



# Watson & Walker zRoadshow

SHARE in Dallas 2022

Frank Kyne (Frank@watsonwalker.com)



# Introduction



- Who are we?
- Watson & Walker founded in 1988 by Cheryl Watson & Tom Walker.
- Publisher of *Cheryl Watson's Tuning Letter* and Cheryl's CPU Charts since 1991.
- We are completely independent, not beholden to any vendor.
- Our focus areas at the moment are convincing upper management that the zSkills issue is *not* fake news, and understanding SMF/AI/ML.
- For more information, see our website, [www.watsonwalker.com](http://www.watsonwalker.com).
- The objective of this session is to make you aware of interesting items and point you to where you can get more information.
- Please remember to do the session evaluation for your *last* session.

# Agenda

- Latest news on zSort Accelerator exploitation and customer experiences
- LPAR Design tool
- What's happening with SMF data access
- Coupling Facility Monopolization Avoidance
- z/OSMF items
- WLM RTCapLeadTime thoughts
- JES2 Spool Compression and zEDC
- Page Data Sets (or not)
- SHARE Tips

# Latest News on Sort Accelerator



# Latest News on Sort Accelerator

The Sort Accelerator was included as part of the z15 announcement in 2019.

- It is designed to reduce CPU and elapsed time for large sorts.


Because it is up to each product that wants to exploit it to do their own thing, it is not easy to get a *comprehensive* picture of where it is or isn't supported.

- Products that exploit the accelerator are expected to handle the long list of restrictions when deciding whether or not to use the accelerator for a given sort.
  - See page 14 of [this document](#) for a list of the restrictions.
  - *Ideally*, this should all be transparent to the user, except that sorts that are a good fit for the accelerator should consume less CPU time.
    - However, different exploiters have different default values regarding zSort enablement, so explicit actions are required in some cases.

The following is the latest information we have:

# Latest News on Sort Accelerator Experiences

- DFSORT support is available on z/OS 2.3 or later with the PTF for APAR [PH03207](#).
- **DFSORT use of zSort is turned OFF by default.**
  - You can enable/disable zSort at the *system* level, and override at the *individual sort* level.
- There are currently no open zSort-related DFSORT HIPER APARs, but there *are* a number of PE APARs.
- *One* customer's experience:



# Latest News on Sort Accelerator Experiences

- To gather zSort statistics in the DFSORT SMF records, make sure that DFSORT is writing 'TYPE=FULL' SMF records (and ignore the obsolete warning in the DFSORT manual about the overhead of using TYPE=FULL).
- There is a flag in the basic records (ICEZSRNU) to indicate if zSort was used or not.
  - If the field contains x'00', zSort was used, and more information is provided in the zSort section.
  - If the field contains a value *other than* x'00', that indicates *why* zSort was not used. For a list of the meanings of the other codes, refer to <https://www.ibm.com/docs/en/zos/2.5.0?topic=messages-ice267i>. (the next release of the DFSORT manual will add this information to the SMF field description).
    - This is how zBNA is able to show which jobs would be valid candidates for zSort.

# Latest News on Sort Accelerator Experiences

- Precisely (Syncsort) support was delivered for MFX R3.1 by APAR TZ04284.
  - Like DFSORT, Syncsort provides the ability to enable/disable zSort at the system level, and override at the individual sort level.
  - However, **Syncsort use of zSort is turned ON by default.**
- No customer feedback so far.
- We believe that the Syncsort SMF record *does* indicate whether zSort was used or not.
  - However, the SMF record doesn't tell you *why* zSort was not used.
  - Reviewing the output messages from thousands of sorts is not a viable option.



# Latest News on Sort Accelerator Experiences

- Db2 can support the Sort Accelerator in two places:
  1. Internal ('RDS') sorts ('ORDER BY' or 'GROUP BY'). This support was delivered by Db2 12 APAR [PH31684](#).
  2. Db2 utilities.
- Db2's decisions about whether or not to use zSort for RDS sorts are based on internal Db2 performance measurements.
  - You have no control over whether zSort is used by RDS sorts (other than an 'undocumented' subsystem-level ON/OFF switch). **The default is ON.**
- According to IDUG presentation [Db2 for z/OS Performance Improvement Opportunities](#) by **Akiko Hoshikawa**, average CPU savings for RDS sorts was 8%.
- What metrics are available to help you analyze *your* Db2 RDS use of zSort?

# Latest News on Sort Accelerator Experiences

- Sort of (sorry) documented in IFCID 95 (Start of sort)/96 ('End' of sort). Db2 APAR [PH31684](#) added information about whether each sort used zSort or not. However:
  - Even though they are not very big, you get 1 for *each* RDS sort – that could be *hundreds of millions* per day, so not feasible to enable on a production system.
  - The 'RECNO' field sometimes shows sorts with 0 (!) or 1 record being passed to zSort. This *might* be because there are cases where the number of records to be sorted is not known until AFTER the IFC96 record (you know, the END OF SORT one) is written.
    - And, there is no flag in the record to tell you if this is one of those sorts, or if it really did try to sort zero records.

# Latest News on Sort Accelerator Experiences

- Extended counter 255 in the SMF Type 113 records reports the number of 'SORTL' calls.
  - We have examples where the only user of zSort is Db2 RDS, on a system running only online transactions.
  - You (I?) would *not* expect *online* transactions to be retrieving huge numbers of rows from Db2.
  - Yet we see intervals where the *average* number of SORTL calls per RDS sort is  $> 1.8$ . Strange. And a potential application tuning opportunity if online transactions are retrieving too much from Db2.
- There is 'room for improvement' in the Db2 instrumentation.
- Having said all that, one of our large customers determined that enabling zSort decreased *overall* Db2 CPU per transaction by 2.5% using the default Db2 SRTPOOL value (roughly 20% of all their Db2 transactions performed a sort, so this infers a saving of about 12.5% for the txns that used zSort).
  - IBM guideline is that increasing the SRTPOOL value can increase the number of RDS sorts that use the Sort Accelerator.

# Latest News on Sort Accelerator Experiences

- The other place Db2 can use zSort is in Db2 utilities:
  - If you have IBM Db2 Sort for z/OS V2, apply the PTF for APAR [PH32374](#).
    - Also, see APARs [PH35059](#) and [PH38747](#).
    - **The default for Db2 Sort for z/OS V2 is that zSort is ENABLED.**
  - If you do NOT have IBM Db2 Sort for z/OS V2, Db2 uses DFSORT for utilities. In that case, support for zSort was delivered by Db2 APAR [PH28183](#).
    - That APAR adds new ZPARM - UTILS\_USE\_ZSORT.
      - Specifying YES tells DFSORT to use zSort.
      - **The default is NO, indicating that zSort should *not* be used.**
      - The value on the UTILS\_USE\_ZSORT ZPARM overrides the system level DFSORT zSort setting – similar to providing an override via JCL at the individual job step level.
      - Note that, with the above APARs installed, the UTILS\_USE\_ZSORT ZPARM has no effect if you are using IBM Db2 Sort for z/OS V2.

# Latest News on Sort Accelerator

If you are interested in enabling Sort Accelerator for selected jobs, use the **FREE** IBM zBNA tool to help you identify good candidates. For more info about zSort see:

- SHARE Summer 2021, [\*Making the Most of zSort\*](#), by **Ryan Bouchard** and **Samuel Smith**.
- SHARE Summer 2021, [\*Z Sort on z15 + Large Memory = Big performance!\*](#) by **Dave Betten**.
- DFSORT Documentation for enabling APAR PH03207: [\*DFSORT User Guide IBM Integrated Accelerator for Z Sort\*](#).
- [\*z/OS Sort Accelerator\*](#) article in *Tuning Letter 2020 No. 3*.
- Precisely knowledgebase article about zSort: [\*IBM z15 Integrated Accelerator and Syncsort MFX \(precisely.com\)\*](#)

**Make sure to monitor the IBM.Function.Zsort FIXCAT.**

# LPAR Design Tool



Some of you might find yourself installing a new CPC before the next SHARE conference.

- IBM sent out a 'Save the Date' email for April 5<sup>th</sup>.

The new z16 uses IBM's new Telum chip that was previewed at the 2021 'Hot chips' conference.

- One aspect of the Telum chip is a complete redesign of the Level 2/3/4 caches.

Because of that cache redesign, it is *especially* important that you are collecting the CPU MF data (SMF type 113 records) for all your LPARs:

- The cost is so low that you can barely measure it, and compared to major subsystems, the volume of data it collects is *tiny*.
- And yet we **still** encounter sites that have not enabled this invaluable resource.

It is well known that some customers have achieved *much* better performance than predicted when they moved to a CPC with more, but slower, cores – for example, z13 7xx to z14 5xx, z13 7xx to z15 6xx, and so on.

*Part* of the reason for the better-than-expected performance was that by having more CPs, you get more cache than a higher-speed model with the same number of MSUs.

On the Telum, each core has 32 MB of L2 cache vs 8MB in z15.

Additionally, unused space in a Telum L2 cache can potentially be used as a virtual L3 or L4 cache by another core.

- This effectively means that the sizes of the L3 and L4 caches are constantly variable, and not fixed as they are on previous generations.



How will this affect performance?

- We've no idea yet, but the chip industry is raving about this as the model for future chips.

The important thing is that if you are upgrading to a z16, you are able to provide IBM with as much information (SMF type 7x and 113 records) as possible about the cache usage of your systems, to allow IBM to more accurately predict how your workloads will interact with the new chip.

Then, after you take delivery of your shiny new z16, you should ensure that your LPAR definitions are optimized for the new CPC. And that is where the (**FREE**) LPAR Design tool comes into the picture.

‘LPAR Design’ is a spreadsheet-based tool created by **Alain Maneville** when he worked in IBM, together with his customer, **Thierry Deleris**.

The tool allows you to model how your LPAR definitions and weights interact with HiperDispatch.

It also offers ‘expert suggestions’, where it will analyze your configuration and make suggestions about fine tuning LPAR weights to optimize the number of Vertical High CPs.

You can define your configuration manually (including soft caps and group caps), using tabs on the spreadsheet, or you can input from a file created by zPCR.

- You can even fine-tune your target configuration and then send the updated (hopefully more efficient) configuration back to zPCR.
  - Remember that LPAR weights generally do NOT act as hard caps on an LPAR’s capacity, so fine adjustments should only have a capacity impact if the CPC is running at 100% busy, but they CAN improve the HiperDispatch topology.

# LPAR Design Tool – Config tab

Remember to click this if you change the config

ID= - LPARDesign-HD-V9-T02		Current zPC	
CFG-LP-VALID?	YES	#PhyProc	36
Machine-type	2964-736	#LPs (non-ICF, non-DED)	254
MSU	4,452	Ratio LP/PP (base)	7.06
Total Weight	421	LSPR-AVG-V2R2-MI	37,972
Max LPAR	85	#LPARs	23

LPARNAME	WEIGHT	#LP	%SHARE (By Pool)	RESERVED	Guaranteed #PP	Wkld LSPR	MinReq#LP	Check#LP	HD-HIGH#	HD-MED#	HD-MED%	HD-LOW#	Active LP	Report LP
WAW0B1	173	30	41.1%		14.79	Average	15	#VL>2	14	1	79.3%	15	15	15
WAW0T0	1	2	0.2%		0.09	Average	1	OK	0	1	8.6%	1	2	1
WAW0T1	4	2	1.0%		0.34	Average	1	OK	0	1	34.2%	1	2	1
WAW0X1	11	10	2.6%		0.94	Average	1	#VL>2	0	1	94.1%	9	2	1
WAW0X5	8	5	1.9%		0.68	Average	1	#VL>2	0	1	68.4%	4	2	1
WAW0Y1	10	10	2.4%		0.86	Average	1	#VL>2	0	1	85.5%	9	2	1
WAW0Y5	3	5	0.7%		0.26	Average	1	#VL>2	0	1	25.7%	4	2	1
WAW0Z0	2	2	0.5%		0.17	Average	1	OK	0	1	17.1%	1	2	1
WAW0Z1	44	20	10.5%		3.76	Average	4	#VL>2	3	1	76.2%	16	4	4
WAW0Z5	10	5	2.4%		0.86	Average	1	#VL>2	0	1	85.5%	4	2	1
WAW0B0	12	4	2.9%		1.03	Average	2	OK	0	2	51.3%	2	2	2
WAW0B3	11	25	2.6%		0.94	Average	1	#VL>2	0	1	94.1%	24	2	1
WAW0B5	36	25	8.6%		3.08	Average	4	#VL>2	2	2	53.9%	21	4	4
WAW0B7	14	30	3.3%		1.20	Average	2	#VL>2	0	2	59.9%	28	2	2
WAW0D0	4	2	1.0%		0.34	Average	1	OK	0	1	34.2%	1	2	1
WAW0D1	11	10	2.6%		0.94	Average	1	#VL>2	0	1	94.1%	9	2	1
WAW0Q1	10	10	2.4%		0.86	Average	1	#VL>2	0	1	85.5%	9	2	1
WAW0Q5	4	5	1.0%		0.34	Average	1	#VL>2	0	1	34.2%	4	2	1
WAW0R0	2	2	0.5%		0.17	Average	1	OK	0	1	17.1%	1	2	1
WAW0R1	29	20	6.9%		2.48	Average	3	#VL>2	1	2	74.0%	17	3	3
WAW0R3	8	5	1.9%		0.68	Average	1	#VL>2	0	1	68.4%	4	2	1
WAW0R5	7	10	1.7%		0.60	Average	1	#VL>2	0	1	59.9%	9	2	1
WAW0R7	7	15	1.7%		0.60	Average	1	#VL>2	0	1	59.9%	14	2	1

# LPAR Design Tool – Expert Tab

Separate recommendations for each processor pool

Separate recommendations for each LPAR

ID=Watson Walker M - LPARDesign-HD-V7-T01 Current zPCR Version-9.1a EXPERT	
<a href="#">Click for EXPERT NOTES - GCI</a> <a href="#">Click for EXPERT NOTES - zIIP</a> <a href="#">Click for EXPERT NOTES - IFL</a>	
LPAR	Suggested Improvement Notes - GCP - Machine Type = 2964
<b>WAWOB1</b>	<p>(R8-GCP) - *WARNING* - The number of VL (13) is above the IBM Best Practice See: <a href="http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/TD106388">http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/TD106388</a> - for this Best Practice</p> <p>[+] NOTE - You have 2VM and 13VL When ALL VIs will be UnParked, each VM and VL will have an entitlement of 3.57%</p>
<b>WAWOT0</b>	<p>(R5-GCP) - Due to the "at least 2 LP rule", the first VL will be always Unparked Your current Guaranteed#PP is 0.08 - you will have 2-VM with an entitlement of 4% and so 2 Active LP</p> <p>[+] NOTE - You have 1VM and 1VL When ALL VIs will be UnParked, each VM and VL will have an entitlement of 4%</p>
<b>WAWOT1</b>	<p>(R5-GCP) - Due to the "at least 2 LP rule", the first VL will be always Unparked Your current Guaranteed#PP is 0.32 - you will have 2-VM with an entitlement of 16.05% and so 2 Active LP</p> <p>[+] NOTE - You have 1VM and 1VL When ALL VIs will be UnParked, each VM and VL will have an entitlement of 16.05%</p>
<b>WAWOX1</b>	<p>(R3-GCP) - You have 1-VM with an entitlement of 88%" Your current Guaranteed#PP is 0.88 - raising the Weight and removing 9-LP would give you 1-VH but with less flexibility" The New Weight should be : 12 - The current Weight is : 11 - So you must decrease another Lpar Weight by : 1 to keep Total Weight(448) constant</p> <p>(R8-GCP) - *WARNING* - The number of VL (9) is above the IBM Best Practice See: <a href="http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/TD106388">http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/TD106388</a> - for this Best Practice</p> <p>[+] NOTE - You have 1VM and 9VL When ALL VIs will be UnParked, each VM and VL will have an entitlement of 8.84%</p>

The tool also automatically checks for common ‘errors’ such as an LPAR weight that is unattainable because the LPAR is not defined with sufficient logical CPs, or if the LPAR has >2 Vertical Low CPs.

- While fine tuning your LPAR configuration, don’t forget the Db2 recommendation that every production Db2 LPAR should have at least one Vertical High zIIP.

And if you attended **Jim Horne’s** [\*A Different "Weigh" to Define LPAR Weights\*](#) session yesterday, you can use LPAR Design to try out Jim’s suggestions on your ‘weigh’ home ;).

**We *strongly* recommend** that all our clients use the LPAR Design tools, not only when upgrading CPCs, but also *any* time they are making a change to their LPAR topology.

# LPAR Design Tool

LPAR Design will be updated with z16 LSPR numbers when they are published by IBM.

For more information about the Telum chip, refer to:

- [Short article and video](#) by **Christian Jacobi**.
- **Christian Jacobi** [Telum Introduction](#) at Hot Chips conference, August 2021.
- Excellent Anandtech analysis article '[Did IBM Just Preview The Future of Caches?](#)' by **Dr. Ian Cutress**, with accompanying [YouTube video](#).

For more information about IBM's zPCR tool, see [Getting Started with zPCR \(IBM's Processor Capacity Reference\)](#).

**Kathy Walsh's** excellent TechDoc, [Number of Logical CPs Defined for an LPAR](#)

To download the LPAR Design and LPAR Design Extended tools, including the User Guides, go to <https://github.com/AlainManeville/z-OS-LPARDesign>. For more information about using the tool, refer to [LPAR Design Tool](#) article in *Tuning Letter 2017 No 4*.

# z16 Announcement

We're sure you've seen this already, but just in case....

z16 announcement will be on Tuesday, April 5<sup>th</sup>.

Announcement event will be live streamed.

Everything you need is here: <https://www.ibm.com/community/z/ibm-z16-day-se-2022/>

- Register for the event.
- List of sessions and speakers (both IBM and customers) and times.
- 3 parallel tracks.
- Starts at 10:00 Eastern, last session *starts* at 6pm Eastern.

# What's Happening with SMF Data Access?





# What's happening with SMF?

We all LOVE the ocean of information that lives in SMF records.

- SMF provides insights and control that other platforms can only dream about.

But SMF record formats are not exactly user-friendly.

And the pool of expertise for processing SMF with traditional tools is not growing.

ONE of the strategies for surviving with fewer z/OS skills in the coming years is to *supplement* today's experts with Artificial Intelligent (AI) and Machine Learning (ML).

But the programming languages that are used for AI and ML are not that great at reading Assembler mapping macros.

We need help.

# What's happening with SMF?

We had two presentations from IBM this week about *future* tools that will read SMF data files and make the information accessible to modern languages such as Python:

- [Take a Rest: The Brand New RESTful-API of the z/OS Data Gatherer!](#), by **Yevgen Karpenko**
- [IBM SMF Explorer - The Modern Way to Work with SMF Records](#), by **Dorian Czichotzki**

Both of these are Statements of Direction for the future, not available yet. Both hope to be included as standard (not chargeable) features in z/OS.

SMF Explorer uses the Data Gatherer API to access the data (from SMF data sets now, hopefully from SMF log streams in the future), and is based on Python and Jupyter Notebooks.

Currently they support only the RMF SMF records, but hope to expand to include other record types in the future.

Recorded session  
– watch before it goes away!

# What's happening with SMF?

Both of these IBM tools are intended to make SMF data accessible to 'AI' and 'ML' programs. It is up to YOU to create those programs.

- You know SMF data, but you don't know AI.
- Data scientists know AI, but they don't know SMF.
- In *our* view, in the future, neither can deliver value without the other.



In relation to SMF and Machine Learning, *the* best session this week for me was [\*Machine Learning Applied to SMF Data to Keep You Ahead of Critical Issues\*](#) by **Machhindra Nale** and **Mikulas Cserge**.

This is a recorded session, make sure you watch it!



# What's happening with SMF?

Other SMF-related items...

Especially for those that are new to SMF, session [LNL: Insights You Can Gain from Integrated Visibility Across All Types of RMF/SMF Data](#) by **Todd Havekost** does an excellent job of illustrating the information and insights that are lurking in your SMF data.

IBM Data Gather in z/OS 2.5 has a new GRBSMFR service.

- Retrieves entire (not parsed) SMF records from sequential data sets or SMF logstream.
- Supports all record types (not limited to Type 70-79).
- For *traditional* programming languages.
- For more info, see [What's New in RMF Data Gatherer](#) by **Peter Mailand**.

ZEBRA is an Open Mainframe (open source) incubator project. Currently feeds off XML data created by RMF – doesn't read SMF data directly.

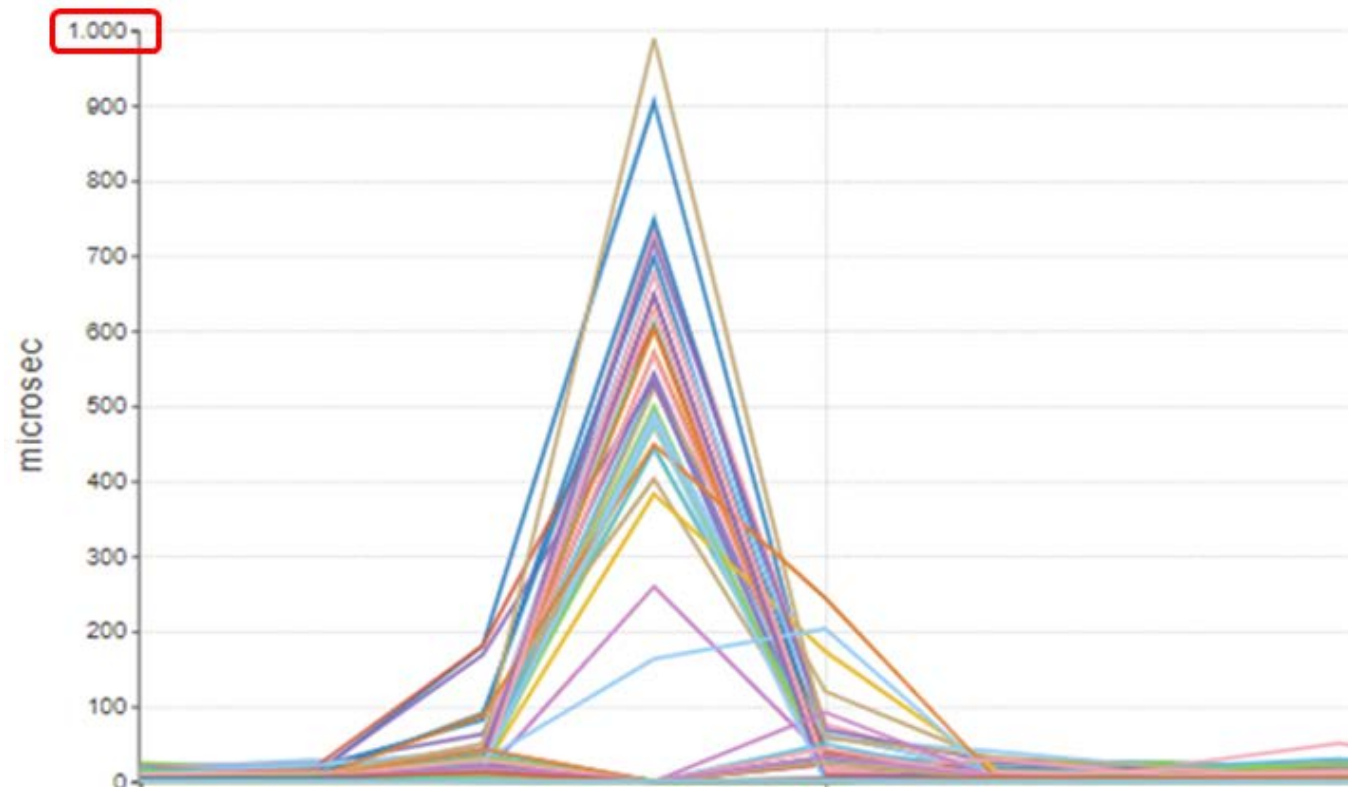
- For more info, see [How to Talk to Your Mainframe About Performance Metrics Utilizing Open Source Software Tools](#).

# Coupling Facility Monopolization Avoidance



# CF Monopolization Avoidance

- It is possible for a CF to be using so much of the CF's CPU capacity that the resulting queueing and latch contention in the CF impacts the service time for ALL requests to that CF:
  - Each line represents the Sync response time for a different structure in a given CF.
  - You can see that nearly all structures are experiencing response times of hundreds of mics during the bad interval.




# CF Monopolization Avoidance

To address this situation, APAR [OA56774](#) for z/OS 2.2/2.3/2.4 introduced a new function called CF Monopolization Avoidance. If a CF finds that a structure is monopolizing its resources, it can signal XCF to throttle the rate at which requests are sent to that structure.

Same situation, after enabling this function:

- Only the 'trouble-making' structure is impacted.
- Event that structure receives better response time than before.



# CF Monopolization Avoidance

This function defaults to being turned OFF in z/OS 2.2-2.4. But in z/OS 2.5, the default changed to ON.

RMF reports on number of impacted requests, in the CF Structure Activity Report, but not very easy to quickly spot if it is kicking on.

But you should monitor for messages IXL062E and IXL063I that indicate that requests are being throttled.

**Our recommendation:** ENABLE this function (see SETXCF FUNCTIONS).

For MUCH more information, see **Mark Brooks's** [z/OS Parallel Sysplex Update](#) presentation.



# z/OSMF Items



- One of the enhancements in z/OS 2.5 is a new CFRM Policy Editor function in z/OSMF (and subsequently rolled back to z/OS 2.4 with APAR [PH40176](#)).
- This delivers powerful new capabilities for people familiar with the traditional method of CFRM management:
  - You can display all attributes of all structures in a matrix, and sort each column, so (for example) you can immediately see which GBP structures are defined to use duplexing.
- For people that are *not* familiar with managing CFRM policies, the policy editor greatly reduces the elapsed time to get to a ‘clean’ policy.
- It is not perfect yet, and there are a few RFEs against it, but for most customers we believe it is a big step forward.
- We reviewed it in detail in [Meet the New z/OSMF CFRM Policy Editor](#) article in *Tuning Letter 2021 No. 4*.

# z/OSMF Management Services Catalog

- One of the many concerns related to preparing the next generation of mainframers to eventually take over the reins is – how to pass on your 40 years' of jobs that do just about any task that you are likely to have to do?
- z/OSMF is here to stay, many next generation mainframers are more comfortable with the z/OSMF interface, so it is a logical vehicle to facilitate the handover.
- APAR [PH40810](#) delivers the new z Management Services Catalog function for z/OS 2.4 and 2.5. We believe that this could be a powerful tool to help you hand over those 40 years' of jobs. But it is brand new, and is *not* yet perfect, so this is your ideal opportunity to play with it, decide what you like and don't like, and influence its direction before it is set in stone.
- For more info, see session [Re-imagine Managing Your z/OS Systems Using the z/OSMF Management Services Catalog!](#), by **Marna Walle** and **Travis Biro**.

# z/OSMF File and Dataset REST API vs FTP

- If you find yourself having to download large files from z/OS, and want to do it as efficiently as possible (especially in these days with everyone working from home), you might be interested in the experiences of **Andrew Rowley** from Blackhill Software in Australia.
- Andrew's EasySMF product downloads SMF data to Windows – potentially a very time (and bandwidth) consuming process.
- He discovered that the z/OSMF file and dataset REST API lets you request that the data is compressed with zEDC before being sent to your PC. In his case, using SMF data, this method reduced elapsed time for the download by 90% compared to FTP.
- For more information, see <https://www.blackhillsoftware.com/news/2021/03/23/loading-data-10-times-faster-using-z-osmf/>.

# WLM RTCapLeadTime



# RTCapLeadTime Value

- WLM Sysplex Routing Services (IWMSRSRS and IWM4SRSC) can be used by exploiting products to do smarter dynamic workload routing in a sysplex environment.
  - Some exploiters are Db2 DDF, TCP/IP, WebSphere traditional, WebSphere Liberty, CICS CPSM, maybe others...
- Depending on the interface and parameters used, WLM returns a weight for each server reflecting attributes such as Displaceable Capacity, LPAR Health, and Performance Index.
- z/OS 2.3 added the ability to request WLM to factor upcoming soft capping into its weights:
  - WLM maintains '*Estimated* remaining time, in seconds, before the system is capped due to the defined capacity limit' and '*Estimated* remaining time, in seconds, before the system is capped due to the group capacity limit.' values.

# RTCapLeadTime Value

- This new capability is controlled by the RTCapLeadTime parameter in IEAOPTxx.
  - By default, the estimated time-to-cap values are *ignored*.
  - A non-zero RTCapLeadTime setting tells WLM “Take capping into account in your recommendations if the estimated time to capping is less than x minutes”
- The effect of this *might* be to route less work to a system that is going to be capped soon.
- WLM updates its recommendations as frequently as every 10 seconds, however, WLM is only *providing information* – it is up to each exploiter to decide how often they want an updated recommendation from WLM, and what they then do with that information.
  - Some exploiters use the information to decide which server to connect to, but could then send 000s of transactions over that connection.
  - Others use it when deciding where to send *each transaction*.
- However, it is not easy to know *which products* are using this service (WLM does not report this), or *how* they are using it – that is up to each exploiter.

# RTCapLeadTime Value

- One of our clients did some testing with this function but found:
  - WLM's estimated times to cap were very volatile – apparently predicting the future is not that easy!
    - *But*, they were testing with test LPARs with small caps and very volatile R4HAs.
  - They tested with RTCapLeadTimes of 0 (turned off) or 60, but didn't notice any obvious change in where DDF requests were routed to.
  - But they also found that capping is much less predictable on development and test systems than on production.
    - The RTCapLeadTime value is unique to each system, so you might decide to use a lower RTCapLeadTime value on a development system than on production.
- In any case, they were concerned with the lack of transparency – it might not be obvious externally why a work manager suddenly alters its distribution of work across systems, potentially in wasted time investigating a 'problem' that is actually the work manager behaving as you asked it to.



# RTCapLeadTime Value

- One other thing...
- If you use a dynamic capping product, the WLM estimated time to cap value might be less accurate, because the capping product keeps changing the cap before it cuts in.
  - Make sure that you understand in what situations the product *increases* the cap, and also when it might *decrease* an LPAR's cap.
- There are a *lot* of moving parts here, so we recommend approaching it cautiously.

# RTCapLeadTime Value

- **Our Recommendation:** Unless you have systems with very predictable behavior, *and you know exactly which products use Sysplex Routing Services and how they use it*, enabling this capability *might* just confuse matters.
  - For example, it is probably not a good fit for erratic Test LPARs, or LPARs with very small caps.
  - But, if you have many times when one production LPAR is capped while another has spare capacity, and your workloads can run on multiple systems, this might deliver some help.
- **IBM Recommendation:** If you enable this service, specify a value of 20 or less. The further in advance of the estimated capping time you use, the more volatile the estimated time to cap values will be.
- Is anyone here using this service? If you would like to share your experiences, please email us at [technical@watsonwalker.com](mailto:technical@watsonwalker.com).

# RTCapLeadTime Value

- For more information on WLM support for sysplex workload balancing, refer to:
  - IBM Redbook, *System Programmer's Guide to MVS Workload Manager*, [SG24-6472](#)
  - IBM Redpaper, *Db2 9 for z/OS Data Sharing: Distributed Load Balancing and Fault Tolerant Configuration*, [REDP-4449](#).
  - IBM Manual, *z/OS MVS Planning: Workload Management*, [SC34-2662](#).
  - [Route Me, Workload Manager](#), white paper by **Horst Sinram**.
  - *MVS Workload Manager, Sysplex Routing Services*, white paper by **Robert Vaupel**.
  - SHARE in Pittsburgh 2014, Session [15507](#), *Sysplex Networking Technologies and Considerations*, by **Gus Kassimis**.
  - Information about the WLM Estimated Time to Cap is shown in the RMF Data Portal CPC report header, or in the SMF type 99.1 SMF99\_Time\_to\_Cap and SMF99\_Time\_to\_Cap\_Group fields if you would like to track how accurate/smooth they are in your system.

# Other WLM-related notes

- Don't forget that you can no longer override the Service Definition Coefficients in z/OS 2.5. And there is a WLM Health check to warn you if you are using SDCs other than the values that z/OS 2.5 will use.
- Both **Brad Snyder** (aka 'Not-Kathy Walsh') and **John Baker** recommend NOT enabling WLM I/O Priority management.
  - See Brad's [Performance Hot Topics](#) and John's [Just Say No: A Case for Eliminating I/O Priority Management](#) sessions for more information.
  - Note that the default is OFF, but IBM had been recommending to turn it ON.
- If you moved to Tailored Fit Pricing, consider using WLM service class-level HONORPRIORITY option to force non-response-time-critical zIIP-eligible work to run on zIIPs.
  - Also, see Brad's presentation for considerations about using ZIIPAWMT to reduce zIIP overflow to GCPs, and situations where you should NOT alter the ZIIPAWMT setting.

# Other WLM-related notes

- Last thing – TURN ON YOUR SMF TYPE 98 RECORDS!
- Even if you don't have Workload Interaction Navigator or similar tools to analyze them, IBM and other vendors might be able to use the Type 98 records to help you analyze performance issues.
  - If you have z/OS 2.4 and RMF, or z/OS 2.5 and IDG, you can use zWIC at no additional charge. See APAR [OA62268](#) for more information.
- **This is now an IBM Best Practice** – see Brad's Performance Hot Topics handouts.

# JES2 Spool Compression

# JES2 Spool Compression

z/OS 2.4 added support for compressing and encrypting JES2 spool files.

IF z/OS is running on a z15 or later, zEDC will be used to perform the compression and decompression.

- If running on an earlier CPC, compress requests will be ignored, and decompress requests will be handled in software → it is probably a good idea to hold off doing anything with this until all your CPCs are z15 or later.

**Note that this use of zEDC does NOT require the chargeable zEDC feature of z/OS.**

For more information, and some sample performance numbers, see **Tom Wasik's** [\*JES2 Spool Compression and Encryption\*](#) session.

# NO Page Data Sets??!!





We recently heard from a customer that has loads of memory in their z/OS LPARs, plus Storage Class Memory (SCM), and that plans to *eliminate* all-but-one of their page data sets.



We *never* heard of anyone doing this before, so we consulted WSC guru **Kathy Walsh**.

Kathy confirmed that it *is* possible, and feasible. But you must keep one page data set, *and* you must have sufficient memory in your z/OS LPARs.

- We don't yet have a formula for determining how much is 'sufficient'.
- If you are interested, email us at [technical@watsonwalker.com](mailto:technical@watsonwalker.com), or contact Kathy directly.

# Misc SHARE Info



A few reminders:

- The recordings of the ‘Live-stream’ sessions will only be available until May 31<sup>st</sup> (**just 8 weeks away**). And when they are gone, they are ... hmmm – watch this space. Unfortunately, it is *not* possible to download them, so make sure to catch any sessions you were double-booked for.
- The next SHARE will be in Columbus, Ohio, from August 21<sup>st</sup>, to the 27<sup>th</sup>. That’s right, a *full-week* event again. Book your flights and hotels *now*, before rates go up.
- Please remember to complete your evaluation for this session, and for *all* the sessions you’ve attended. Your evaluations help the organizers select sessions for future events – if you want more sessions on your favorite topics, communicate that through your evaluations.

# Your feedback is important!

## Submit a session evaluation for each session you attend:

[www.share.org/evaluation](http://www.share.org/evaluation)

