Watson & Walker ZedRoadshow

Mario Bezzi, Frank Kyne
Watson & Walker

November 2021

Session 6AW

GSE UK Conference 2021 Charity Raffle

- The GSE UK Region team hope that you find this presentation and others that follow useful and help to expand your knowledge of z Systems.
- Please consider showing your appreciation by kindly donating to our charities this year, Royal National Lifeboat Institution (RNLI) & Guide Dogs for the Blind. Then follow the link on your receipt to enter your receipt number & amount donated into the GSE Raffle. You will get a raffle entry for every pound donated.
- Follow the link below or scan the QR Code:

http://uk.virginmoneygiving.com/GuideShareEuropeUKRegion

Supporting

Welcome

Who are we?

- Watson & Walker founded in 1988 by Cheryl Watson & Tom Walker.
- Publisher of *Cheryl Watson's Tuning Letter* and *CPU Charts* since 1991. If you haven't seen our *Tuning Letter*, please send us an email for a sample copy.
- In addition to our publications, our small team with over two hundred years of IT experience provides consulting on performance, capacity and upgrade planning, sysplex, and software pricing (such as TFP evaluations).
  - In photo above, see Alan Murphy (Dublin), Frank Kyne (President, New Hampshire), Tom Walker and Cheryl Watson (the short one, both from Florida), and Mario Bezzi (Milan). Not shown – Graham Horn (Johannesburg, SA).

Agenda

Mario:

- A VERY quick introduction to z/OS Language Environment
- IBM's SMFpy
- Reminder – Our free tools for the z/OS community

Frank:

- RMF Changes in z/OS 2.5
- WLM Service Unit Coefficient Changes in z/OS 2.5
- Db2 and z15 Sort Accelerator
- Recent XCF and CF-related enhancements
- Various Helpful websites
- Tidbits:
  - z/OSMF
  - COBOL

Mario's 'half'

A VERY quick introduction to z/OS Language Environment

IBM's SMFpy

Reminder – Our free tools for the z/OS community

A VERY quick introduction to LE

What is Language Environment

Common Runtime Library compilers rely on to provide functionalities to applications

- Defines a consistent, language-agnostic management model for storage, programs, error conditions etc.
- Provides languages with an opaque interface to underlying operating system services and hardware.
- LE exploited by Assembler (optionally), Fortran, PL/I, COBOL, C/C++, Java, Node.JS, Python, GO . . .
- Is the foundation for seamless Inter-Language-Communications (ILC).
- Generally a Language Environment (environment) is initialized for each new main program's execution.

Language Environment Run-Time Options

- Run-Time Options allow to specify how Language Environment should behave when an application runs.
- Defined via CEEPRMxx parmlib member, SETCEE command, CEEOPTS ddname and more. [1]
- Major areas of customization :

  Storage management, AMODE management, Error handling.

[1] https://www.ibm.com/docs/en/zos/2.1.0?topic=options-methods-available-specifying-runtime

ABC of LE program and storage mgmt models

Program management

- Application programs runs under LE threads. Within a thread execution is serial.
- One or more LE threads are associated to an enclave.
- Resources (storage, programs ..) are owned by the enclave and shared with daughter threads.
- Multi-threading under an enclave is not common among traditional applications.

Storage Management

- LE manages two types of storage for use by applications:
- HEAPs – used for COBOL WORKING-STORAGE, C malloc, and PL/I ALLOCATE requests.
- STACKs – used for save areas plus COBOL LOCAL-STORAGE , C and PL/I automatic variables.
- Initial storage is obtained with a single GETMAIN and managed internally by LE.

Important Runtime Options – HEAP/HEAP64

HEAP

HEAP controls the allocation of the initial heap, controls allocation of additional heaps created with the CEECRHP callable service, and specifies how that storage is managed.

Heaps are storage areas where you allocate memory for user-controlled dynamically allocated variables such as:

- C variables allocated as a result of the `malloc()`, `calloc()`, and `realloc()` functions
- COBOL WORKING-STORAGE data items
- PL/I variables with the storage class CONTROLLED, or the storage class BASED

Performance considerations

To improve performance, use the storage report numbers that are generated by the RPTSTG runtime option as an aid in setting the initial and increment size for HEAP.

Important Runtime Options – STORAGE

STORAGE

STORAGE controls the initial content of storage when allocated and freed. It also controls the amount of storage that is reserved for the out-of-storage condition. If you specify one of the parameters in the STORAGE runtime option, all allocated storage processed by that parameter is initialized to the specified value. Otherwise, it is left uninitialized.

You can use the STORAGE option to identify uninitialized application variables, or prevent the accidental use of previously freed storage. STORAGE is also useful in data security. For example, storage containing sensitive data can be cleared when it is freed.

Performance considerations

The use of the STORAGE runtime option to set values within heap and stack storage can have a negative impact on the performance of a Language Environment application.

IBM strongly recommends that you use STORAGE(NONE,NONE,NONE,0K) when you are not debugging, especially with any performance-critical applications. Do not set the STORAGE runtime option to values other than NONE at the system or region level if at all possible, because settings at those levels affect many more programs than necessary. Instead, use language mechanisms to ensure that automatic variables and data structures, including those contained within COBOL LOCAL-STORAGE and WORKING-STORAGE, are properly initialized.

Important Runtime Options – HEAPPOOLS / HEAPPOOLS64

HEAPPOOLS (C/C++ and Enterprise PL/I only)

Derivation: HEAP storage POOLS

The HEAPPOOLS runtime option is used to control an optional heap storage management algorithm known as *heap pools*. This algorithm is designed to improve performance of multithreaded C/C++ applications with high usage of `malloc()`, `__malloc31()`, `calloc()`, `realloc()`, `free()`, `new()`, and `delete()`. When active, heap pools can eliminate contention for heap storage.

Note:

HEAPPOOLS/HEAPPOOLS64 are disabled by default. Using them may benefit multi-threaded applications

A major example:

- AT-TLS (TCP/IP Application Transparent Transport Layer Security) exploits z/OS System SSL.
- AT-TLS is a 64 bit LE application which makes use of multi-threading to support encryption parallelism.
- Using HEAPPOOLS64 for AT-TLS can significantly improve networking latency

See – CICS APAR PH30197: ACTIVATE HEAPPOOLS64 OPTION FOR TLS PROCESSING

Important Runtime Options – RPTSTG

RPTSTG

Derivation: RePorT ST or aGe

RPTSTG generates, after an application has run, a report of the storage the application used. Language Environment writes storage reports only in mixed-case US English.

Use the storage report information to help you set the ANYHEAP, BELOWHEAP, HEAP, HEAP64, HEAPPOOLS, HEAPPOOLS64, IOHEAP64, LIBHEAP64, LIBSTACK, STACK, STACK64, THREADHEAP, THREADSTACK, and THREADSTACK64 runtime reports for the best storage tuning.

Performance considerations

This option increases the time it takes for an application to run. Therefore, use it only as an aid to application development.

The storage report generated by RPTSTG(ON) shows the number of system-level calls to obtain storage that were required while the application was running. To improve performance, use the storage report numbers generated by the RPTSTG option as an aid in setting the initial and increment size for stack and heap. This reduces the number of times that the Language Environment storage manager makes requests to acquire storage. For example, you can use the storage report numbers to set appropriate values in the HEAP *init_size* and *incr_size* fields for allocating storage.

AMODE 31 and AMODE 64 interoperability

PH28966: NEW FUNCTION

A fix is available

Obtain the fix for this APAR.

APAR status

Closed as new function.

Error description

New Function

Local fix

Problem summary

```
**************************************************************
* USERS AFFECTED:                                            *
*  z/OS Language Environment users who need AMODE 31 and AMODE *
* 64 programs interoperability support.                      *
**************************************************************
* PROBLEM DESCRIPTION:                                       *
* APAR PH28966 provides AMODE 31 and AMODE 64 programs       *
* interoperability support. The support provides the         *
* capability to run AMODE 31 programs and AMODE 64 programs  *
* together in an application in the same address space.      *
**************************************************************
* RECOMMENDATION:                                            *
**************************************************************
See problem description.
```

Document Information

Software version:
7B0

Operating system(s):
z/OS

Document number:
6462449

Modified date:
31 July 2021

Subscribe

You can track all active APARs for this component.

Notify me when an APAR for this component changes.

IBM z/OS V2.5: Enabling innovative development to support hybrid cloud and AI business applications

Table of contents

z/OS V2.5 supports the scale and simultaneous deployment of agile business use cases for hybrid cloud and AI capabilities and delivers the following values, features, and capabilities to help organizations succeed in their modernization efforts:

- Enterprise modernization with more seamless COBOL-Java interoperability. This gives application developers full application transparency by extending application programming models.

SMFpy

SMFpy – Introduction

- IBM just published a z/OS Statement of Direction about new capabilities to help clients analyse SMF data in an easy and modern way.

- Given our interest about SMF past, present and future, IBM development kindly offered us to be sponsor users to provide feedback.

- We thought it would make sense for us to SHARE™ with you what we learned so far.

Applying Artificial Intelligence to IT Operations

- IT infrastructures' complexity ever increasing.

- Larger and larger amounts of data to be timely processed to spot issues/saving opportunities

.

- Tuning and Problem Determination skills in shortage, especially in the mainframe environment. At the same time data analysis is emerging as a discipline per se.

- All of the above requires moving from an analysis approach based on the knowledge of the specific subject to one based on statistics and pattern recognition which can be used by non IT people, or automated.

SMFpy SoD in z/OS 2.5 announcement

- It is in light of the above scenario that in the z/OS 2.5 announcement IBM made the following Statement of Direction (can you read the fine print ?):

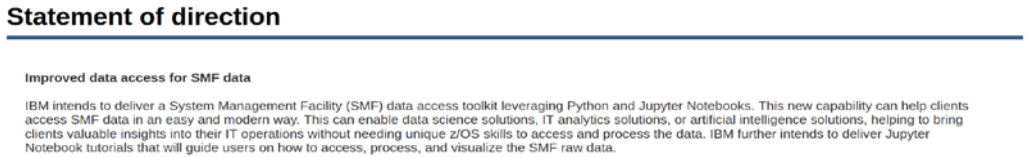

Statement of direction

Improved data access for SMF data

IBM intends to deliver a System Management Facility (SMF) data access toolkit leveraging Python and Jupyter Notebooks. This new capability can help clients access SMF data in an easy and modern way. This can enable data science solutions, IT analytics solutions, or artificial intelligence solutions, helping to bring clients valuable insights into their IT operations without needing unique z/OS skills to access and process the data. IBM further intends to deliver Jupyter Notebook tutorials that will guide users on how to access, process, and visualize the SMF raw data.

Improved data access for SMF data

IBM intends to deliver a System Management Facility (SMF) data access toolkit leveraging Python and Jupyter Notebooks. This new capability can help clients access SMF data in an easy and modern way. This can enable data science solutions, IT analytics solutions, or artificial intelligence solutions, helping to bring clients valuable insights into their IT operations without needing unique z/OS skills to access and process the data. IBM further intends to deliver Jupyter Notebook tutorials that will guide users on how to access, process, and visualize the SMF raw data.

https://www-01.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/0/897/ENUS221-260/index.html&lang=en&request_locale=en

Why Python

ZDNet

MUST READ: Supercomputing can help address blockchain's biggest problem. Here's how

Programming languages: Python could soon overtake C and Java as most popular

Python is battling for pole position, but Rust and TypeScript have made notable gains in popularity over the past year.

By Liam Tung | July 6, 2021 -- 11:16 GMT (12:16 BST) | Topic: Developer

C and Java remain the most popular languages in the Tiobe community index, but Python is stalking them and will likely take top spot in the future.

According to Tiobe's July 2021 index, the three most popular programming languages are C, Java and Python.

While the order hasn't change, Tiobe CEO Pau Jansen notes that the difference in apparent popularity is remarkably small, with just 0.67% between C and Python.

"This means that the next few months will be exciting. What language is going to win this battle? Python seems to have the best chances to become number 1, thanks to its market leadership in the booming field of data mining and artificial intelligence," Jansen noted.

https://www.zdnet.com/article/programming-languages-python-could-soon-overtake-c-and-java-as-most-popular/

Python – Batteries included!

Not logged

Article Talk Read Edit View hist

WIKIPEDIA
The Free Encyclopedia

Main page
Contents
Current events
Random article
About Wikipedia
Contact us

Python (programming language)

From Wikipedia, the free encyclopedia

For other uses, see Python (disambiguation).

Python is an interpreted high-level general-purpose programming language. Python's design philosophy emphasizes code readability with its notable use of significant indentation. Its language constructs as well as its object-oriented approach aim to help programmers write clear, logical code for small and large-scale projects.[30]

Google

python batteries included

All Images Shopping Videos News More Tools

About 9,640,000 results (0.53 seconds)

Motto of the Python programming language, meaning it comes with a large library of useful modules. Batteries Included (slang), in a product usability (mostly in software) it states that the **product comes together with all possible parts required** for full usability.

https://en.wikipedia.org › wiki › Batteries_Included

Batteries Included - Wikipedia

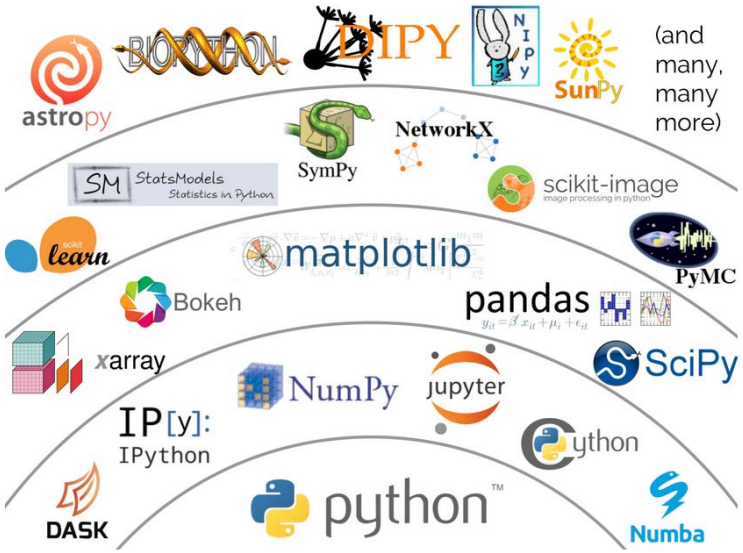

Credit: Jake vanderPlas, "The Unexpected Effectiveness of Python in Science", PyCon 2017

Jupiter Notebook

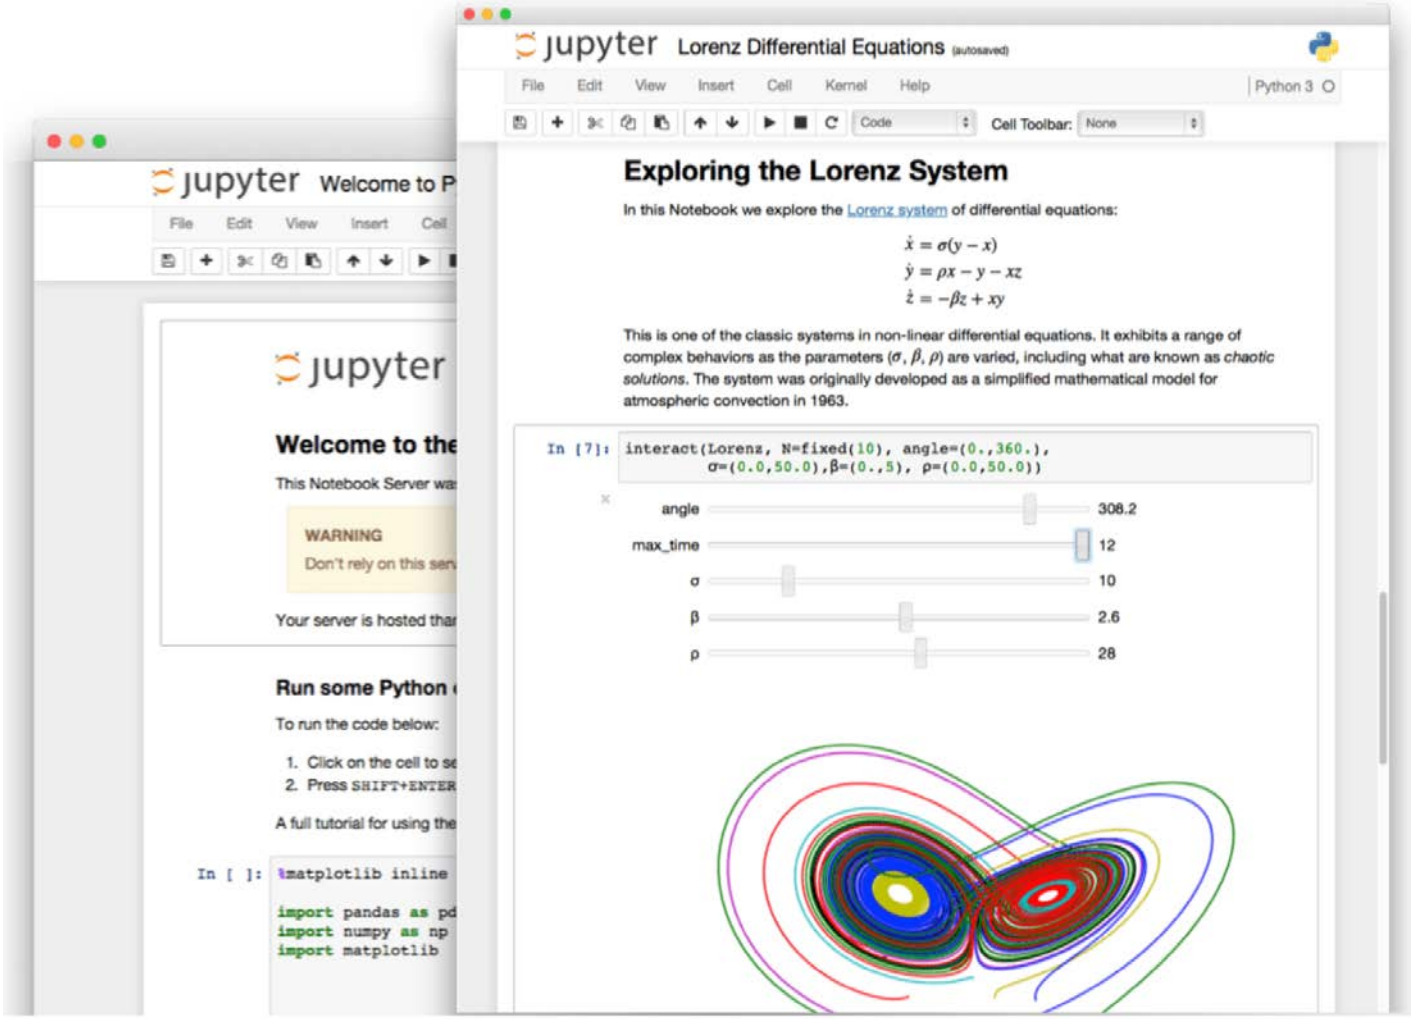

The Jupyter Notebook

The Jupyter Notebook is an open-source web application that allows you to create and share documents that contain live code, equations, visualizations and narrative text. Uses include: data cleaning and transformation, numerical simulation, statistical modeling, data visualization, machine learning, and much more.

Try it in your browser | Install the Notebook

https://jupyter.org/

SMFpy Value Proposition

SMFpy – Toolkit for modern access to z/OS Systems Management Data

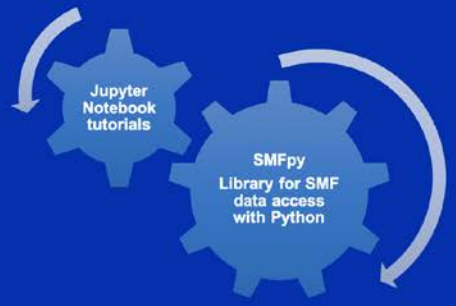

The components

SMFpy

A library for accessing Systems Management Facility (SMF) data with the help of Python

Jupyter Notebook tutorials

Guide how to fetch, process, and visualize SMF raw data to enable graphical analysis

Accessing SMF records with a modern programming language

Making data processing transparent

No z/OS skill needed to access and process data

Quick prototyping

Positioning SMFpy

- There are several other tools / products playing in this area. Most of them are part of larger, turnkey solutions.

- SMFpy is different, as it simply allows to consume SMF data using popular, state of the art technologies like Jupyter Notebooks and Python to extract valuable insights which can be leveraged for AI.

- With SMFpy YOU have the responsibility, and the power, to leverage your Python coding skills to easily build analytics assets addressing your specific needs.

SMF data science with SMFpy

[20]:
```
import smfpy
import plotly.express as px
from smfpy.fields import SMF99S1
from smfpy import names

ctx = smfpy.new_context("WLM.DATA.SMF")

df = ctx.samples.p_utilization().run()

plot = px.line(df,x=names(SMF99S1.timestamp), y=names(SMF99S1.cp_util, SMF99S1.ziip_util, SMF99S1.total_util),
               title="System Utilisation",
               labels={"value": "%","variable": "Processor Type", names(SMF99S1.timestamp): "Time" })

display(plot)
```

Summary

- SMFpy just appeared on the horizon as an interesting Statement of Direction.
- As far as we understand it, it is a powerful toolkit, not a packaged solution.
- Its purpose in life is to make it easy for non Mainframers and novice administrators to apply data science to SMF data using the tools they know and love.
- IBM kindly offered us to be sponsor users to provide feedback, we will keep you posted on our experience with it!

Watson and Walker free tools

Free?!? Are you kidding me?

Our Important Messages Health Check

- IBM recommends passing certain system messages to your alerting infrastructure so that the right people can be informed. However, identifying the specific messages that require this treatment, and using an available local automation package to detect and handle them is the customer's responsibility.
- The Watson and Walker IMPORTANT_MESSAGES check tracks specific system messages, raising a Health Checker exception if these messages are detected.
- It is provided with a default list of critical messages we build from Redbooks, manuals, APARs, customers' experiences. Users can extend this list with their own groups of environment-specific messages.
- We plan to periodically deliver updated versions of the important messages list, based on our continuous analysis of new software functions and on the feedback of users of the check.

200+ users downloaded it!

Our Important Messages Health Check(s)

- More than 150 individual messages identified for z/OS plus the ones tagged as critical by z/OS.

- Some 50 individual messages identified for Db2 (thanks to IBM's Adrian Burke).

```
  Display  Filter  View  Print  Options  Search  Help
-------------------------------------------------------------------------------
SDSF HEALTH CHECKER DISPLAY  S0W1                          LINE 66-84 (212)
COMMAND INPUT ===>                                            SCROLL ===>
NP    NAME                                CheckOwner       State
      IMPORTANT_MESSAGES_CUSTOM           WATSONWALKER     ACTIVE(ENABLED)
      IMPORTANT_MESSAGES_SYS_CRIT         WATSONWALKER     ACTIVE(ENABLED)
      IMPORTANT_MESSAGES_WWD21911         WATSONWALKER     ACTIVE(ENABLED)
      IMPORTANT_MESSAGES_WWIM1905         WATSONWALKER     ACTIVE(ENABLED)
```

- You can easily add your own.

Important Messages Health Check in action – I

- Let us see what happens if an important messages is generated..

```
00000281  IEF126I IBMUSER - LOGGED OFF - TIME=10.38.31
00000281  $HASP395 IBMUSER  ENDED - RC=0000
00000281  $HASP250 IBMUSER PURGED -- (JOB KEY WAS D6899D49)
00000281  IEA989I SLIP TRAP ID=X33E MATCHED.  JOBNAME=*UNAVAIL, ASID=0020.
00000090  *IGW048A IMF LRU Stalled SMSPDSE 876
00000090  last completed:08/07/2019 11:15:57.029629
00000090  BUFFER SPACE USED:16MB
00000090  HZS0001I CHECK(WATSONWALKER,IMPORTANT_MESSAGES_SYS_CRIT): 877
00000090  WWIMCK9E In the last monitoring interval, System S0W1 issued 1
00000090  important message(s) of those tracked by message monitoring group
00000090  SYS_CRIT.
00000290  LOGON
00000281  $HASP100 IBMUSER  ON TSOINRDR
00000090  $HASP373 IBMUSER  STARTED
00000281  IEF125I IBMUSER - LOGGED ON - TIME=12.04.42
```

1. A IGW048A message is generated because of an issue with SMSPDSE
2. The Important Message Check traps it and issues an Health Checker Exception

The Exception report – I

```
  Display  Filter  View  Print  Options  Search  Help
 ------------------------------------------------------------------------------
 SDSF OUTPUT DISPLAY IMPORTANT_MESSAGES_SYS_CRIT   LINE 0       COLUMNS 02- 81
 COMMAND INPUT ===>                                           SCROLL ===> CSR
********************************** TOP OF DATA ***********************************
CHECK(WATSONWALKER,IMPORTANT_MESSAGES_SYS_CRIT)
SYSPLEX:    ADCDPL    SYSTEM: S0W1
START TIME: 08/07/2019 12:04:33.985545
CHECK DATE: 20190402  CHECK SEVERITY: LOW
VERBOSE MODE: YES
CHECK PARM: WARN(ANY)


* Low Severity Exception *

WWIMCK9E In the last monitoring interval, System S0W1 issued 1 important
message(s) of those tracked by message monitoring group SYS_CRIT.

  Explanation:  CHECK(IMPORTANT_MESSAGES_SYS_CRIT) detected that since
    when last monitoring interval started (08/07/2019 11:15:58.82),
    system S0W1 issued 1 important message(s) of those tracked by
    message monitoring group SYS_CRIT.

    1 of these messages have been issued by system S0W1 since when the
    Important Messages Health Check was initialized (08/07/2019
    04:28:06.93).

    In total 1 of these messages have been issued by system S0W1 since
    when message monitoring was initialized (08/07/2019 04:28:06.95).

    Note that message monitoring is only active while at least one
    instance of the IMPORTANT_MESSAGES check is active.
```

Continues on next slide

The Exception report – II

```
  Display  Filter  View  Print  Options  Search  Help
 ------------------------------------------------------------------------------
 SDSF OUTPUT DISPLAY IMPORTANT_MESSAGES_SYS_CRIT   LINE 51       COLUMNS 02- 81
 COMMAND INPUT ===>                                            SCROLL ===> CSR
  Check Reason:  Health Check for msgs with desc code 1, 2, 3 or 11

WWIMCK3I The Check ran in verbose mode. Following a list of all the 1
occurrence(s) of messages in group SYS_CRIT since when message
monitoring was initialized (08/07/2019 04:28:06.95).

------------------------------------------------------------------------

MSG# 1 - 08/07/2019 12:04:33.62 -

*IGW048A IMF LRU Stalled SMSPDSE
 last completed:08/07/2019 11:15:57.029629
 BUFFER SPACE USED:16MB

------------------------------------------------------------------------


END TIME: 08/07/2019 12:06:34.412597  STATUS: EXCEPTION-LOW
****************************** BOTTOM OF DATA ******************************
```

Our WWUNTERSE utility

IBM proprietary TERSE compression format is still popular among z/OS customers, although the availability of industry standard zlib and zEDC may change this.

WWUNTERSE allows to uncompress on distributed systems, F, FB, V, VB, VS, and VBS sequential data sets created on z/OS using AMATERSE or TRSMAIN on z/OS. It is based on IBM's tersedecompress, the open source port to Java of the uncompress function of terse.

It supports the inclusion of the Block Descriptor Word (BDW) in the output file, meaning that the output files can used by products expecting them.

Because of its support for tersed VB and VBS files, we expect this tool to be particularly of interest to people looking to efficiently move SMF data to distributed platforms.

WWUNTERSE is available for x86-64-bit Windows and x86-64-bit Linux environments.

150+ users downloaded it!

Running WWUNTERSE

```
mario:~/Downloads/Test\ wwunterse -h
** wwunterse  terse decode unpack v1.3.0. Compiled on: Mar  2 2021 17:30:34   **
** wwunterse  derived from github.com/openmainframeproject/tersedecompress    **
**            wwunterse is a free program provided by Watson and Walker inc.  **
**            by using it you agree with https://watsonwalker.com/disclaimer/ **

**            usage: wwunterse infile [-v] [outfile]
**                   only provide infile to show input file characteristics
**                   use the -v option to force RECFM=V output for RECFM=VB input files
```

Running WWUNTERSE – continued

```
mario:~/Downloads/Test\ wwunterse -h
** wwunterse  terse decode unpack v1.3.0. Compiled on: Mar  2 2021 17:30:34   **
** wwunterse  derived from github.com/openmainframeproject/tersedecompress    **
**            wwunterse is a free program provided by Watson and Walker inc.  **
**            by using it you agree with https://watsonwalker.com/disclaimer/ **

**            usage: wwunterse infile [-v] [outfile]
**                   only provide infile to show input file characteristics
**                   use the -v option to force RECFM=V output for RECFM=VB input files
```

```
mario:~/Downloads/Test\ wwunterse ./SMF101.D201030.TRS ./SMF101.D201030.SMF
```

Running WWUNTERSE – continued

```
mario:~/Downloads/Test\ wwunterse -h
** wwunterse  terse decode unpack v1.3.0. Compiled on: Mar  2 2021 17:30:34   **
** wwunterse  derived from github.com/openmainframeproject/tersedecompress    **
**            wwunterse is a free program provided by Watson and Walker inc.  **
**            by using it you agree with https://watsonwalker.com/disclaimer/ **

**            usage: wwunterse infile [-v] [outfile]
**                   only provide infile to show input file characteristics
**                   use the -v option to force RECFM=V output for RECFM=VB input files
```

```
mario:~/Downloads/Test\ wwunterse ./SMF101.D201030.TRS ./SMF101.D201030.SMF
** wwunterse  derived from github.com/openmainframeproject/tersedecompress    **
**            wwunterse is a free program provided by Watson and Walker inc.  **
**            by using it you agree with https://watsonwalker.com/disclaimer/ **

** 08/11/2021 10:40:17 starting terse decode unpack v1.3.0
** wwunterse  input  file "./SMF101.D201030.TRS"
** wwunterse  input  recfm=VBS blksize=32760 lrecl=32763 packtype=PACK terse version flags=0x02.
** wwunterse  output file "./SMF101.D201030.SMF" will have RDWs and BDWs.
```

Running WWUNTERSE – continued

```
mario:~/Downloads/Test\ wwunterse -h
** wwunterse  terse decode unpack v1.3.0. Compiled on: Mar  2 2021 17:30:34   **
** wwunterse  derived from github.com/openmainframeproject/tersedecompress    **
**            wwunterse is a free program provided by Watson and Walker inc.  **
**            by using it you agree with https://watsonwalker.com/disclaimer/ **

**            usage: wwunterse infile [-v] [outfile]
**                   only provide infile to show input file characteristics
**                   use the -v option to force RECFM=V output for RECFM=VB input files
```

```
mario:~/Downloads/Test\ wwunterse ./SMF101.D201030.TRS ./SMF101.D201030.SMF
** wwunterse  derived from github.com/openmainframeproject/tersedecompress    **
**            wwunterse is a free program provided by Watson and Walker inc.  **
**            by using it you agree with https://watsonwalker.com/disclaimer/ **

** 08/11/2021 10:40:17 starting terse decode unpack v1.3.0
** wwunterse  input  file "./SMF101.D201030.TRS"
** wwunterse  input  recfm=VBS blksize=32760 lrecl=32763 packtype=PACK terse version flags=0x02.
** wwunterse  output file "./SMF101.D201030.SMF" will have RDWs and BDWs.
** 08/11/2021 10:42:39 terse complete decode unpack
** wwunterse  input file size: 1013035884 bytes output file size: 8426656834 bytes compression ratio:  12%
```

My little PC untersed 8GB of SMF data in 2 minutes!

Where to get our free tools

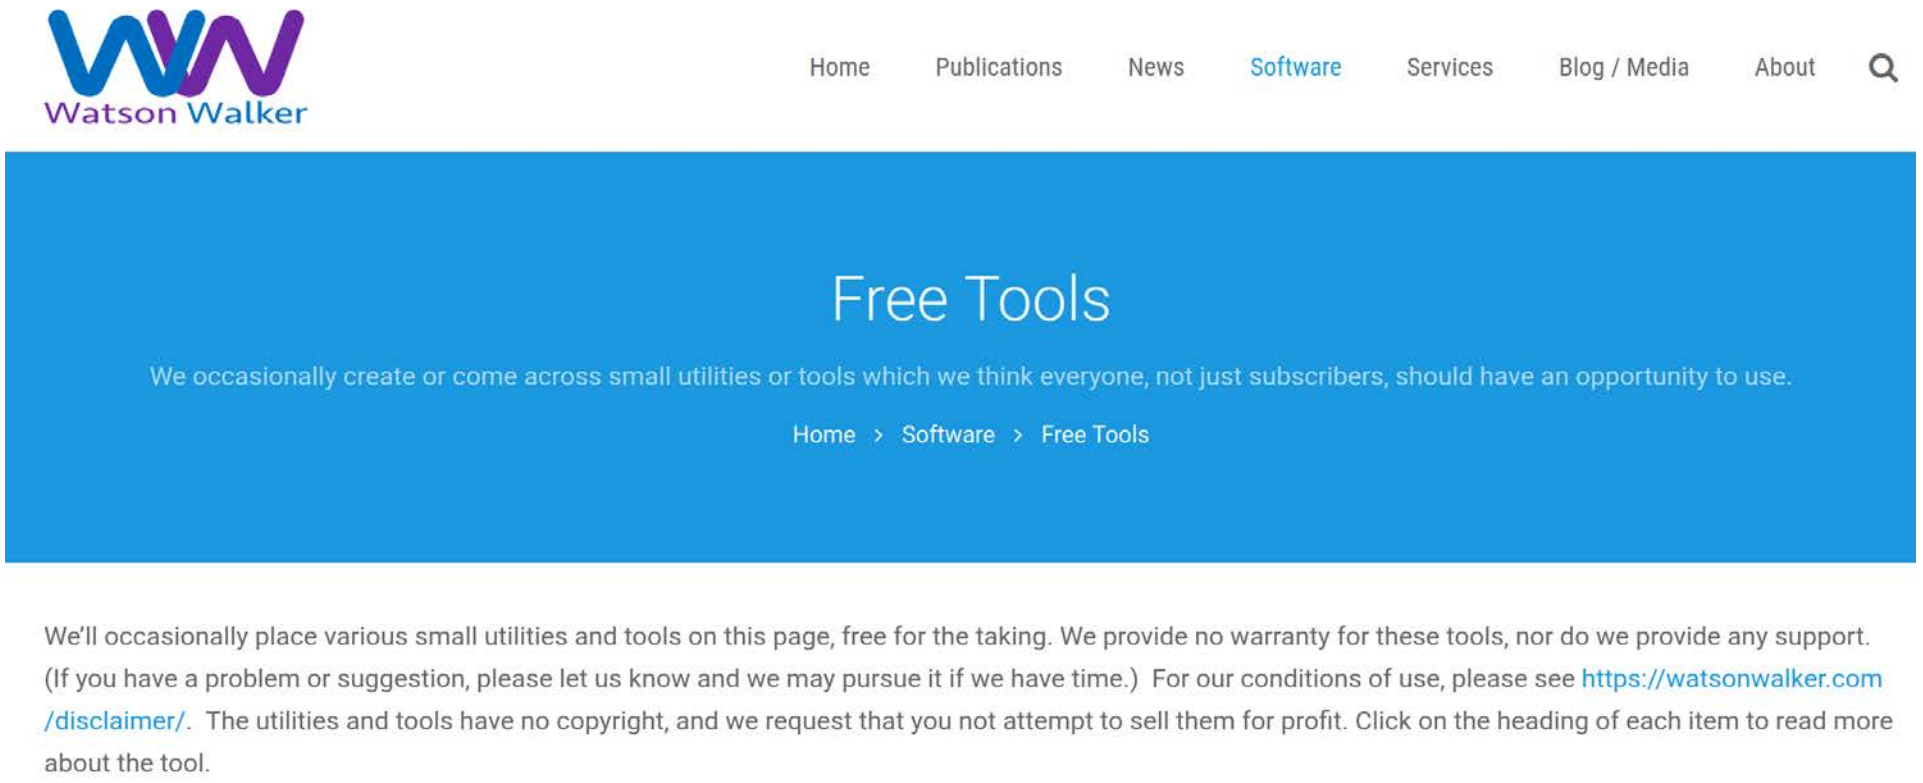

https://watsonwalker.com/software/free-tools/

Frank's 'half'

RMF Changes in z/OS 2.5

WLM Service Unit Coefficient Changes in z/OS 2.5

Db2 and z15 Sort Accelerator

Recent XCF and CF-related enhancements

Tidbits:

- User catalog extents
- COBOL
- z/OSMF

RMF 'Restructure'

- The z/OS 2.5 announcement includes information about a restructuring of RMF into two z/OS features:
  - ADG (Advanced Data Gatherer) – this consists of the part of the traditional RMF product that *gathers* performance information.
  - RMF – this consists of the *reporting* parts of the traditional RMF product – this is now managed by Rocket Software.
- **If you are an existing RMF customer, there is no net change** – licensing the RMF feature entitles you to use ADG as well, and price is unchanged compared to z/OS 2.4.
  - There *are* some small technical changes, to accommodate new SGRBaaaa ADG data sets.
  - Also, note that in 2.5 there are **2** bookshelves – z/OS Data Gatherer, and z/OS RMF.

RMF 'Restructure'

- Because there are *so many* parts to traditional RMF, some of which combined data gathering and data reporting (RMF Mon III, for example), you should refer to **Peter Muench's** excellent *What's New in RMF Data Gatherer* presentation (not available yet, but *will be* in a new zOS-DataGatherer directory on the IBM Github repository https://github.com/IBM/IBM-Z-zOS)
  - **Marna Walle's** *Upgrade to z/OS 2.5 Technical Actions: Part 2 of 2* (session 7AP) includes helpful info about IFAPRD and LNKLST considerations.
- For Tuning Letter subscribers, you can find more information in What RMF's Restructure Means To You, including a helpful matrix that shows which functions are in which product.

Preparing for WLM SDC Changes in z/OS 2.5

The IBM z/OS 2.4 announcement included a Statement of Direction that:

- In z/OS 2.5, WLM will no longer allow customer override of service definition coefficients.
- The *default* values will change from CPU=10, SRB=10, IOC=5, MSO=0 to CPU=1, SRB=1, IOC=0, MSO=0.
  - The IBM *recommended* values have been CPU=1, SRB=1, IOC=0.5 (or, more recently, 0), and MSO=0.
- If you have been using the IBM recommended values *in all your sysplexes*, this will probably have little or no effect on you.
- APAR OA59066 delivered a migration health check to warn you if your SDC values are other than the new z/OS 2.5 ones - ZOSMIGV2R4_NEXT_WLM_ServCoeff – make sure to activate this check.

Preparing for WLM SDC Changes in z/OS 2.5

How is this change likely to impact you?

- Changed SDC values come into effect, *sysplex-wide*, automatically as soon as the *first* z/OS 2.5 system is IPLed into the sysplex.
- Multi-period WLM Service Classes use 'weighted' service unit consumption to specify the DURATION of each period.
  - Note that WLM Resource Group definitions use **UN**-weighted service units, so no change.
  - Enqueue Recommendation Value (ERV) in IEAOPTxx is also defined in terms of UNweighted service units, so no change there either.
- Most chargeback processes use *weighted* service units.
- SMF type 30, 72.3, 79, 90.34, 99.1, 120.9, and 120.11 records report weighted service units. *Most* 'service unit' fields in 70.1 records are actually MSUs.

For more information, see excellent *How long is too long? Multi-period Service Classes*, session 6AV, by **John Baker**, and Tuning Letter article *WLM Service Definition Coefficient Changes* in Tuning Letter 2021 No. 2.

Db2 and Z15 Sort Accelerator (Z Sort)

- IBM measurements of DFSORT use of Z Sort show impressive CPU and elapsed time savings – up to 40%.
- However, there have been a number of HIPER DFSORT Z Sort-related APARs which have delayed widespread production rollout.
  - There is a FIXCAT called `IBM.Function.Zsort`, however as of Nov 8, 2021, it is not assigned to any APARs. Hopefully we can get that fixed.
  - In the meantime, the following are all the Z Sort-related HIPER DFSORT APARS: PH34177, PH13891, PH32037, PH33837, PH32571, PH34436*, PH35418*. PTFs are available for all of them.
- There *are* open DFSORT Z Sort-related APARs at the moment, but none of them are data integrity-related.
- **Terminology**: The Z15 Sort Accelerator function is accessed by programs (such as DFSORT, Db2, Syncsort, etc) using the SORTL instruction.

Db2 and Z Sort

There are a number of aspects to **Db2** exploitation of Sort Accelerator:

- Db2 12 *internal* exploitation of SORTL (for example, when an SQL statement specifies ORDER BY or GROUP BY) was added by Db2 APAR PH31684.
  - For this processing, Db2 does *not* use an external sort product – it uses a Db2 component called RDS (Relational Data Services), and RDS contains its own support for the Sort Accelerator function.
    - RDS decisions about whether to use SORTL are based on Db2 measurements, meaning that those decisions could change based on user experiences ➔ **monitor for new APARs.**
- There is 'no' way to disable this capability. If Db2 is running on a z15 and "the necessary conditions are met", RDS will decide if the current situation would be a good fit for Sort Accelerator (but see very interesting Db2 APAR PH36930 for more info).
  - The Db2 IFCID 2 and IFCID 96 records report on the use of this capability:
    - QXSTSRT # of times RDS Sort was performed
    - QXSTSRTL # of times RDS Sort used Z Sort (SORTL)

ICEPRMxx setting has NO impact on this

Db2 and Z Sort

- The *other* side of Db2 exploitation of Z Sort is use by Db2 *utilities*:
- Db2 APAR PH28183 added a Db2 ZPARM (UTILS_USE_ZSORT) that is used to control whether Db2 tools should try to use Z Sort (Note - this does NOT control Db2's own internal use of Z Sort).
- There are two cases with these utilities if UTILS_USE_ZSORT is specified:
  - If using DFSORT, Db2 will tell the utility that you would like it to use Z Sort, and the utility would then add the appropriate parm to the sort. Ultimately, *DFSORT* decides if this is a good place to use zSort.
    - The only utility I'm aware of that supports this so far is REORG.
  - If using the DB2 Sort for z/OS product (an alternative to DFSORT):
    - This currently this does NOT support Z Sort (CNK526E ... INTERNAL ABEND – 09).
    - DB2 Sort for z/OS APAR PH35059 stops it from trying to use Z Sort, even if UTILS_USE_ZSORT is specified in Db2 zPARMs.
    - Monitor FIXCAT `IBM.DB2.Db2Sort` for changes to this situation.

Precisely Syncsort and z Sort Accelerator support

- If you use Syncsort, contact Precisely to get the current situation regarding Syncsort support for the z Sort Accelerator. A recent webcast stated that the support would be delivered 'very soon'.
- You should also check the status of Db2 Sort for z/OS – Precisely claim that this is working, but Db2 APAR PH35059 forces sort calls to NOT use Z Sort.

Coupling Facility Monopolization Avoidance

- It is possible for a CF to be bombarded with so many requests for a single structure that the resulting queueing and latch contention in the CF impacts the service time for ALL requests to that CF.
- For example:

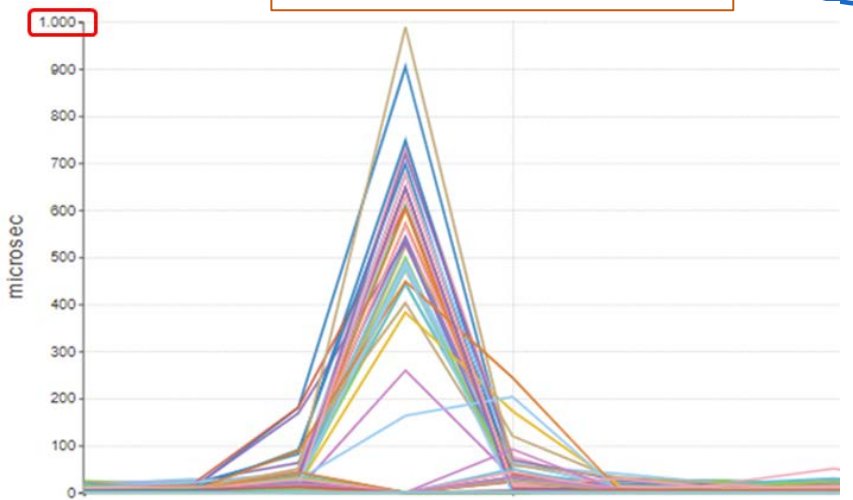

Coupling Facility Monopolization Avoidance

- To address this situation, a new function called CF Monopolization Avoidance was delivered for z/OS 2.2 to 2.4 by APAR OA56774 and CF Level 24 (z15 T01 or T02).
- The new function is controlled by the XCF FUNCTION CFMONOPAVOID.
  - This is DISABLED by default in z/OS 2.4 and prior releases.
  - It is ENABLED by default in z/OS 2.5.

Coupling Facility Monopolization Avoidance

- After enabling this FUNCTION, the previous interval looked like this instead:
- The spiking structure is the one that was monopolizing the CF. All other structures were unaffected.

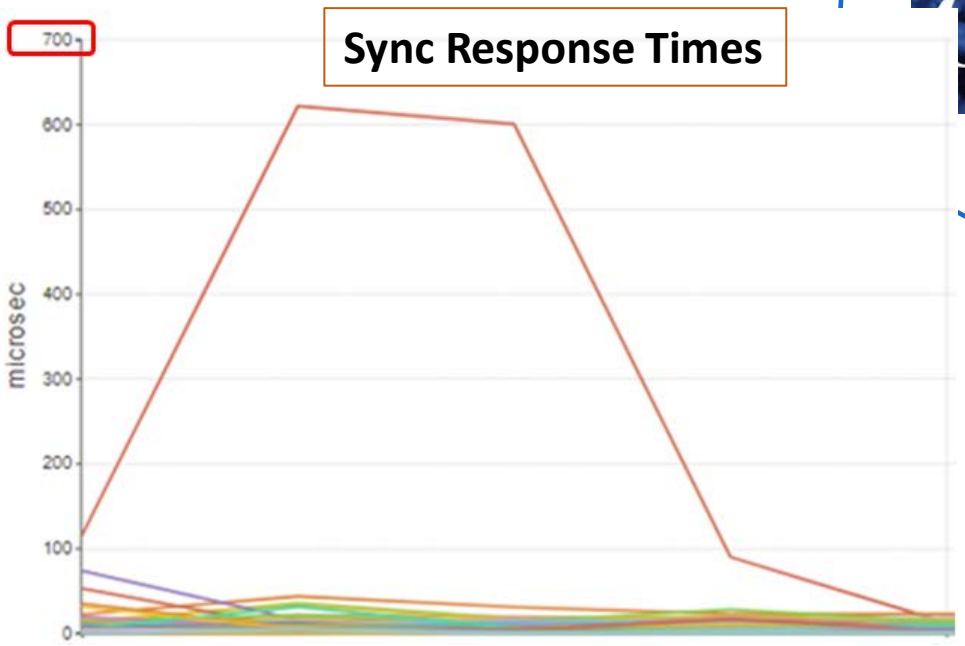

Coupling Facility Monopolization Avoidance

- For much more information, including a list of the messages that indicate this function is being invoked, refer to **Steve Warren's** *Parallel Sysplex Update*, Session 5AD.

- We highly recommend installing and enabling this function.

210915 Machine Alert for 8561, 8562

- The above alert provides information about a CFCC bug that can result in high CF CPU time for requests to Db2 Group Buffer Pool structures.
- Because there are so many structures in a Coupling Facility, it can be challenging to determine if this problem is affecting you (assuming that you even saw the alert).
- One of our clients was experiencing high CF CPU utilization, with one of the GBP structures being identified by RMF as using >80% of the used CF CPU time.

210915 Machine Alert for 8561, 8562

- To determine if that structure might be a victim of the CFCC bug, we calculated the CF CPU time per request for each interval over a day.
- It transpired that the CF CPU time per request in those peak intervals was up to 6x higher than in a normal interval, with a clear correlation between high request rates and high CPU time per request.
- We highly recommend that sites with busy GBP structures install z15 Bundle 48 or 49A before year end change freeze.

MQ CF CPU Usage

- With our new-found metric (CF CPU per request), we investigated other large CPU-consuming structures and found intermittent spikes in an MQ shared queue structure where the *average* CPU time was 12x that of 'normal' intervals.
  - If the periods of high CPU usage can be tied back to a particular application or job, it might be possible to alter the application to achieve the same net effect in a more CPU-efficient manner.
- We also recommend adding this metric as a Key Performance Indicator to help detect unusual activity or potential bugs, such as the one addressed by Bundle 48.

Something to watch out for

- I'm sure it is purely coincidental, but recently we have seen a number of user catalogs with 30 or more extents.
- Having a large number of extents isn't necessarily an issue. However, reorging a catalog can be very disruptive, so it is prudent to keep an eye on catalog usage.
- The secondary extents *might* be a lot smaller than the primary extent, so you *could* go from 1 extent to 50 extents very quickly.
- Use the F CATALOG,ALLOCATED command to display percent of maximum extents (123) for each allocated catalog – easy to add to automation.
- Might be able to slow growth by not defining additional HLQs in that catalog (*if* you know that it is getting larger than you would wish).
- See IBM Techdoc "ICF Catalog Management Recommendations & Guidelines"
- **Bonus** – z/OS 2.5 adds the ability to *dynamically* switch to a new Master Catalog.

Tidbits

- If you have z14 or z15, make sure you update CPENABLE to (5,15).
- If running z/OS 2.4, aim for XCF MAXMSG values not larger than 6000, and apply HIPER APAR OA60480 (UJ06017).
- Don’t forget that COBOL V4 compiler goes EOS on April 30, 2022.

Speaking of COBOL EoS…

- Have you seen the discussion about ‘Software Drag Racing’ on IBM-MAIN?
- We always assumed that COBOL is *way* faster than Java.
- Well, it turns out that that isn’t always the case – have a look at *Java vs C++ : Drag Racing on z/OS* paper by **Andrew Rowley**.

Speaking of COBOL EoS…

- This is especially interesting if you use Sub-cap CPCs (4xx, 5xx, 6xx), where the zIIPs are between 1.7 and 7x faster than the general purpose CPs.
  - Java uses zIIPs, COBOL does not.
  - Even if Java isn't faster for a given program, faster zIIPs make up for a *big* performance disadvantage.
  - AND, CPU time used on a zIIP doesn't count towards your software bill.

Speaking of COBOL EoS…

- To make this topic even more interesting, IBM recently introduced APARs and supporting documentation to improve inter-operability between COBOL and Java programs.
  - Requires APARs PH28966 (LE), PH37101 (LE COBOL), and PH40798 (Java 8 SR7 FP36 64-bit).
  - See Technote *An enhancement in COBOL/JNI interface*.
- For a mapping of Java service and fixpack levels to APAR numbers, refer to https://www.ibm.com/support/pages/apar/II14735.

And on the topic of Java

- IBM recently announced Java 11 for z/OS:
  - "IBM Semeru Runtime Certified Edition for z/OS, Version 11, formerly known as IBM 64-bit SDK for z/OS, Java Technology Edition, is certified with the Java Compatibility Kit as a fully compliant Java product."
  - Semeru will be available on Nov 19, 2021.
- Information about Java SDK products on z/OS, including links to download non-SMP/E version, is available on https://www.ibm.com/support/pages/java-sdk-products-zos

z/OSMF

- In most presentations and publications since 2010, Cheryl has been STRONGLY encouraging our customers to install and use z/OSMF. That's 11 years!
- Today's z/OSMF is (literally) unrecognizable from the 2010 version, with features that are ONLY available in z/OSMF, greatly improved performance, and you will soon need it to install a ServerPac.
- All of that is great. But THE best reason to start using z/OSMF is for the new CFRM Policy Editor in z/OS 2.5 (and, maybe?, rolled back to previous releases).
- Watch out for a dedicated article about the Policy Editor in an upcoming Tuning Letter. Email us if you would like to know more about this cool new function or would like to work with us on the article.

Thank you!

- For more information:
  - See our website at www.watsonwalker.com.
  - Send an email to technical@watsonwalker.com.
  - View our YouTube video about our new Application Profiler for Z.
  - Check out our Free Tools at www.watsonwalker.com/software/free-tools.
  - See over 110 free previous presentations at www.watsonwalker.com/publications/presentations/.
  - See our free SMF Reference Summary at www.watsonwalker.com/publications/#smf. We'll update it again soon now that z/OS 2.5 has gone GA.
  - Access Out-of-Print IBM Systems Center Books at www.watsonwalker.com/publications/ibm-systems-center-books/.

HOW Could you say NO to this face??

Please dig deep and help someone that is not as fortunate as you.

Become a member of GSE UK

- Company (880 Euro) or individual membership (388 Euro) available.
  - The 'Individual' membership also applies to vendors and consultants.
- Benefits include:
  - GSE Annual Conference: Receive 5 free places + 2 free places for trainees
  - 20% discount on fees for IBM Technical Conferences
  - 20% on IBM Training Courses in Europe
  - 15% discount for IBM STG Technical Conferences in the USA
  - 20% discount on the fee for taking the Mainframe Technology Professional (MTP) exams
  - European events – via GSE HQ
- Get more information at https://www.gse.org.uk/join/
- Contact membership@gse.org.uk for details

GSE GUIDE SHARE EUROPE UK REGION

Please submit your session feedback!

- Do it online at https://conferences.gse.org.uk/2021/feedback/6aw
- This session is 6AW

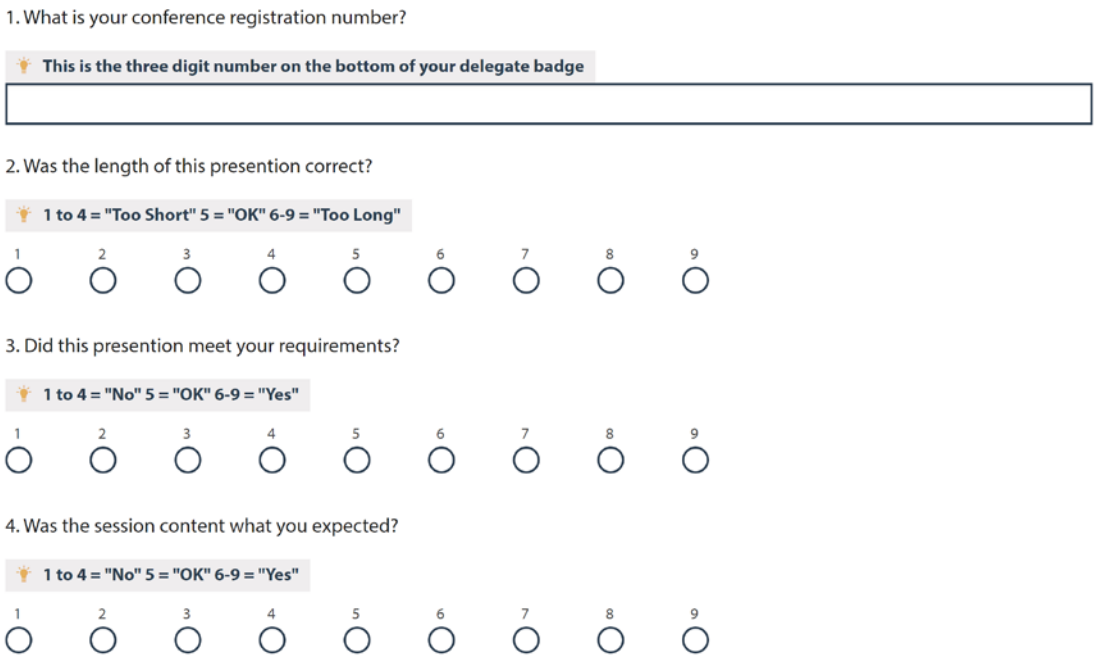

1. What is your conference registration number?

This is the three digit number on the bottom of your delegate badge

2. Was the length of this presention correct?

1 to 4 = "Too Short" 5 = "OK" 6-9 = "Too Long"

1 2 3 4 5 6 7 8 9

3. Did this presention meet your requirements?

1 to 4 = "No" 5 = "OK" 6-9 = "Yes"

1 2 3 4 5 6 7 8 9

4. Was the session content what you expected?

1 to 4 = "No" 5 = "OK" 6-9 = "Yes"

1 2 3 4 5 6 7 8 9

GSE GUIDE SHARE EUROPE UK REGION