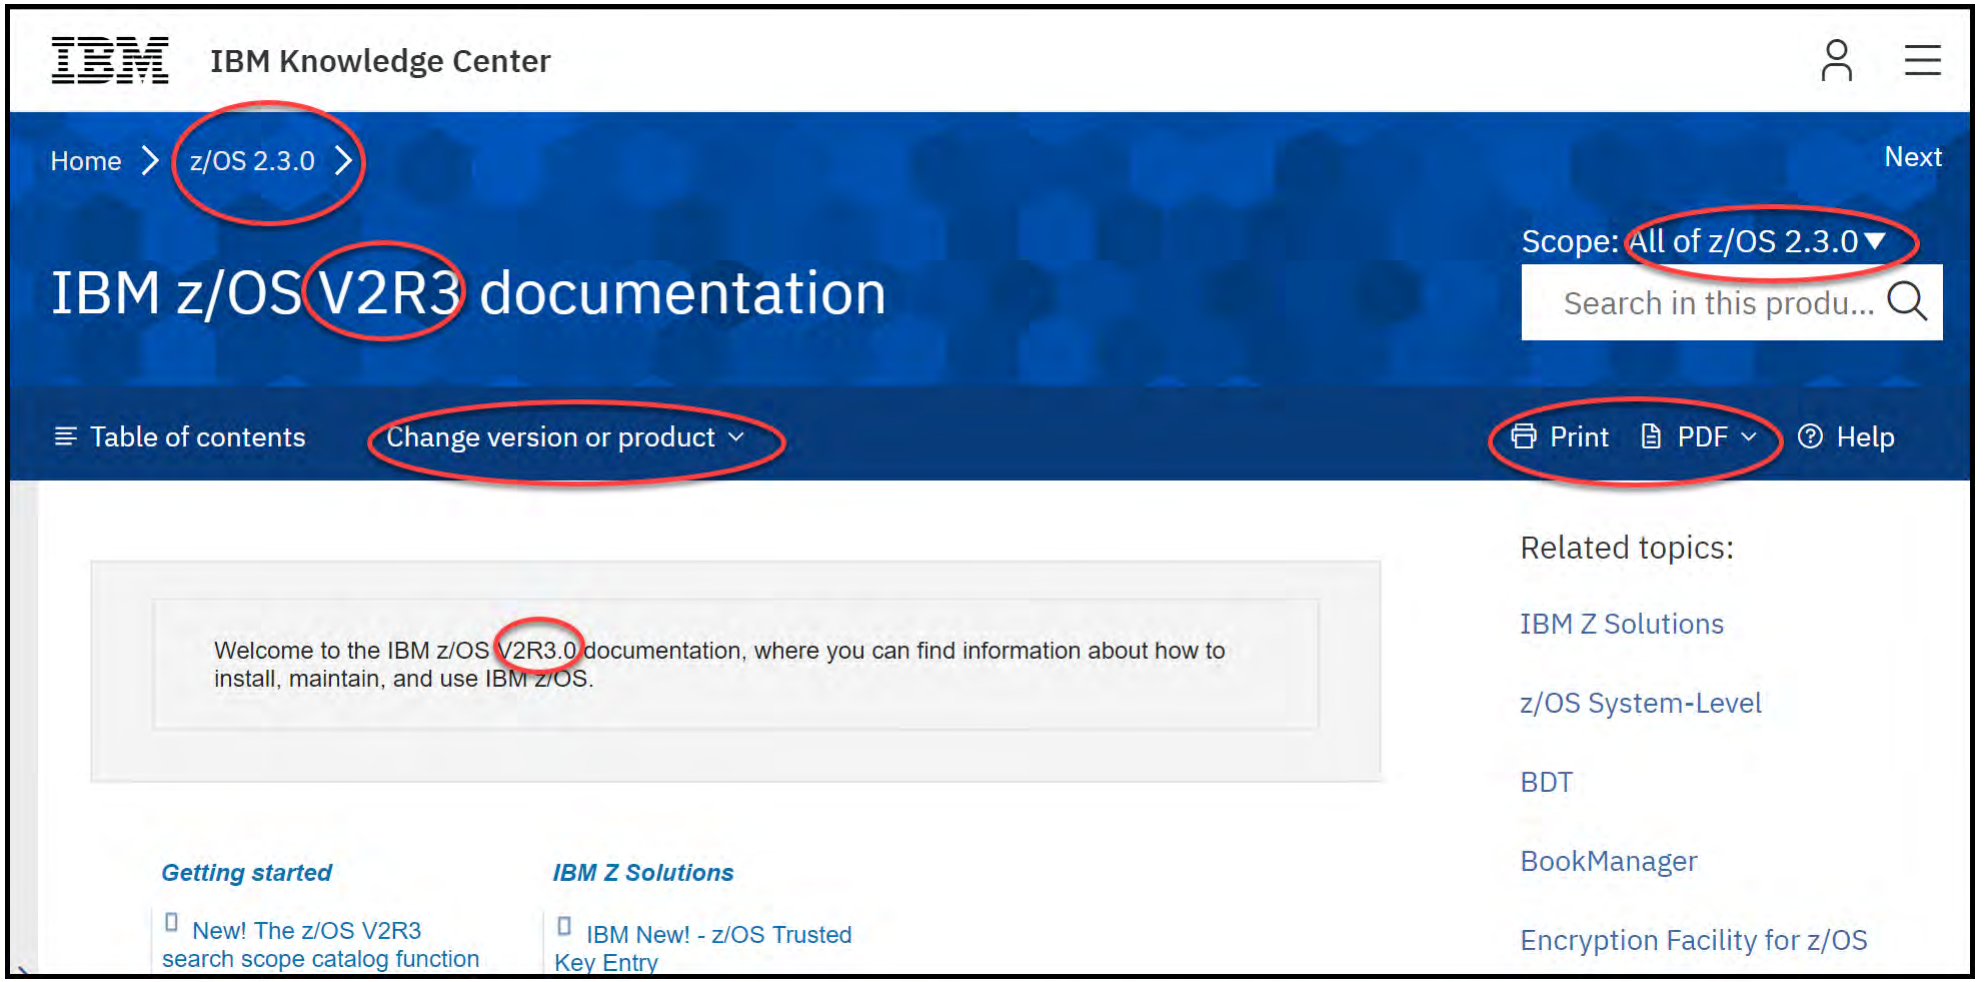

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

What's New in IBM Documentation

- Google 'z favorites' for a new and very improved zFavorites list

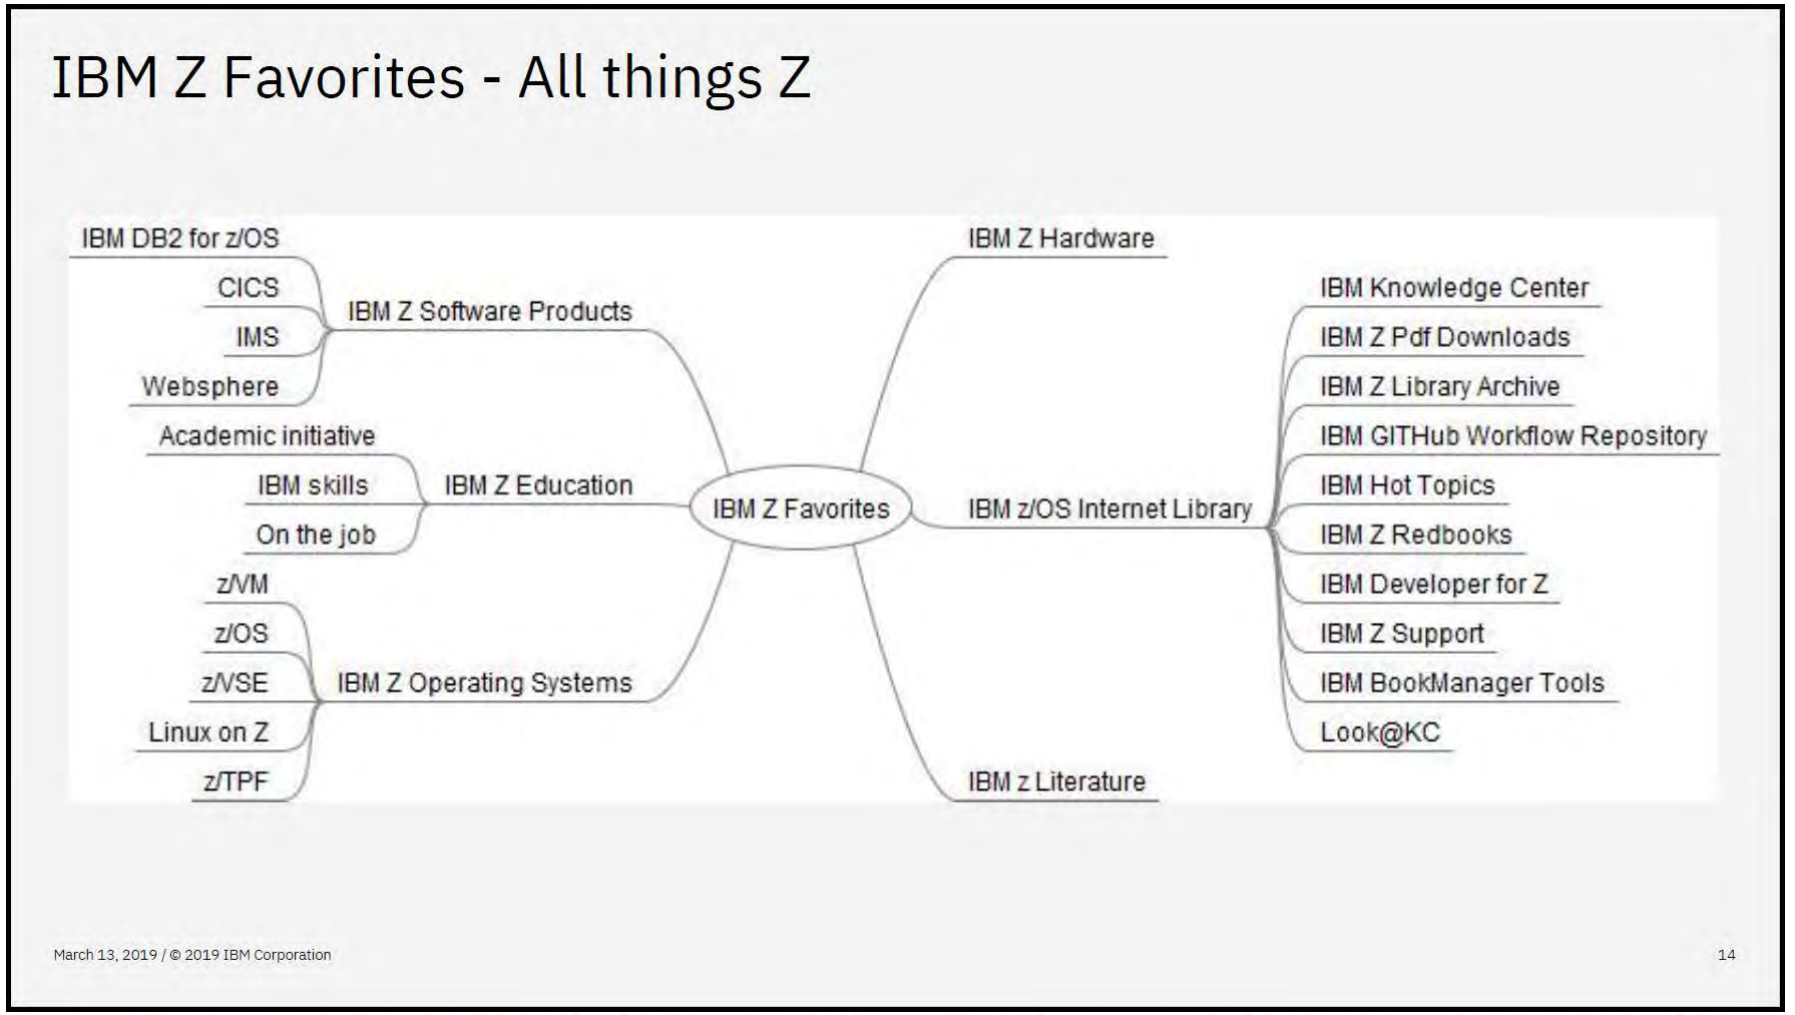

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

What's New in IBM Documentation

- IBM Z Content Solutions
  - All content for a function in a single location (documentation, workflows, videos, KC, etc.)
  - Currently found (but not identified) on first page of KC (see next slides)
- Currently working on a dedicated 'landing page' for Z Content Solutions.

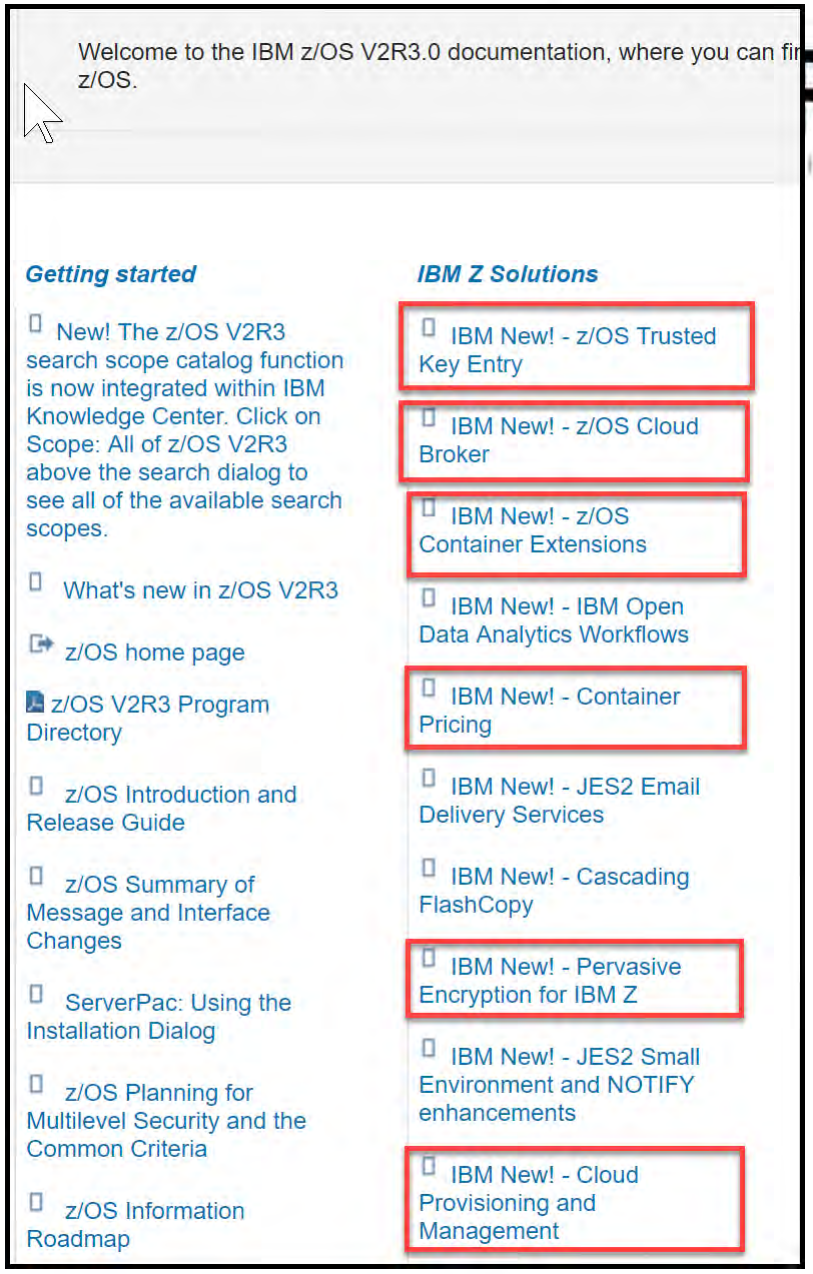

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

What's New in IBM Documentation

- IBM Z Content Solution - What is Container pricing? (Part 1)

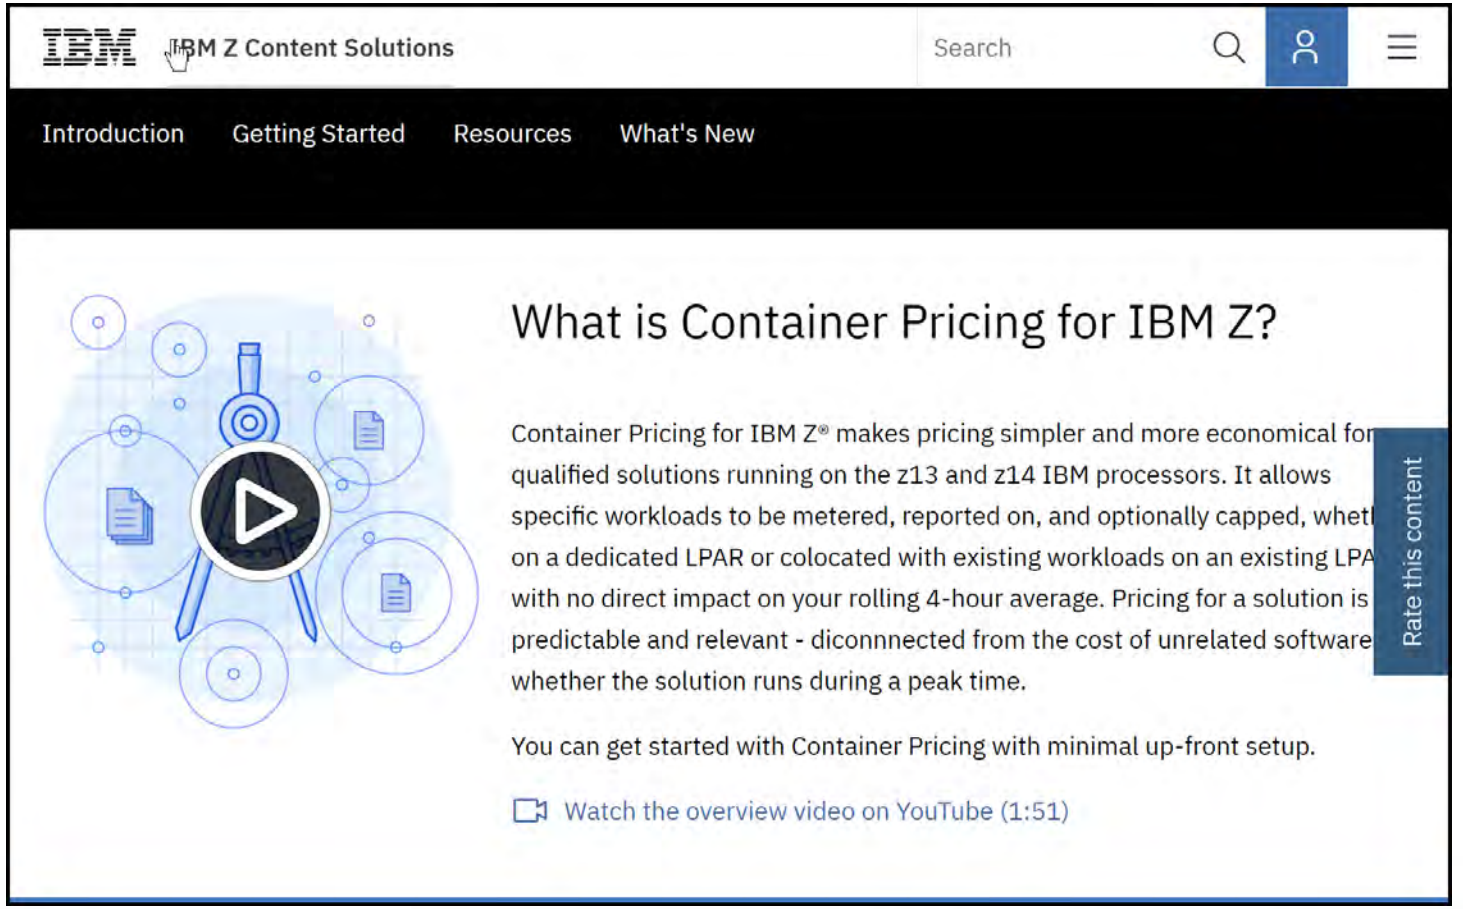

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

What's New in IBM Documentation

- IBM Z Content Solution - What is Container pricing? (Part 2)

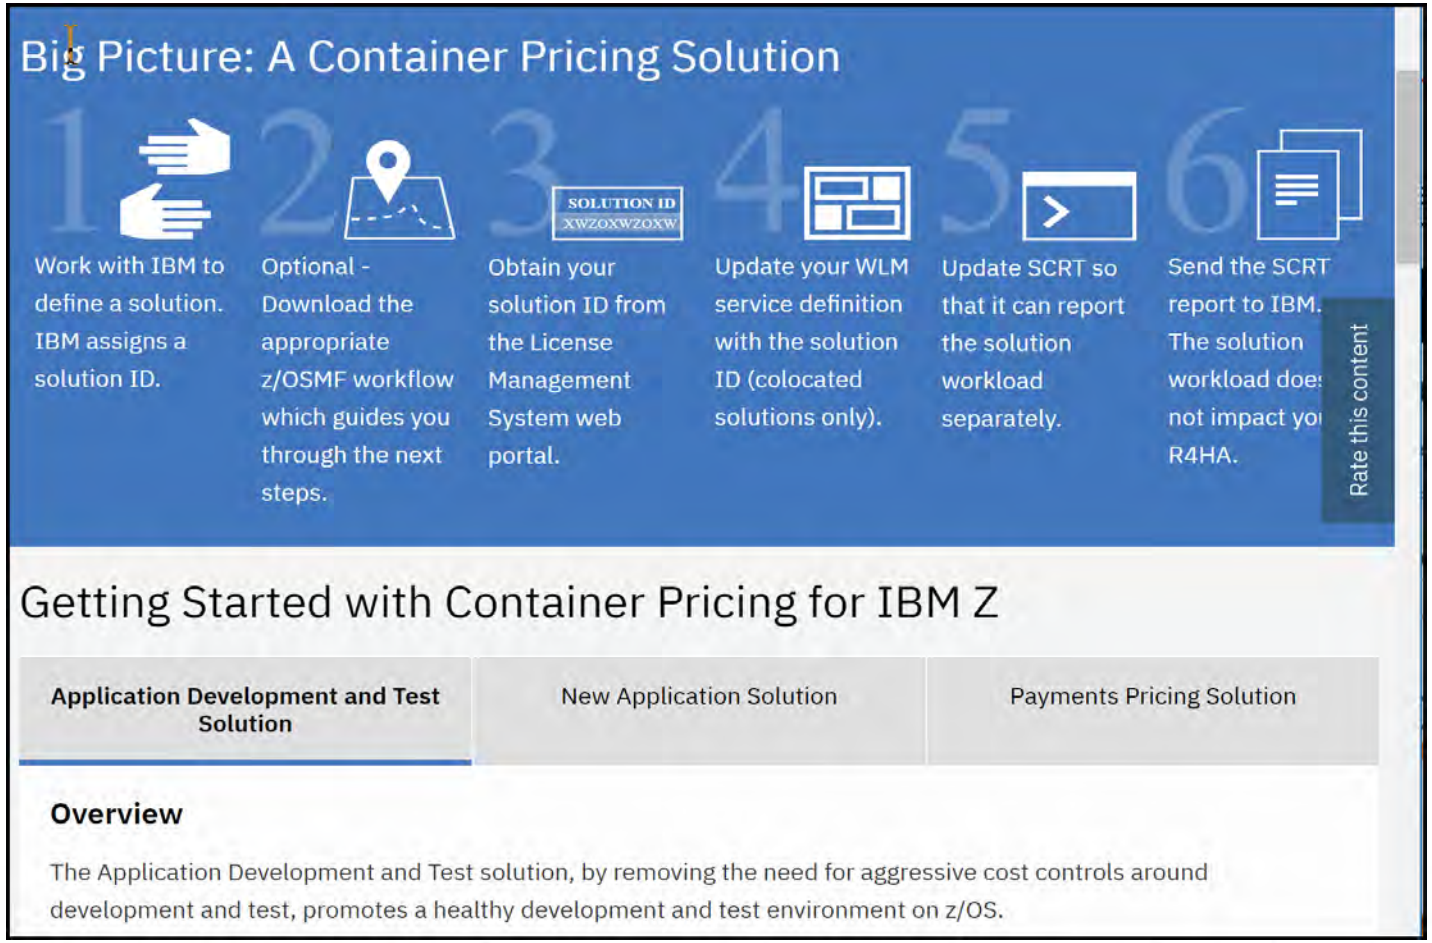

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

What's New in IBM Documentation

- IBM Z Content Solution - What is Container pricing? (Part 3)

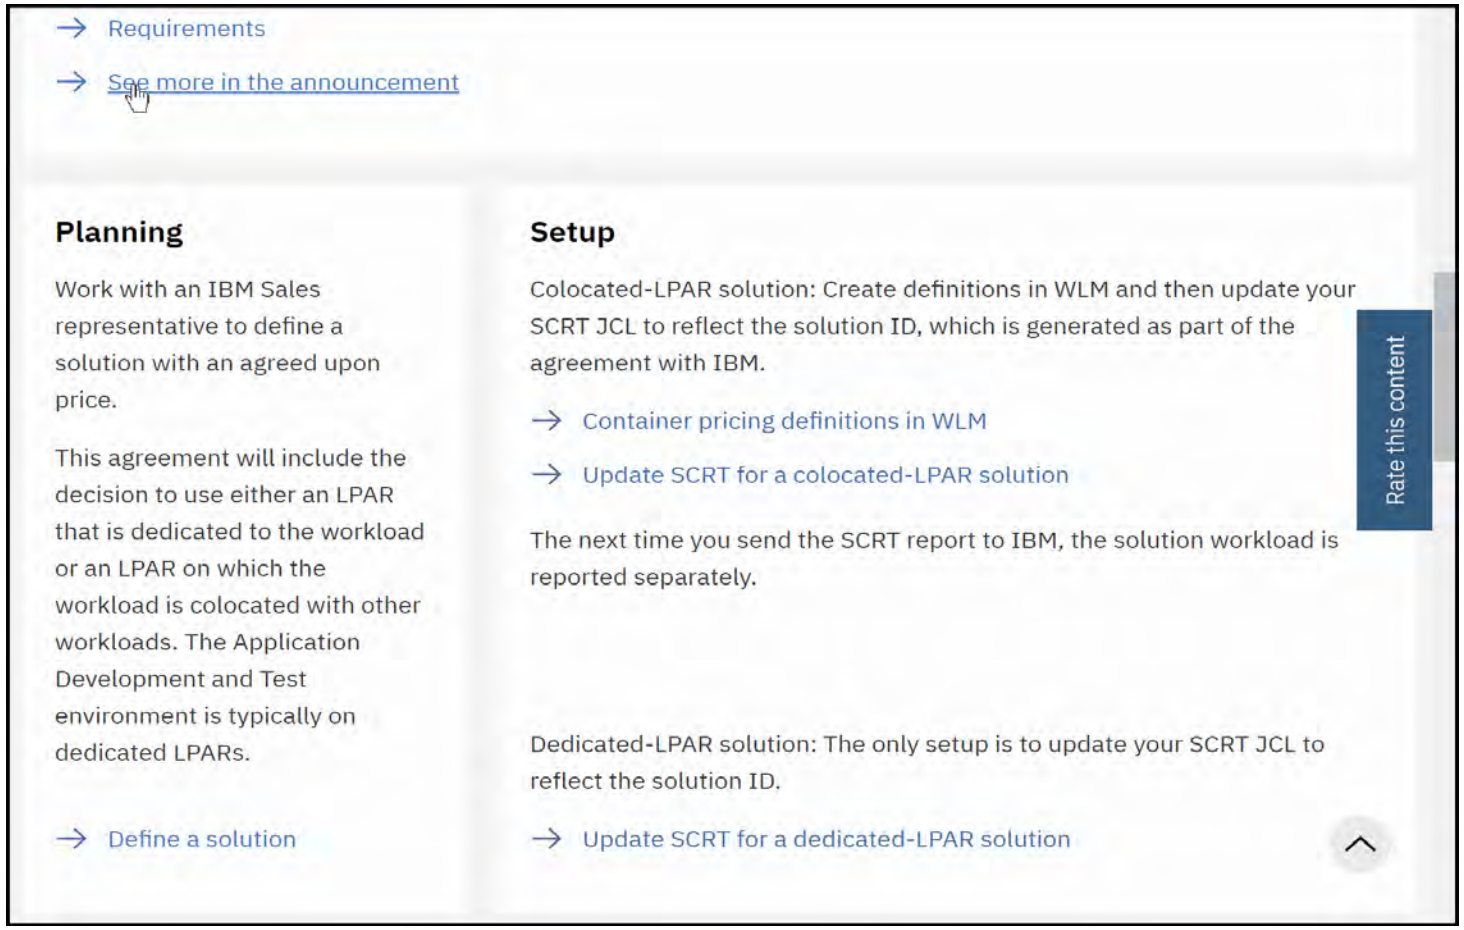

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

What's New in IBM Documentation

- IBM Z Content Solution - What is Container pricing? (…last part)

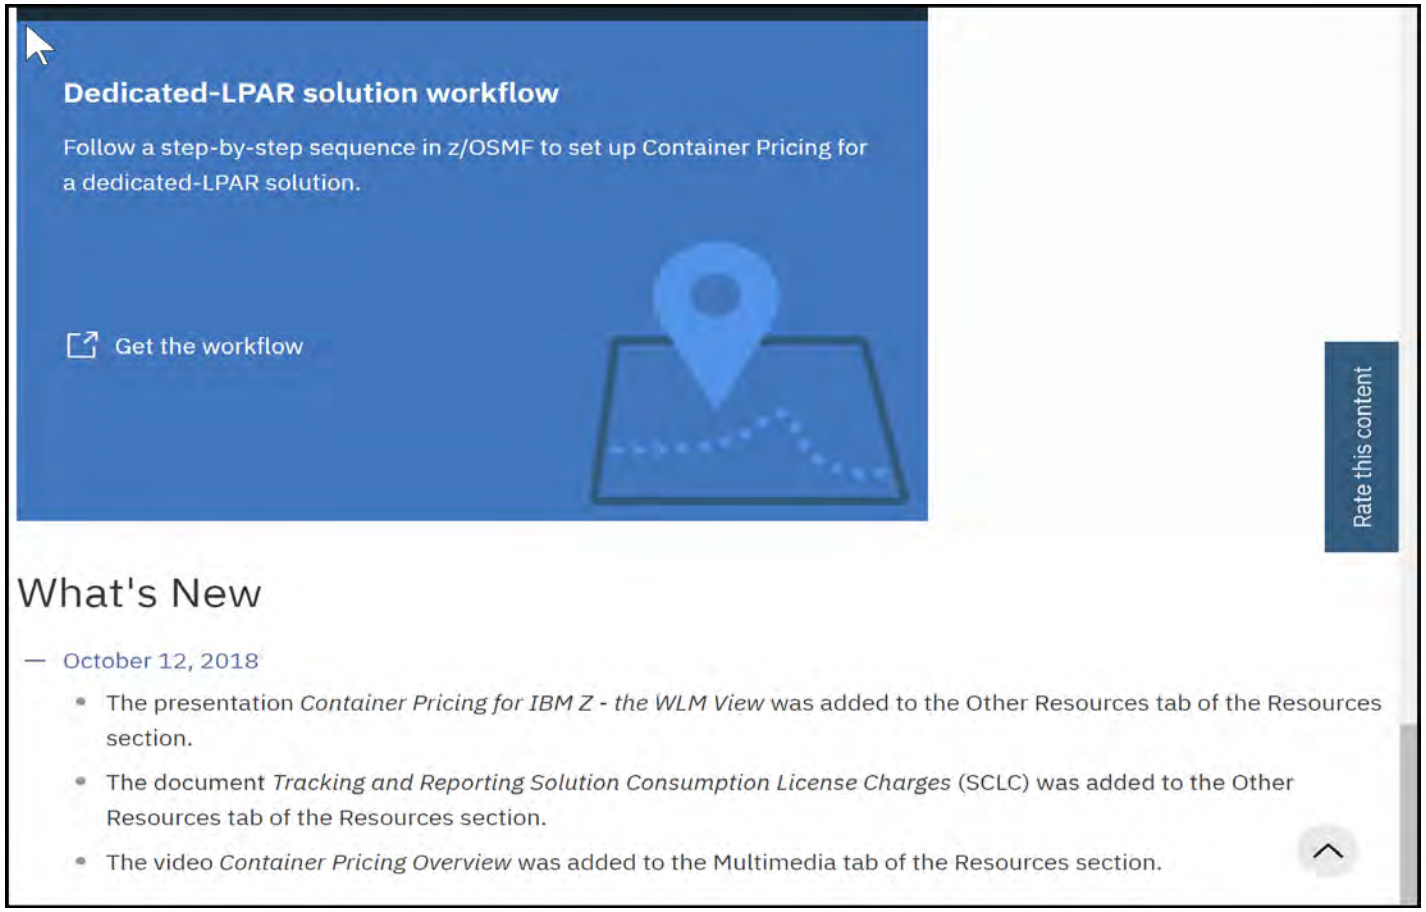

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

What's New in IBM Documentation

- SHARE Sessions
  - Systems Programmer's Information Navigator – 24647 – **Bill Keller**, **Sue Shumway**
- Please take their survey - http://www.surveygizmo.com/s3/4504455/IBM-z-OS-Product-Documentation-Survey-2018
- A few questions for you:
  - Would it be helpful to have a list of related z/OSMF Workflows on the Content Solutions landing page?
  - Can you access YouTube from work?
  - Does your company have any issue with using code from github on your production z/OS?

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

Edumacation

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

What's It Like To Be The New Kid At SHARE

- Reflections from a Newbie:
  - The atmosphere of learning and sharing is *amazing*.
  - I can't believe how many young people are here!
  - I can't believe how willing all these experts are to answer all my questions.
  - I would like to have a logical plan laid out for a person new to the platform:
    - The Intro to z/OS classes and Demystifying Mainframe Jargon would be helpful to have on the first day to set up subsequent sessions.
    - There are so many sessions, and it's not always intuitive as to which can be understood by a new person versus those that need some background. Can a set of sessions be recommended for the new person, so we don't feel so overwhelmed? And can they be sequenced so that they run from most basic to most advanced?

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

IBM Training

- Cheryl's Favorite Tip of Week
  - New website as of January - IBM Z Skills Employer Resource Center
    - Can also access through ibm.biz/zskills and select 'Sign in to use the Resource Center'
  - Item #4 when you get on that site (Train talent) provides many suggestions for no-cost z/OS training:
    - IBM Master the Mainframe Learning System
    - z/OS Introductory videos (27)
    - Introduction to the New Mainframe: z/OS Basics Redbook
    - IBM z/OS Bootcamp for Beginners (10-day instructor-led remote)
    - IBM z/OS Introduction & Workshop for the IT Professional (5-day online or remote)
    - Mainframe Playground
    - Youtube videos
    - Interskill classes

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

https://www.onlinedigitallearning.com/course/view.php?id=3532

Apps | Pricing | IBM | WW | Our Sites | SHARE | Travel | Movies | News | Misc. | Financial | Links | Sarasota | Politics

IBM Training and Skills

IBM Z Skills Employer Resource Center

- 1. IBM Z Skills Employer Resource Center Introduction
- 2. Develop a Talent Strategy
- 3. Attract & Hire talent
- 4. Train talent
- 5. Grow and Retain talent
- 6. Transition and Succession
- 7. Talent experts who can help

My Enrollments

My courses

IBM Z Skills Employer Resource Center

My Enrollments / My courses / IBM Z Skills Employer Resource Center

1. IBM Z Skills Employer Resource Center Introduction

Your progress

Enterprise computing runs on talent. This comprehensive resource center is designed to provide organizations with the tools and guidance needed to develop a solid strategy for building a robust pipeline of mainframe talent that will propel your company's enterprise infrastructures into the future. Throughout this resource center, we will share best practices to help you create a thorough, thoughtful talent strategy for attracting talent and cultivating them for a career in enterprise computing. Whether you need to find, train, or grow Z talent, we have the resources to help!

Review the guide below to understand how to navigate through this course and then begin your learning journey to discover how you can start building your pipeline of Z talent today. For a high-level summary of this resource center, access the Employer Reference Guide on the IBM Z Skills website.

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

IBM Training

- From Employer Resource Center
  - 10 day, no-charge, instructor-led online z/OS Bootcamp for those new to z/OS

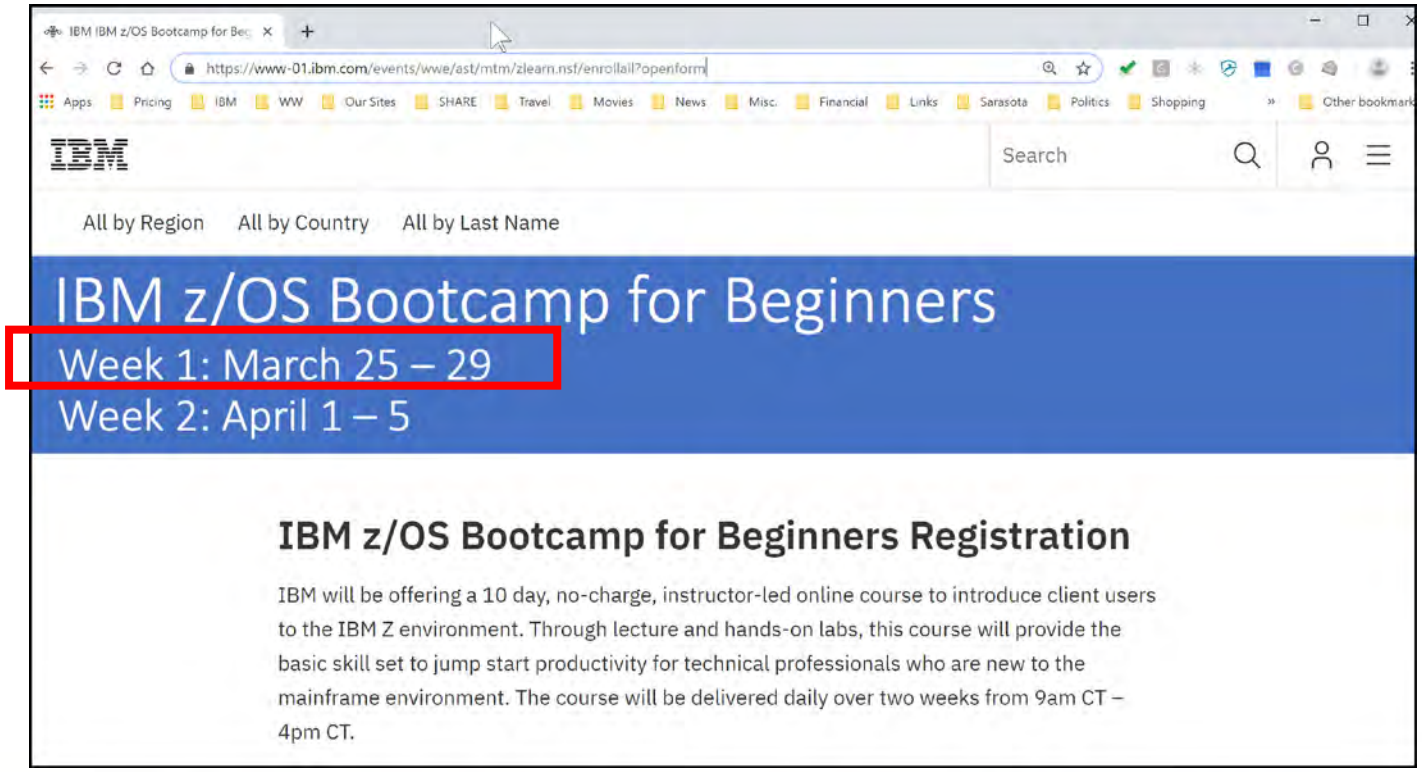

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

z/OSMF
(Have We Mentioned This Before?)

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

z/OSMF

- Watson & Walker has been promoting and supporting the use of z/OSMF since its inception because we believed in its potential
- Its adoption was slow because of its poor performance, resource consumption, and installation complexity (and, quite frankly, because old-timers could do things faster with their own tools)
- But it was really designed for new sysprogs, and IBM has continued to develop it to where it is now a requirement for all installations (and much easier for new sysprogs)
- Starting with z/OS 2.3, z/OSMF will start automatically unless you take explcity action to stop it from doing so.

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

z/OSMF

- This is NOT bleeding edge any more!
- Resource usage is a fraction of what it used to be.
- Installation and customization is much easier than it used to be.
- *Every* SHARE has a number of z/OSMF labs – go along and learn from IBM's experts. And take the opportunity to tell them what you like and what you don't like.
- z/OSMF now has functions that are not available elsewhere.
  - z/OSMF Workflows are not just nice-to-have, in z/OS 2.5, the Upgrade (was Migration) Guide will be replaced with a z/OSMF workflow.
  - Rumor has it that a post-z/OS 2.4 CICS ServerPac install will *require* z/OSMF.
- Marna and Cheryl are working on an article about z/OSMF Workflows for the next Tuning Letter.
- If you *don't* start using z/OSMF, you are going to have to answer to "Mean Cheryl".

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

z/OSMF

- Workflows are the future for IBM, ISVs, and users
- Phoenix Software seems to be ahead of the pack
  - See their webinar on how to install a product using an ISV-created workflow
  - https://www.phoenixsoftware.com/webinars by the one and only **Ed Jaffe**! It's excellent and shouldn't be missed!
- z/OSMF sessions:
  - 24119 – What's New in z/OSMF V2.3? – **Joey Zhu**
    - The new Console app and new desktop views look very neat
  - 24120 – z/OSMF V2.3 Implementation and Configuration – **Joey Zhu**
  - 24565 – z/OSMF User Experiences: Workflow Editor and Workflows – **Ed Webb**
  - 24460 – Ready Set – z/OSMF Security Setup The Basics of Getting z/OSMF Configured Securely – **Richard Faulhaber**
  - 24661 – Accelerate Provisioning of z/OS Middleware – **Hiren Shah**

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

z/OSMF

- z/OSMF sessions (cont.)
  - 24155/24154 – Pervasive Encryption – IPsec AT-TLS Using Configuration Assistant for z/OS V2R3 – **Linda Harrison**
  - 24152/24151 – TCP/IP Stack Configuration with Configuration Assistant for z/OS V2R3 CS – **Mike Fox**
  - 24123 – Installing a Product Using z/OSMF Software Management – **Kurt Quackenbush** (this was a particularly well-attended session)
  - 24121 – z/OSMF Hints and Tips, and Common Problems – **Kenneth Irwin**

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

Thank you!

- If you have any questions, suggestions, comments, or general abuse, please email us at technical@watsonwalker.com
- Have a safe trip home, thank you for coming, and we hope to see you in Pittsburgh in August.

- Please complete the online evaluation!

Thanks!!!

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

ADDITIONAL MATERIAL

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

Container Pricing

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

Container Thoughts (no, not THOSE Containers)

- This week, we heard more about the Dev/Test container and a new new (not a stutter) application container called SCLC.
- Last October, IBM announced a refinement to the Dev/Test container, providing a sizing tool, clarification as to what is a qualifying workload, and the availability of Limited Use Licenses for IPLA products “at an attractive price”.

  Our thoughts:

  - For many customers, this model is a very attractive offering, so be sure to check it out!
  - Customers can get significantly more space to develop or test on the platform at a very low cost (per IBM, no additional **MLC** costs).
  - Be sure to check the real costs before signing up. There may be costs associated with IBM IPLA products, additional hardware, maintenance, and ISV software.
  - Be aware of how your workload runs. If current dev/test work does not drive the R4HA peak, you may already be getting this capacity for “free”.
  - Traditional capping can still be used to control both production and dev/test costs, but caps should be revised to reflect the change in pricing models.

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

Container Thoughts (no, not THOSE Containers)

- On Wednesday, IBM also clarified that a Dev/Test container can be used even when the workload is co-located in a single LPAR with production workload.

  Our thoughts:

  - The process to set this up will take some time, so be sure to plan at least 6 months prior to when you would like to sign up for the model.
  - Capping may not be as effective under this model. If you increase the LPAR's softcap so that your DevTest workload has more capacity, there is no effective way to stop the production work from using some or all of the additional capacity. This would not affect the MLC for your DevTest work, but it *could* increase the MLC for your production work.
  - Also, all CPU consumption by system tasks that are shared between production and test (CATALOG, HSM, GRS, JES2, automation, and so on) are charged to production.
  - The net is that while it is possible to use the DevTest container offering for mixed Production/Development workloads, it is much easier and more effective to use with dedicated production LPARs.

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/

Container Thoughts (no, not THOSE Containers)

- IBM also announced Solution Consumption License Charges touted as a Pay As You Go Model for workloads new to the platform. IBM's take is that this simplifies software pricing, because customers will be able to match increase/decreases in charges directly to increases/decreases in capacity consumption.

  Our thoughts:

  - SCLC will likely be priced attractively when adding an application to the platform.
  - Customers have asked to pay for exactly what they used, so this is IBM's response.
  - Customers have been tuning to the R4HA for almost 20 years. This model requires a shift in the way we've been thinking for all that time.
  - This new model charges for every MSU, so there are no longer "free" MSUs during hours outside of the R4HA peak.
  - Customers will lose control over software costs. How would capping apply?
  - Customers will still need to manage their R4HA peaks for several reasons:
    - Used for IBM MLC for existing workloads
    - Used for IBM IPLA for both new and existing workloads
    - Often used for ISV software

Copyright© by SHARE Association Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. http://creativecommons.org/licenses/by-nc-nd/3.0/