

Watson & Walker zRoadshow Summer 2022 Edition

Mario Bezzi – Watson & Walker (Mario@watsonwalker.com)

Frank Kyne – Watson & Walker (Frank@watsonwalker.com)



Session 46885 – August 2022

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 helping clients prepare for z16, and convincing senior management that the zSkills issue is not fake news.
- For more information, see our website, www.watsonwalker.com.
- The objective of *this* session is to make you aware of interesting tidbits that we have come across recently, and point you to where you can get more information.

- Frank:
 - zSkills Drought - Preparing for a hassle-free Career 2.0:
 - z/OS Management Services Catalog
 - CFRM Policy Editor
 - Reducing reliance on Assembler skills for JES2 and MVS exits
 - Misc.
 - Tips for decommissioning software
 - RMF-related changes
 - RRS/Logger tip
 - XCF Signaling status and tips
 - z16 CF Link performance improvements
 - JES2 spool compression
 - Reporting on Z Sort exploitation progress
 - Pointers to interesting information

Topics

- Mario:
 - Information about important performance-related SMF records.
 - Db2 DDF for the Rest of Us
- But first....

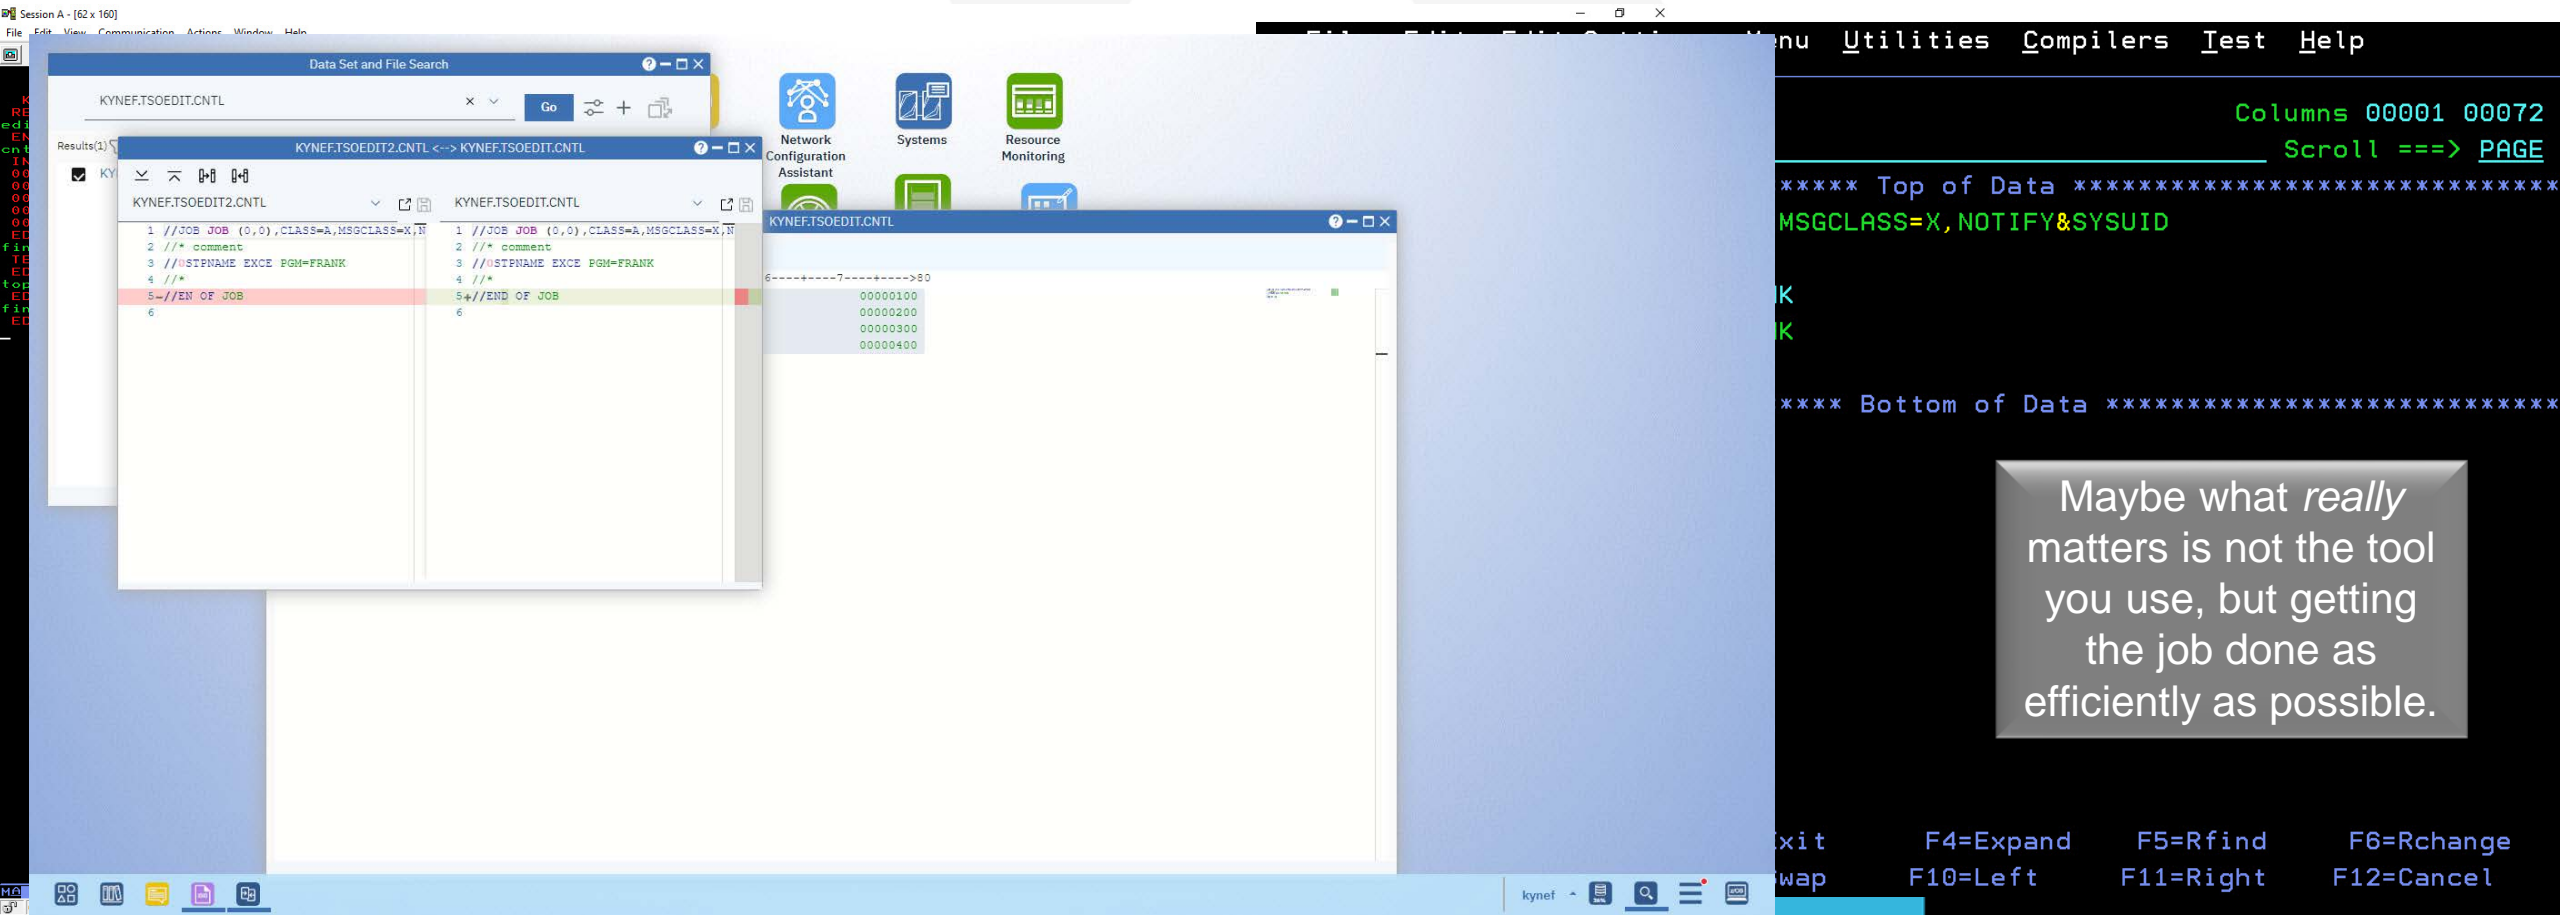
Show-of-Hands Survey Questions

- Anyone running on z16 yet?
- How many are using zCX in production?
- How many are using OpenShift on z/OS?
- Anyone installed a z/OS ServerPac using z/OSMF yet?
- Anyone have a z15/z16 but are NOT licensed for zEDC?
- Have you enabled Z Sort with your sort product yet?
- Is there anyone that does NOT have a skills or demographics issue in your company?

Addressing the Looming Skills Shortage

zManagement Services Catalog


- Does anyone remember this – Line Mode TSO Editor?
 - When I started working, ‘real’ sysprogs didn’t believe in ISPF...



Maybe what *really* matters is not the tool you use, but getting the job done as efficiently as possible.

- Most of us will not have time to pass all our years' of experience over to our young colleagues before we are promoted to 'Career 2.0'. But there *are* things you can do *now* to make their and your lives easier..
- ONE of them is to use the z Management Services Catalog to help hand over your PDS full of useful (but cryptic and uncommented) jobs that you have been building for the last 40 years.
 - As an incentive... Another potential (labor-saving) use of zMSC would be to give users access to a zMSC service and let them kick off processes that are currently semi-automated.
- zMSC is built on top of z/OSMF Workflows, and lets you add variables, logic, access controls, and logging to workflows, making it easier for less experienced staff to perform tasks with fewer opportunities for 'learner's mistakes'.

z/OS Management Services Catalog

- What z/OS releases does it run on? z/OS 2.4+.
- How do you install it? Initial release was delivered by z/OSMF APAR [PH40810](#).
- Where is it documented? z/OSMF product manuals, but mainly online help.
- Recent enhancements? APAR [PH44234](#) delivers Release 1.1.
 - APAR [PH48001](#) Provides a fix for help messages.
 - The Comp ID to identify additional APARs is 5655S28MS. 
- Is it perfect? Of course not! MVS also isn't perfect (yet), and it has been around a little longer.
- Install it, try it out, and let IBM know what enhancements you would like.
- Watch out for **Hiren Shah's** upcoming z/OS Community blog post ([here](#)) and videos about how to quickly turn an existing job into a z/OSMF workflow, and a workflow into a zMSC service.

z/OS Management Services Catalog

- Reference material:
 - SHARE in Dallas 2022, Session [Re-imagine Managing Your z/OS Systems Using the z/OSMF Management Services Catalog!](#), by **Hiren Shah**
 - SHARE in Columbus 2022, Session [Modernize z/OS System Management with As-a-Service Approach](#), by **Hiren Shah**
 - z/OS Management Services Catalog [demo video](#).
 - [Lifecycle of a Service](#) video
 - z/OS Management Services Catalog [Content Solution page](#)
 - [z/OS Management Services Catalog Lab](#) slides
 - [z/OSMF Guild z/OS Management Services Catalog session](#)
 - [‘What’s New in z/OS 2.5 \(Part 2\)?’](#) article in *Cheryl Watson’s Tuning Letter 2021 No. 4*

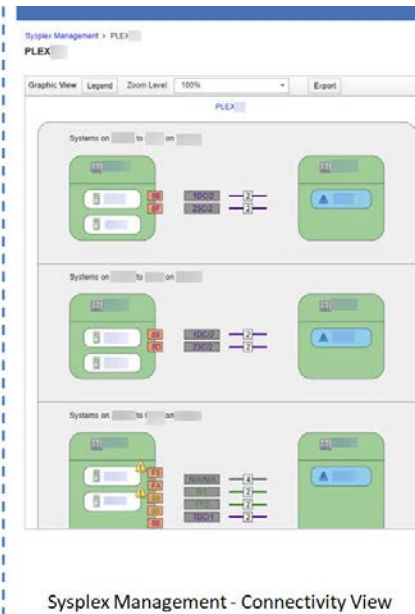
- Most importantly....
- **Cheryl** says zMSC is key to a less painful transition to the next generation of techies.
- Remember what Cheryl said about z/OSMF when everyone else was writing it off.
- I'm not saying that Cheryl is *always* right, but ...
- This is the IT equivalent of “buying Amazon at 4”, so get in there now, and be one of the people that influences the direction of this key function.
 - IBM is looking for sponsor users to work with zMSC development. If you are interested, contact Hiren or [email us](#).

CFRM Policy Editor

- One of the enhancements in z/OS 2.5 was a new CFRM Policy Editor function in z/OSMF (and rolled back to z/OS 2.4 with APAR [PH40176](#)).
- This delivers powerful new capabilities for people familiar with traditional 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.
- And 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.



Sysplex Management - Structure View



Sysplex Management - Connectivity View

- APAR [PH44343](#) delivers the support to let you do side-by-side comparison of two policies, and the ability to export policy data into CSV file.
 - For more z/OSMF Sysplex Management app APARs, track COMPID 5655S2809.
- See **Dave Surman's** excellent [What's New in the Sysplex Management Plug-In?](#) session.
 - One of the enhancements that is being worked on is for the policy editor to show the defaults for all the attributes that you didn't specify, along with an indication of whether the default was hard-coded or not. MUCH easier than checking manual.
 - Also, note Dave's request for enhancement requests.
- We reviewed the Policy Editor in detail in [Meet the New z/OSMF CFRM Policy Editor](#) article in *Tuning Letter 2021 No. 4*.

- Every site relies on system exits written in Assembler, but every site is also facing a reduction in the number of people with Assembler skills.
- There are some *free* options for replacing exits:
 - JES2 2.4 delivered JES2 Policies to replace some JES2 exits, with more policies on the way in future z/OS releases.
 - SMFLIMxx Parmlib member maybe enable you to eliminate IEFUSI exit.
- However, the people that drive the replacement process *must* understand Assembler.
 - For example, the biggest challenge in replacing IEFUSI exits with SMFLIM is not lack of capability in SMFLIM, it is in trying to understand horrendously complex logic in some IEFUSI exits. Who better to simplify that logic, or even replace it with SMFLIM, than the people that wrote and maintain the IEFUSI exit?


Assembler Skills

- A chargeable option is Trident Services' zOSEM product.
- *One* of the functions it provides is to replace Assembler exits with parmlib-like definitions.
- If you still have in-house Assembler skills, those people can create zOSEM definitions to provide the exit functionality.
- If you *don't* have people with those skills, Trident will analyze your Assembler exits for you and help you create the zOSEM replacements.
- Trident provide a no-charge 'Exit Explorer' tool that identifies which exits are in use in your systems – an excellent first step.
 - For more information about the Exit Explorer see https://www.triserv.com/exit_explorer_utility.pdf.

Assembler Skills

- You can find information about JES2 policies in:
 - SHARE in Columbus 2022 Session [Implementing JES2 Policies from an IBM and Customer Perspective](#), by **Tom Wasik** and **Mike Shorkend**.
 - Chapter 74, [JES2 Policies](#), of [z/OS JES2 Installation Exits](#).
 - [‘Introduction to JES2 Policies’](#) article in *Tuning Letter 2022 No. 1*, by **Mike Shorkend**.
- You can find information about SMFLIM in:
 - [Virtual Storage Overview](#) in MVS Initialization and Tuning Guide.
 - [SMFLIM description](#) in MVS Initialization and Tuning Reference.
 - Enterprise Tech Journal article [Recent SMF Parameters](#), by **Cheryl Watson**.
 - IBMMAIN [“SMFLIM Sample”](#) discussion.
 - Upcoming Tuning Letter article.

Demand for z/OS Skills

- Just in case you (or your managers or HR friends) are in any doubt about the demand for people with z/OS system programming skills, we heard of a system programmer that received a \$40,000 *increase* when they changed jobs recently.
- (Rumor has it that the ‘smart’ money is getting out of Bitcoin and stockpiling system programmers instead )

Miscellaneous Tidbits

(Non)Reader Beware

- Given the ultra-competitive mainframe software market, and corporate focus on cost reduction, it is not unusual to see products being decommissioned, either because they are no longer required, or to be replaced by some other product.
- It is natural that the owner of any software product would want assurances that it is no longer installed when a client terminates a contract.
- Most of us would expect 'no longer installed' to mean all product libraries and data sets have been deleted from all systems.

This is important

(Non)Reader Beware

- One of our eagle-eyed clients noticed a clause in a software contract stating that all product data sets AND ALL BACKUPS or copies of those data sets must be deleted if the product is decommissioned.
- Could *you* identify every backup of every data set in your site?
- If you could, do you have the ability to delete selected backup data sets?
 - How would you handle full volume dumps containing a product's data sets?
 - How would you handle backups of ISV.**?
 - How long do you keep those backup data sets for?

(Non)Reader Beware

- Some suggestions (from our readers):
 - Amend your contract review process to ensure that no new contracts contain such clauses.
 - Review *existing* contracts NOW to determine if any contain such clauses.
 - Review your existing system software backup processes to ensure you can delete backups at the individual product level if necessary.
 - Use TADz or similar to show the pattern of product use BEFORE the product is uninstalled, and again after it is uninstalled, and **archive those reports**.
 - Full volume dumps (or HLQ.** dumps) of system software volumes should have reasonably short retention periods.
 - Consider using a Management class with a short backup retention period for product data sets.

RMF-related Changes

- If you are working on optimizing your HiperDispatch topology, it is very helpful to see where logical CPs are assigned. This information is available in the SMF 99.14 records, which is great. However, many sites suppress all SMF type 99 records, meaning that this information is not available to them.
- To make this information more widely available, APAR [OA62064](#) enhances ADG/RMF to include this information FOR ALL LPARS ON THAT CPC in the type 70 records if z/OS is running on a z16.
- Also related to RMF is announcement [221-338](#), stating that the z/OS Workload Interaction Correlator is entitled (FREE) to licensed RMF or ADG customers.
 - Previously this was a chargeable z/OS feature code (see [OA62268](#)).
- And don't forget recent best practice '[Always Collect Correlator SMF Records](#)'.
 - This relates to the HFTSINTVL parm in SMFPRMxx – default is NOHFTSINTVL.

RRS/Logger Tip

- One of our clients recently encountered a problem trying to empty out an old disk subsystem – there was one Logger offload data set that was ALWAYS allocated, meaning that it couldn't be moved.
- Data set was part of the RRS DELAYED.UR log stream.
- IBM Service recommendation was to stop RRS on every system and then move the data set.
 - But RRS is used on *all* their systems, so stopping every RRS at the same time was akin to performing a sysplex IPL – not an option.
- Fix was:
 - D GRS,RES=(SYSDSN,dsname) to see which system(s) is connected.
 - SETLOGR FORCE,DEALLOC,lsname on that system(s).
 - Immediately migrate the data set.

XCF Signaling Tip

- **Mark (“MAXMSGs are the bain of my life”) Brooks** says that MAXMSG of 2000 should be large enough for nearly any sysplex when using XCF Transport class enhancements in z/OS 2.4 or later.
- If you want the gory details, see Mark’s outstanding [z/OS Parallel Sysplex Update](#) session.
- And make sure that the PTF for APAR [OA60480](#) (**HIPER**) is applied.
- And subscribe to OPEN APAR [OA62980](#).
- And make sure that your PATHIN MAXMSG values on any system don’t total to more than 800MB.
- And, just in case you’re still not sure about MAXMSG values – from Mark’s presentation:

But the real problem is overcommitting MAXMSG.
Don’t do that.

z16 Coupling Performance

- Per z16 [Announcement letter](#):
 - For ICA-SR links “While the amount of improvement clients will see is highly dependent on various workload and configuration specifics, service times for short-distance CF read requests or lock requests may **improve by up to 10%**, write requests may **improve by up to 25%**, and duplexed write requests may **improve by up to 25%**, compared to CF service times on IBM z15 systems.”
 - “These efficiency improvements are maximized when both ends of the coupling link connection reside on IBM z16 systems. However, some improvements are also realized when connecting an IBM z16 system to an IBM z15, IBM z14, or IBM z14 ZR1 generation system.”
 - For CE-LR links (using new, z16-only, CE2-LR adapters), expect improved link throughput (most of the service time is due to distance (speed of light), so don’t expect to see noticeable response time changes).
 - Note that Coupling Express LR adapters from previous machine generations cannot be carried forward to IBM z16. Only Coupling Express2 LR new-build adapters may be used to provide long-reach coupling connectivity for IBM z16 systems.

JES2 Spool Compression

- JES2 spool file compression uses zEDC:
 - Must be running on z15 or later
 - z/OS 2.4 or later
- Does NOT require the chargeable zEDC feature code on z/OS.
- Testing by JES2 group found 11:1 compression ratio AND better performance.
- According to **Tom Wasik**, ‘this is a no-brainer’.
- See [Compressing Spool Data Sets](#) for more info.

Z Sort Exploitation by DFSORT

- After a slow start, we are seeing a steady increase in the number of customers using DFSORT Z Sort support.
- But this immediately generates some questions:
 - Which sorts ARE using Z Sort, and what benefit are they getting?
 - Which sorts are NOT using Z Sort, and why not?
 - At the system level, what is the breakdown between using, not eligible to use, eligible-but-not-using?
 - And what is the most common cause of eligible-but-not-using?
- To address this, we recently worked with **Sri Hari Kolusu** in DFSORT devt to create some [sample ICETOOL report jobs and accompanying documentation](#).
 - All the jobs need is ICETOOL and SMF 16 TYPE=FULL DFSORT SMF records.

DFSORT Z Sort Exploitation Summary Report

Summary report of zsort usage/nonusage 2022/07/10 10:00:47

ZSORT Code	System-id	Description	Num Steps	Total CPU(HS)	Total SRB(HS)
000	FPKA	zSORT is used	953	1,158.56	224.41
194	FPKA	zSORT is not enabled/activated	184	32.89	0.78
195	FPKA	SORT FIELDS=COPY or OPTION COPY is used	21,895	5,107.40	235.13
196	FPKA	MERGE FIELDS= is specified	428	3.82	0.00
197	FPKA	INREC is specified	15	6.13	0.12
198	FPKA	OUTREC is specified	27	78.99	9.45
199	FPKA	OUTFIL is specified	1,199	407.50	6.17
200	FPKA	Program Invoked Sort or Db2	7,037	1,158.56	20.93
201	FPKA	Failed to get the required memory object	1	0.59	0.00
209	FPKA	Minimum memory requirements are not met	2	0.02	0.00
211	FPKA	Unsupported sort fields	1	0.00	0.00
212	FPKA	SUM FIELDS= is specified	1,417	20.24	20.24
213	FPKA	JOINKEYS Maintask, Subtask1 or Subtask2	380	5.22	5.22
214	FPKA	MODS Statement is specified	351	0.01	0.01
231	FPKA	Size of the file to sort < 32 MB	10,860	0.00	0.00
Totals :			44,750	10,361.49	523.07
000	FPKB	zSORT is used	689	1,028.30	119.89
194	FPKB	zSORT is not enabled/activated	2	2.26	0.07
195	FPKB	SORT FIELDS=COPY or OPTION COPY is used	3,024	352.31	43.91
196	FPKB	MERGE FIELDS= is specified	15	0.10	0.01
197	FPKB	INREC is specified	6	3.02	0.06
198	FPKB	OUTREC is specified	13	9.15	0.20
199	FPKB	OUTFIL is specified	10	25.77	1.07
200	FPKB	Program Invoked Sort or Db2 REORGS FLR Sort	985	4,246.78	298.59
201	FPKB	Failed to get the required memory object	1	0.31	0.01
211	FPKB	Unsupported sort fields	4	1.89	0.34
212	FPKB	SUM FIELDS= is specified	512	493.28	8.36
213	FPKB	JOINKEYS Maintask, Subtask1 or Subtask2	542	227.93	2.92
214	FPKB	MODS Statement is specified	39	94.38	2.29
217	FPKB	Input or output is a VSAM cluster	1	0.35	0.01
231	FPKB	Size of the file to sort < 32 MB	5,918	516.64	0.25
240	FPKB	Failed to get memory obj of 75% of file size	6	2.12	0.09
Totals :			11,767	7,004.59	478.07

Why Z Sort wasn't used?

How many sort job steps?

TCB and SRB time for each Zsort Code

Summary by system

zBNA Z Sort Support

- ZBNA V2.2.4 added support for identifying DFSORT job steps that could potentially benefit from Z Sort.
 - Note that not all sorts are eligible to use Z Sort, so you should not expect to see *all* 'SORT' job steps using Z Sort.
- The joblog from a sort step will report the Z Sort status on message IEC267I.
 - It worked, or
 - ONE reason why it was not eligible.
- zBNA is able to see ALL the potential reasons why a step might not use zSort, so it can give you a more comprehensive report.
- However, it currently is not very valuable for helping you track your progress once you start enabling Z Sort.
 - zBNA team are working on enhancements, to be delivered by end of September.

zBNA Z Sort Support

- Note that both zBNA and the new ICETOOL reporting programs only report on DFSORT usage of Z Sort.
- Db2 internal sorts are another potentially major user of Z Sort, but that info is not reflected in the SMF Type 16 records.
- However, Extended counter 255 in the z15 SMF Type 113 records reports the total number of 'SORTL' calls.
- Not perfect, but it can be used to get a *little* more insight into Z Sort exploitation.

Hot diggity dog

- 2022 Hot Topics articles:
 - [How to turn your SMF data into valuable insights without z/OS expertise](#)
 - [Real-time network crypto enforcement with zERT](#)
 - [Managing Compliance with Ease on IBM Z](#)
 - [RedHat Open Shift Container Platform on IBM z/OS Container Extensions](#)
 - [Changes coming to ftp servers and sending documentation to IBM](#)
 - [IBM Z and LinuxONE Content Solutions](#)
 - [BCPii - A RESTed development](#)

Upcoming conferences

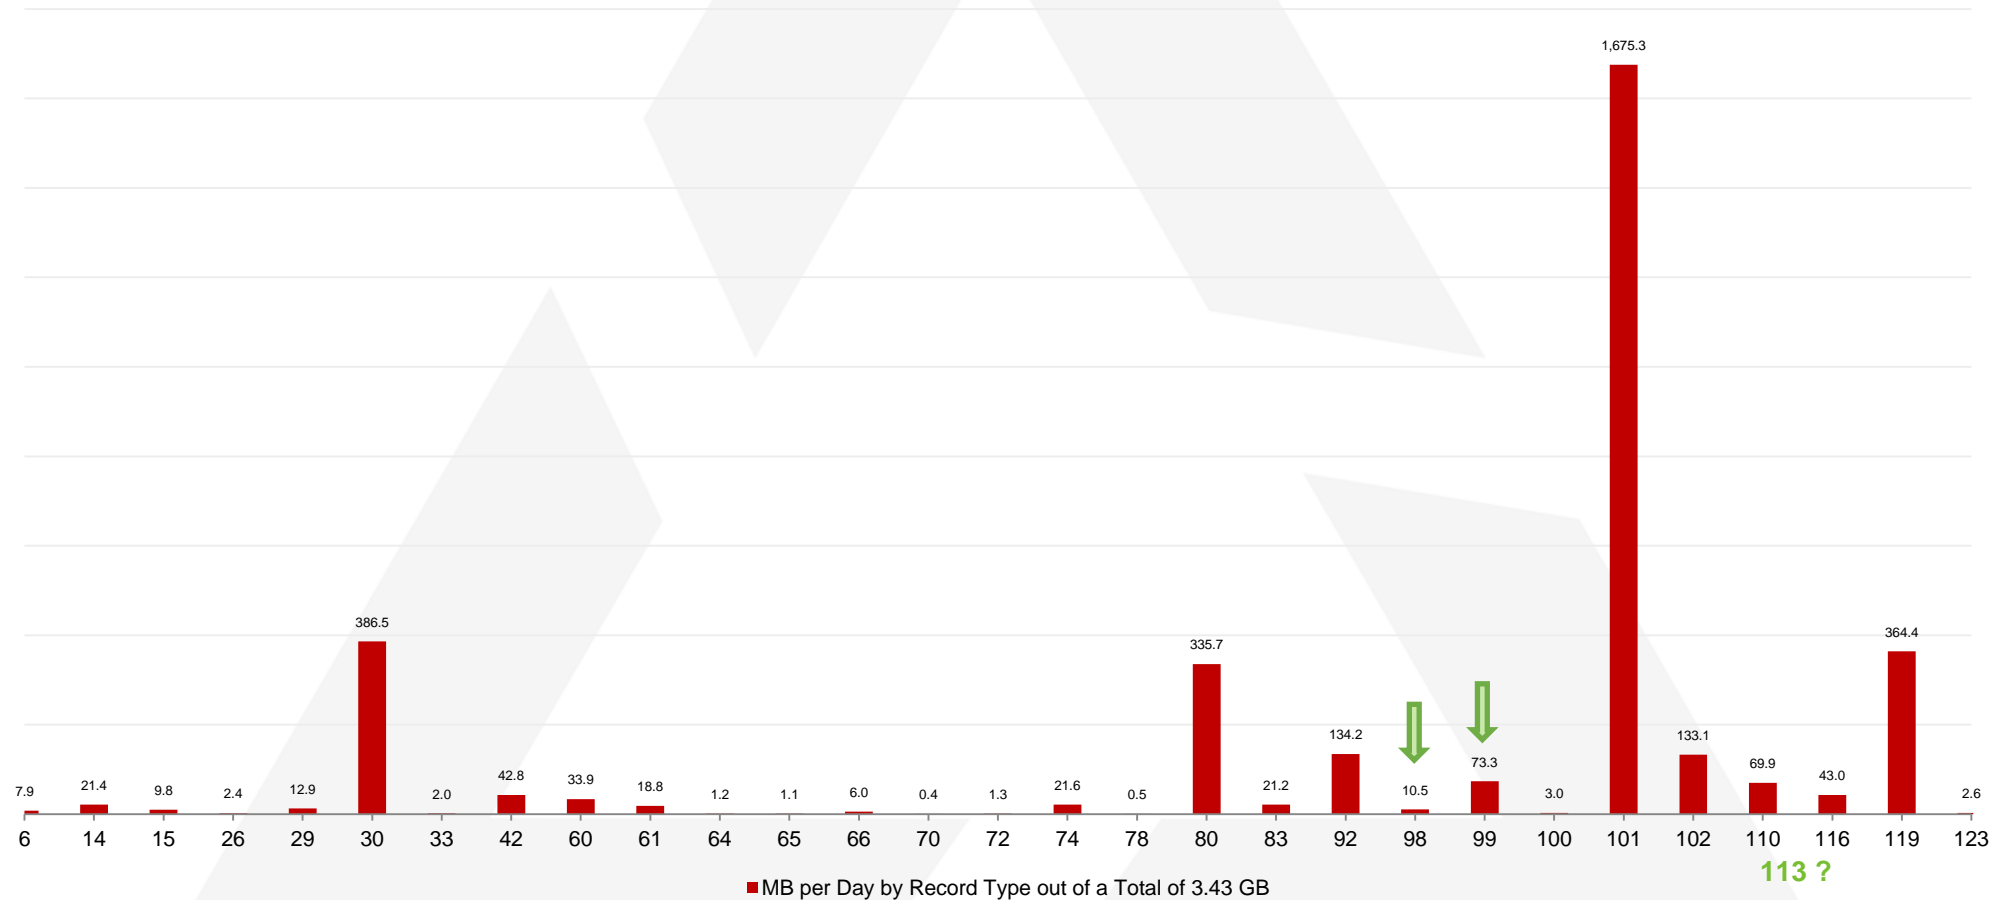
- Broadcom [Mainframe Technical Exchange](#)
 - 2nd to 4th October, virtual, no charge
- [UK GSE annual conference](#)
 - 1st to 3rd November 2022 (In-Person)
 - 7th to 9th February 2023 (Virtual)
- [SHARE Winter 2023, Atlanta](#)
 - 5th to 8th March, 2023 **Note shorter week – Sunday to Wednesday**
- [SHARE Summer 2023, New Orleans](#)
 - 13th to 18th August, 2023

Mario's Section

Gentle reminder - Important SMF records


- Please make sure that you collect the following important SMF records:
 - SMF 98 – High Frequency Throughput Statistics
 - SMF 99 – WLM Decisions
 - SMF 113 – Hardware capacity, reporting, and statistics
- For SMF 98 set the HFTSINTVL in SMFPRMxx to 5 seconds.
- For SMF 99 consider omitting subtype 13 (Peter Enrico's suggestion)
- For SMF 113 make sure to activate the CPU Measurement Facility
- They are VERY low volume if compared with what CICS and DB2 write. For example..

Large system – 3.43 GB of SMF per day,



Record types contributing less than 0.01% of the total have been omitted


Db2 DDF – A Quick Reference for the Rest of Us



Why DDF

- DDF – Distributed Data Facility is a popular technology which allows an application running on a given system to remotely access a Db2 data base hosted by another system.
- DDF can be used from Db2 on z/OS to Db2 on z/OS, or more commonly from the distributed environment to Db2 on z/OS.
- Many customers use DDF and Java to [re] implement business logic on distributed while keeping data on Db2 for z/OS.
- Applications on distributed can also use DDF and Stored Procedures to reuse existing mainframe based business logic.
- DDF workload is eligible for *some* zIIP offload.

Customer X – DDF Trxs/sec

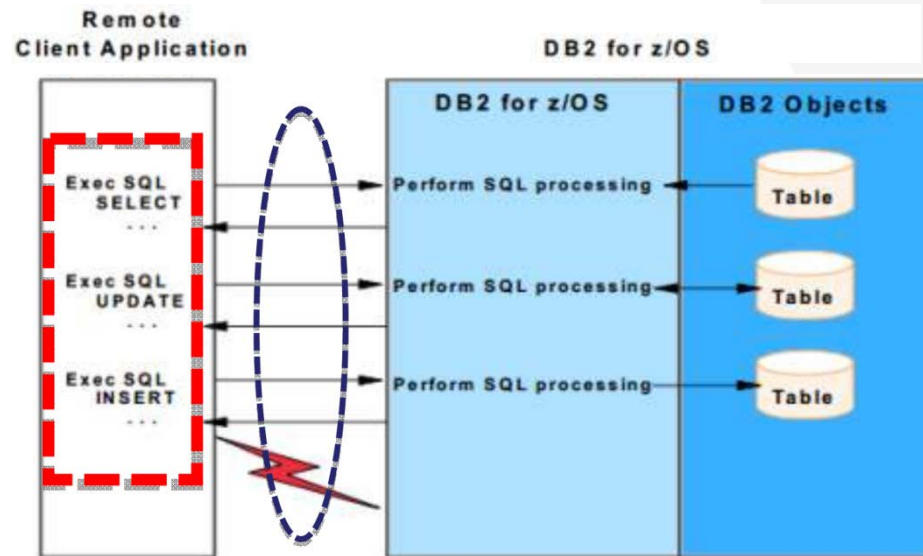


Java and Db2 for z/OS

- Java accesses Db2 on z/OS via a data source object, this is the abstraction layer used by Java to access any relational data base. Remote connectivity to Db2 for z/OS is provided by the JDBC Driver – officially "[The IBM Data Server Driver for JDBC and SQLJ](#)"
- The JDBC driver is installed on the *remote client*, and works with Db2 using a series of generic packages to be installed on Db2 for z/OS.
- JDBC based Java applications make use of dynamic SQL. Dynamic SQL don't require a package, as statements are prepared at execution time.
- A remote client normally network-connects to Db2, executes some remote SQL activity and eventually remote-disconnects

Move business logic off z/OS using Db2 DRDA

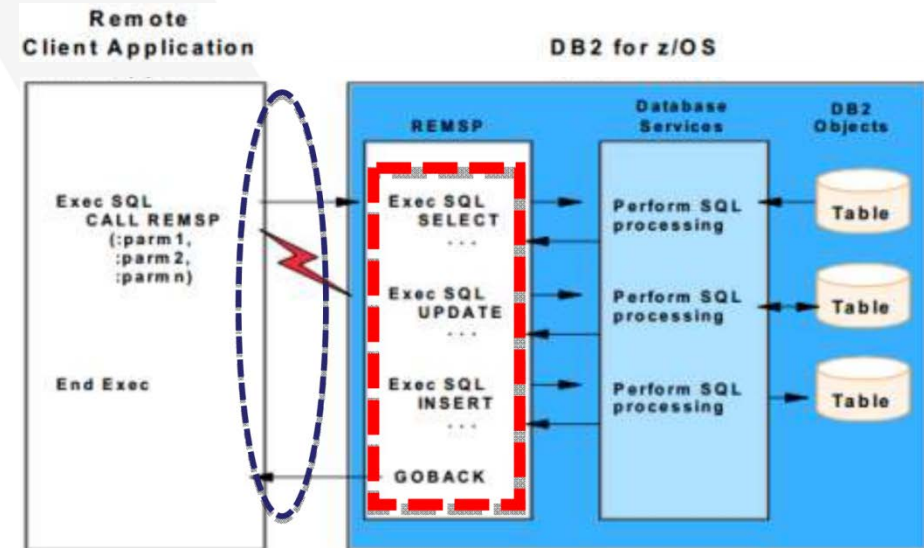
DDF



All business logic moved to distributed

55-60% of ALL SQL processing is zIIP eligible

DDF + External Stored Procedures



Reusing some existing mainframe business logic

55-60% of Remote SQL processing is zIIP eligible

JDBC connection pooling

- The JDBC driver optionally supports pooling connections. This is used to avoid the CPU cost of creating and destroying a new network connection for each transaction (both client and server side).
- With connection pooling the data source determines whether a suitable physical connection already exists, and in case reuses it. When the JDBC application closes the connection, the data source returns the physical connection to the connection pool for future reuse.
- Connection pooling is configured in the JDBC driver on the remote client. We recommend you to make sure it is enabled.

DDF connection routing

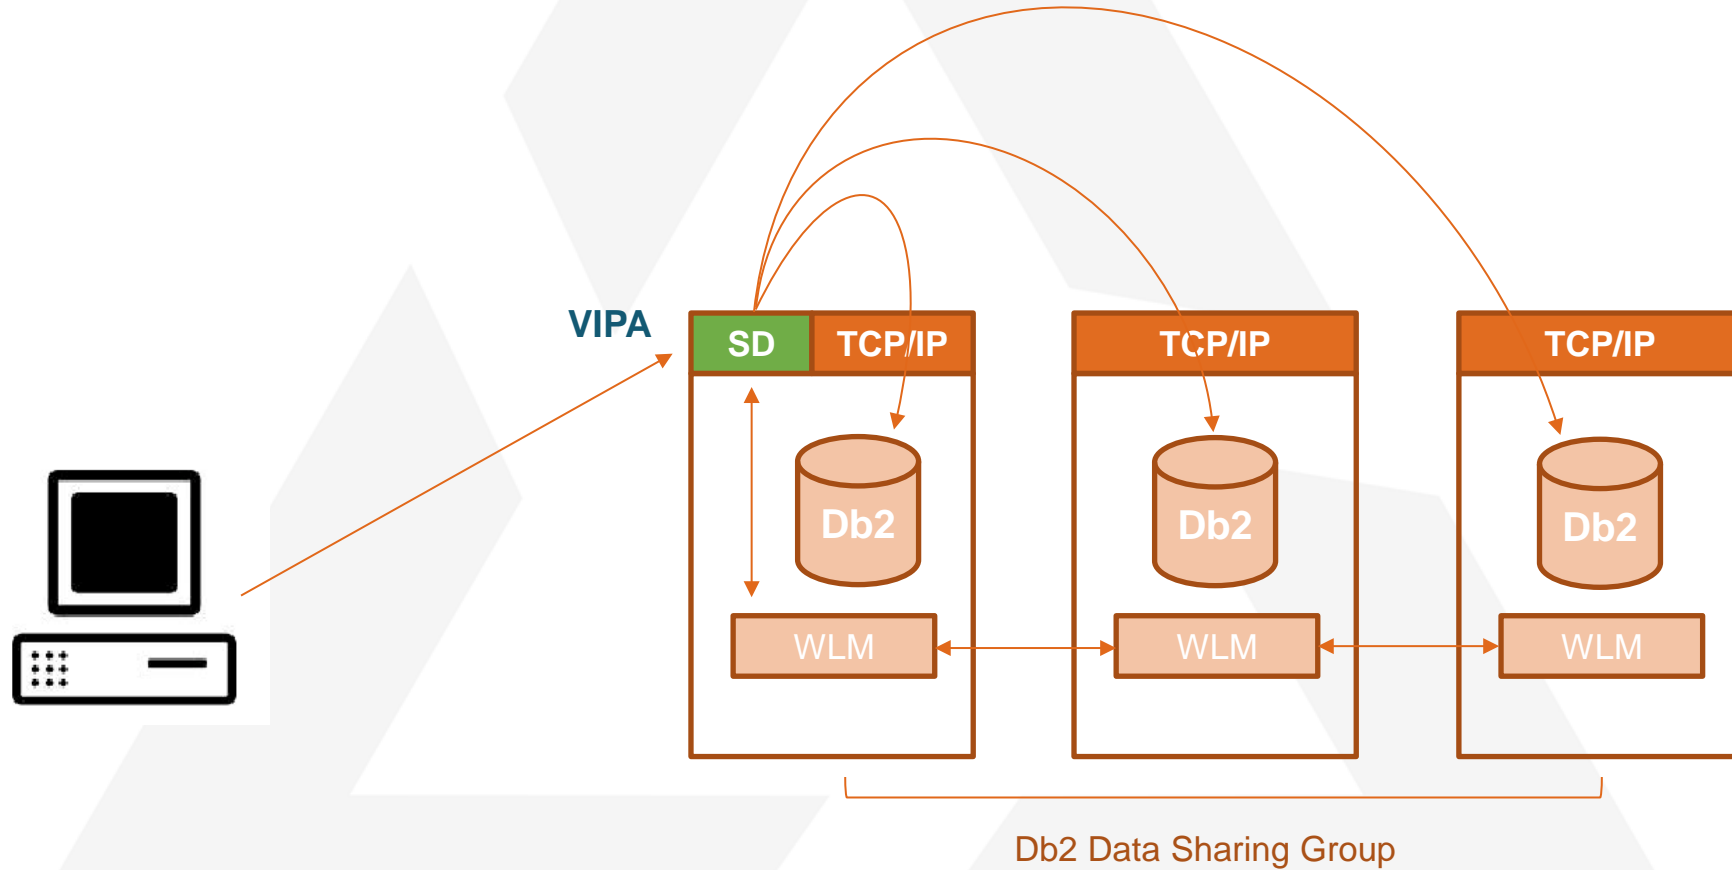
- From a workload balancing perspective, if running data sharing, you can use Dynamic VIPA and Sysplex Distributor (SD) to balance incoming connections among Db2 members. **Connections, NOT transactions.**
- SD offers multiple routing algorithms, we will focus on the two most commonly used: BASEWLM and SERVERWLM. From their names you can see they are both influenced by WLM. **Influenced, NOT driven.**
- For both the above algorithms original WLM recommendations **are modified by SD** to also consider target server responsiveness (TSR) as seen by TCP/IP. TSR values < 100 mean that the target server struggles keeping up with incoming connections.

- The BASEWLM algorithm routes connections simply considering free capacity on each of the target LPARs.
- By default the BASEWLM algorithm doesn't take into account free zIIP capacity, but you can optionally set a relative weight for CP and zIIPs for WLM to use in its calculations.
- DDF workload is zIIP eligible. If you use the BASEWLM algorithm make sure to assign a weight to zIIP processors. See the PROCTYPE option of your [VIPADISTRIBUTE](#) TCP/IP configuration statement.

VIPADISTRIBUTE SERVERWLM

- The SERVERWLM algorithm routes connections considering displaceable capacity as seen from the standpoint of each target Server Address Spaces.
- In case of DDF the target Server Address Spaces are the xxxxDIST address spaces of each data sharing group member.
- With SERVERWLM, connections are routed basing on the ability of the target Server Address space - xxxxDIST - to accept the connection request, **not to achieve the objective transaction response time.**

DDF and Sysplex Distributor connection routing



VIPADISTRIBUTE SERVERWLM

- Specialty engines are included by default by WLM in its calculations for SERVERWLM, also considering the actual amount of zIIP capacity used by the applications.
- Optionally Server Address Spaces can inform WLM about their health status and about the number of transactions abnormally completed. These will also be taken into account by WLM when making its routing recommendations.

Checking how connections are distributed

```
SYSA      2022237  15:42:16.47      -RO SYSB,D TCPIP,,N,VDPT,IPADDR=10.252.237.102
SYSB      2022237  15:42:16.48  STC96264  EZZ2500I NETSTAT CS V2R4 TCPIP 903
                                     DYNAMIC VIPA DESTINATION PORT TABLE FOR TCP/IP STACKS:
                                     DEST IPADDR      DPORT DESTXCF ADDR      RDY TOTALCONN  WLM TSR  FLG
                                     10.252.237.102  03340 10.113.145.234  001 0000474262  10  100
                                     10.252.237.102  03340 10.113.145.238  001 0000433060  10  100
                                     10.252.237.102  03340 10.113.145.239  001 0000466328  10  100
                                     10.252.237.102  03340 10.113.145.240  001 0000430361  10  100
                                     4 OF 4 RECORDS DISPLAYED
                                     END OF THE REPORT
```

Checking how connections are distributed

```
SYSA      2022237  15:42:57.08
SYSB      2022237  15:42:57.09  STC96264  -RO SYSB,D TCPIP,,N,VDPT,DETAIL,IPADDR=10.252.237.102
EZZ2500I NETSTAT CS V2R4 TCPIP 267
DYNAMIC VIPA DESTINATION PORT TABLE FOR TCP/IP STACKS:
DEST IPADDR      DPORT DESTXCF ADDR      RDY TOTALCONN  WLM TSR FLG
10.252.237.102  03340 10.113.145.234  001 0000474265 10  100
DISTMETHOD: SERVERWLM
TCSR: 100 CER: 100 SEF: 100
WEIGHT: 41
RAW          CP: 41 ZAAP: 00 ZIIP: 30
PROPORTIONAL CP: 39 ZAAP: 00 ZIIP: 01
ABNORM: 0000          HEALTH: 100
ACTCONN:      0000000167
QOSPLCACT:   *DEFAULT*
```

W/Q: 10

DDF Transaction Routing

- When you use JDBC driver's ability to pool connections, these tend to stay open, effectively reducing the value of SD connection balancing features.
- Transaction level routing is optionally made available by the cooperation of Db2, WLM and the JDBC driver.
- Setting the JDBC property `enableSysplexWLB` to `TRUE` makes the JDBC driver balance workload at transaction boundary. When a transaction completes (`COMMIT` or `ROLLBACK`), the next one may be sent to a different Sysplex member. To do this the JDBC driver makes use of routing recommendations provided by Db2. The `xxxDIST` address space obtains them from WLM querying it periodically, and sends them back to the JDBC driver.

WLM Routing recommendations and Soft Capping

- If you are using SoftCapping, you may want SD to start routing new connections away from a given LPAR if they are getting close to being capped.
- WLM routing recommendations can optionally take SoftCapping into account if you set the IEAOPTxx RTCAPLEADTIME option. RTCAPLEADTIME specifies, in minutes, how long in advance an upcoming soft capping condition should influence WLM's sysplex routing recommendations.
- The default for WLM is to **NOT** consider upcoming softcap conditions.

DDF Transaction execution – Db2 view

- From a z/OS perspective a remote request enters the xxxxDIST address space, there it is associated by Db2 to an available Data Base Allied Thread (DBAT) if any, and runs into a WLM independent enclave.
- Execution parallelism is controlled using the Db2 MAXDBAT option (Default 200).
- "*Regular*" DBATs are normally pooled, and reused for incoming requests as needed.
- You can use High Performance DBATs to reduce the z/OS CPU cost of running DDF transactions. The use of High Performance DBATs has ramifications, and a complete coverage of the topic is out of scope here.

DDF Transaction execution – WLM view

- Normally a new independent WLM enclave is created for every transaction. This allows WLM to classify transactions individually. DDF transactions run inside the DDF Address Space but have their own WLM importances and goals, different from those of xxxxDIST.
- Service Classes used by DDF transactions can be multi-period and have either velocity or response time goals. For short running transactions a response time goal is recommended.
- While WLM offers a large number of classification criteria for DDF, finding the ones able to identify the actual transaction type is usually tough, and may require changes in the data source or even in the remote application.
- With an individual enclave representing a single transaction WLM is able to track and report the number of executed transactions, their response time, their resource consumption.

What happened lately with High Performance DBATs



- In 2021 IBM discovered an issue with Db2 v12 and high volume DDF transactions executed using High Performance DBATs. The issue was due to a very high enclave creation, rate and caused excessive CPU consumption in the xxxxDIST Address Space.
- APAR [PH34378](#) (PTF UI74139) was created to address the above issue, and the provided solution, made available in February 2021, was to reuse the same enclave for up 200 individual transactions. This is how things work nowadays if you have Db2 v12 with PTF UI74139 installed.
- With such solution the transaction counts and response times collected by WLM are wrong. As a result IBM recommended to change all the response time based service classes used by DDF to velocity goals.

Db2 APAR PH34378 – HOLD(ACTION) ?

PH34378: EXCESSIVE CPU UTILIZATION DURING HIGH CONCURRENT DBAT TRANSACTION PROCESSING

A fix is available

[Obtain the fix for this APAR.](#)

```
***** ATTENTION *****  
With this change, the effective response time of the enclave  
will be longer than the response time to process a single  
transaction when using high performance DBATs.  
It is highly recommended that the WLM service classes, used  
to manage the processing of the client's connection when  
using high performance DBATs, be changed to use velocity  
goals. Existing WLM service classes utilizing response time  
or percentile response time goals will generally not meet  
their performance objectives.  
***** ATTENTION *****
```

- In March 2022 the functionalities of the Enclave Delete services (IWM4EDEL) were extended to allow Db2 to report back to WLM the total number of executed transactions together with the total execution time.
- The new function was made available by APAR [OA61811](#) with the associated PTFs for z/OS 2.3 and up - UJ07942 / UJ07943 / UJ07944, and apparently it is used for reporting purposes only.
- Db2 v12 APAR [PH41024](#) (PTF UI77711) exploits the new WLM function, but as far as we understand if you are running High Performance DBATs **the original recommendation to switch to velocity goals is still valid.**

Do I use High Performance DBATs ?

Ask your Db2 colleagues, but, in order to use High Performance DBATs you must:

- Set the CMTSTAT Db2 subsystem parameter to INACTIVE. This is the default.
- Bind the packages used by the JDBC driver as RELEASE(DEALLOCATE).
- Run DDF with the PKGREL option set to BNDOPT or BNDPOOL. Can be checked with a -DIS DDF DETAIL command
- If you are at Db2 v12, use High Performance DBATs, have the fix for APAR [PH34378](#) installed, AND make use of response time goals for DDF, then **you should check your DDF Service Classes achievements and consider switching them to single period with a velocity goal.**
- To get back to proper WLM reporting about DDF transaction counts and response times **make sure to install the PTFs** for WLM APAR [OA61811](#) and Db2 v12 APAR [PH41024](#).

End

Your feedback is important!

Submit a session evaluation for each session you attend:

www.share.org/evaluation

