

Watson & Walker zRoadshow Summer Virtual SHARE

Thursday, August 12, 2021

Mario Bezzi – Mario@watsonwalker.com

Frank Kyne – Frank@watsonwalker.com

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.
- We also have a small suite of software products, including one (['AP4Z'](#)) that we announced just this week, plus some free tools that can be downloaded from our [website](#).
- In addition to our publications and software, 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 (New Hampshire), Tom Walker and Cheryl Watson (the short one, both from Florida), and Mario Bezzi (Milan). Not shown – Graham Horn (Johannesburg, SA).

The clock is ticking

If you did NOT agree to sharing your information with the conference sponsors, you will only have access to the 'live' AND 'On-Demand' session recordings until the end of THIS MONTH (August 2021). If want to change your mind, contact registration@share.org.

If you DID agree to share your information, both sets of recordings will be available until the end of December.



EVENTS

PROCEEDINGS FILE LIBRARY

Access the archive of SHARE Proceedings to view content from past SHARE events, including past session presentations and handouts. The SHARE Proceedings library is a **SHARE membership benefit**, and you must be a **SHARE member** to view the full content.

Session recordings for past events are only available to those who registered for the event. Full event PDF proceedings are available to **ALL SHARE members**.

The clock is ticking

You might not have heard that z/OSMF Software Management is going to be the tool to drive *all* ServerPac installs from January 2022??



Whether you have heard or not, make sure to watch **Kurt Quackenbush's** [Installing z/OS V2.5 Using z/OSMF Software Management](#) session.

And while you eagerly await the opportunity to order your z/OS 2.5 ServerPac, why not take a test drive using a sample z/OSMF Portable Software Instance that you can download from [here](#).

- The z/OSMF Software Management install process is NOT just a prettier face on the old CustomPac dialogs, so you will need some time to get familiar with this new way.

Agenda

Frank:

- RMF Changes in z/OS 2.5
- WLM Service Unit Coefficient Changes in z/OS 2.5
- Db2 and z15 Sort Accelerator
- Various Helpful websites
- Tidbits
- z/OSMF
- COBOL

Mario:

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

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.

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](#) On-Demand session.

- [RMF: The Latest and Greatest](#) session by **Allison Moshier** and **Paul Stewart** focuses on the enhancements to the new RMF feature, and doesn’t say much about the restructure.
- **Marna Walle’s** [Upgrade to z/OS 2.5 Technical Actions: Part 2 of 2](#) session 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 Upcoming WLM SDC Changes

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, this will probably have little effect.
- 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 Upcoming WLM SDC Changes

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 defined in terms of weighted service units.
- 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 [Performance Infrastructure Improvements in z/OS V2.5 WLM](#), by **Andreas Henicke**, and Tuning Letter article [WLM Service Definition Coefficient Changes](#).

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%. See **John Burg's** [“The Future is Now With System Z New Technology Exploitation Using IBM Z Batch Network Analyzer \(zBNA\)”](#) session and **Ryan Bouchard's** [Making the Most of zSort](#) on-demand presentation for *much* more info.

- However, there have been a number of HIPER DFSORT Z Sort-related APARs which have delayed widespread production rollout. Make sure to monitor FIXCAT IBM.Function.Zsort to be aware of any important APARs or PTFs.
- 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 using the SORTL instruction.

Db2 and Z Sort

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

- 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.

IBM Docs Changes and Outstanding Items


IBM replaced IBM KnowledgeCenter (KC) with '[IBM Documentation](#)' website on April 1.

The new site initially had performance issues and regressed functions that were previously available in KC.

As a result, customers were quite vocal in airing their displeasure on IBM-Main.

Following that customer feedback, enhancements were made, the most visible of which were:

- Performance (with the exception of displaying the Table of Contents) was dramatically improved.
- Performing a search no longer requires you to click on the magnifying glass before you can enter your search term. *And*, the search string is retained.
- The “Change Version” button finally works!



IBM Docs Changes and Outstanding Items

But, there is still more work to be done to return to the level of functionality provided by the (little-loved) KnowledgeCenter:

- It takes up to 10 seconds for the ToC to appear, after the manual ‘page’ is displayed.
 - And the ToC doesn’t always reflect where you are within the book/library
- Doing a search within IBM Docs results in the manual name being displayed IF you initiate the search from the Docs z/OS 2.4 home page or from a 2.4 manual, but not otherwise:

Virtual storage

z/OS / 2.4.0 / z/OS DFSORT / z/OS DFSORT Tuning Guide

Logically, DFSORT (like any other application) works within a
 aults such as TMAXLIM and run-time options such as REGIO
 this address space. With the exception of dataspace sorting (

Virtual storage

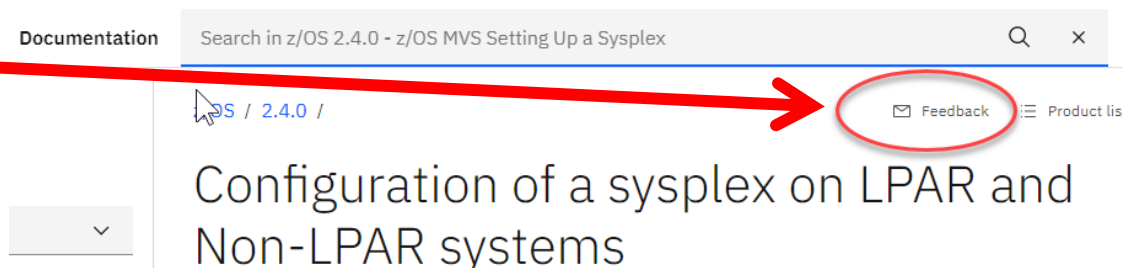
z/OS / 2.2.0

Logically, DFSORT (like any other applicati
 aults such as TMAXLIM and run-time optic
 this address space. With the exception of c

IBM Docs Changes and Outstanding Items

- We would like to see the ‘breadcrumb trail’ from KC restored – this allows you to immediately see which manual and which section you are in – much faster than trying to work that out from the ToC.
- We would like to see the Prev/Next function restored. Now, if you want to go from one section to the next, you need to move back up to the higher-level heading, then back down into the desired section.

- There IS a way to submit feedback about the site (rather than the text you were looking at), but it doesn't jump out at you.




Perhaps the most positive outcome of this event was that it shows the power of customer opinion. IBM made a change, customers reacted very vocally, and IBM addressed at least some of the concerns.

IBM Docs Changes and Outstanding Items

A few usability tips:

- When you do a search from inside the IBM Documentation site, you get this at the top of the list of hits:



- If you want to limit the books, or bookshelves, for the search, click on the dropdown arrow.

IBM Docs Changes and Outstanding Items

You can download a zip file containing all z/OS books in PDF format, INCLUDING the drafts of the z/OS 2.5 books, from here: <https://www-01.ibm.com/servers/resourceLink/svc00100.nsf/pages/zOSV2R5Library?OpenDocument>



2.5 Announce Disclaimer

All 2.5 announce publications are PDF only drafts and will contain all 2.5 technical content for GA.

PDF files for download

+ links open PDF file lists. Click or tap a link to show its list.

+ [Show all PDF files](#)

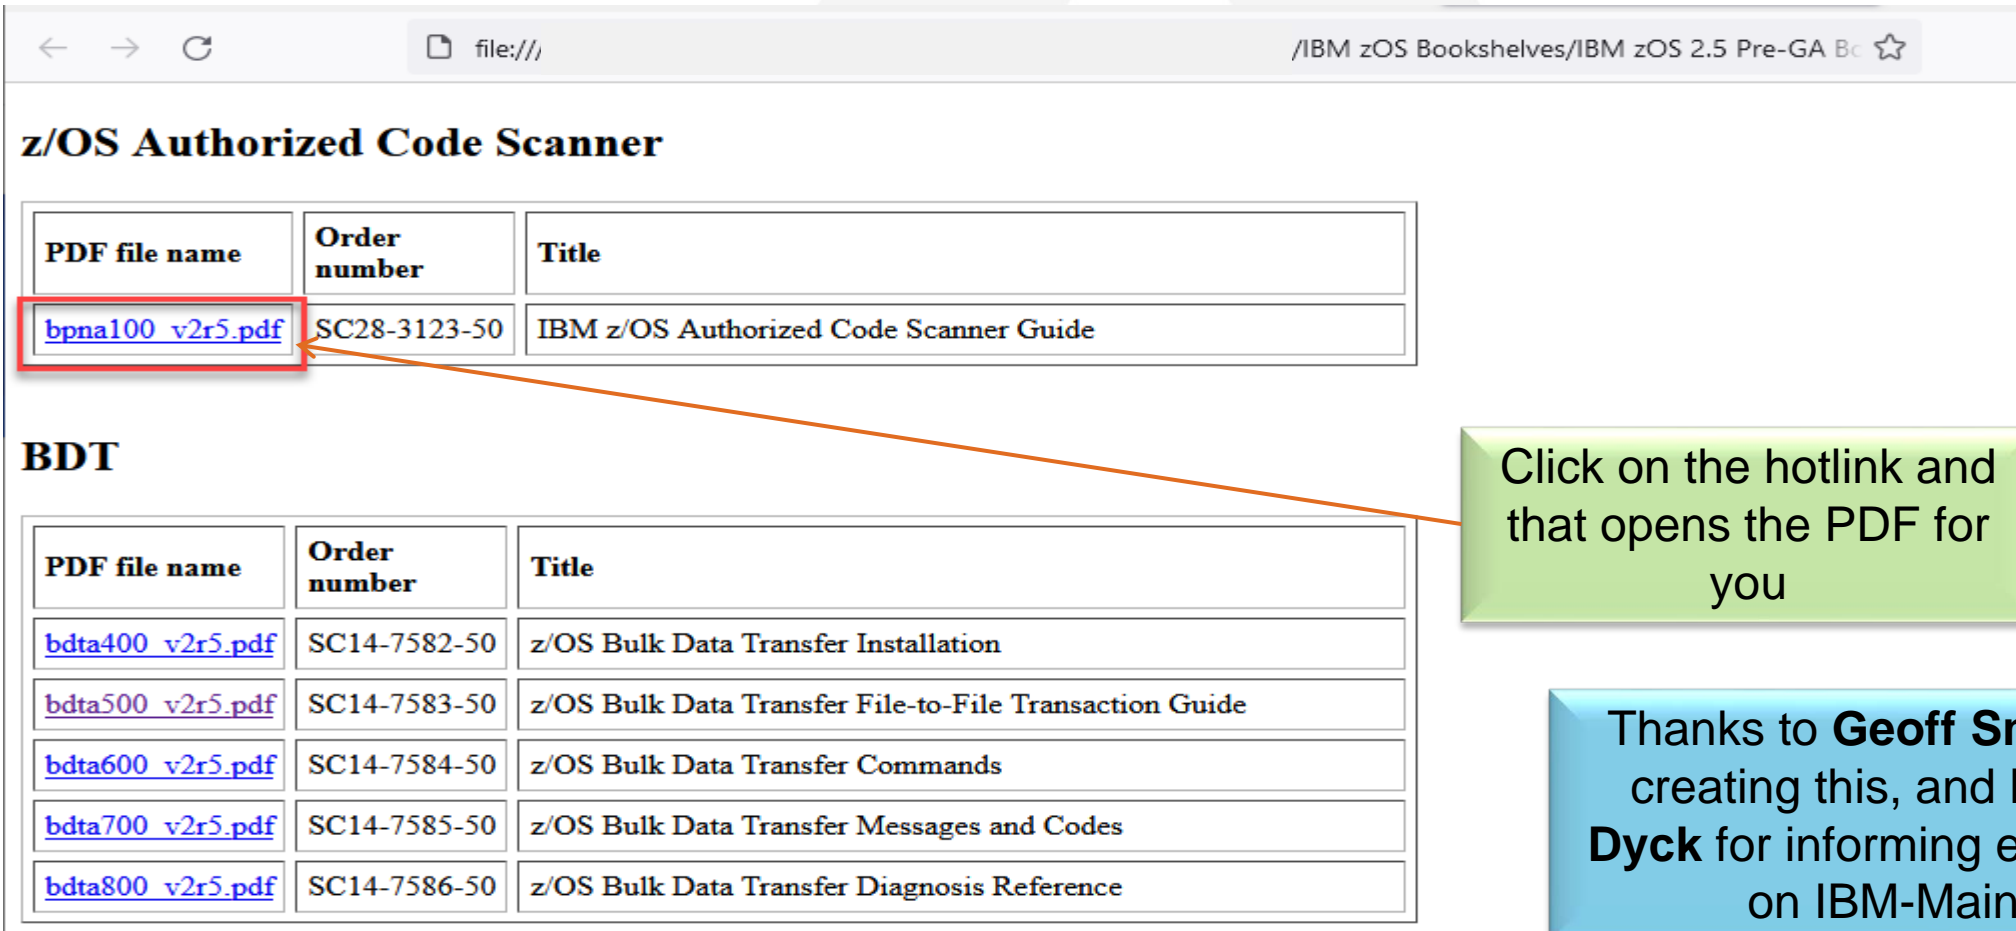


Download all z/OS V2R5 Library publications to ZIP file

309 files, 830.9MB total size
(uncompressed)

IBM Docs Changes and Outstanding Items

Unzip that into a folder and bookmark the index.html file. Then you get this:



z/OS Authorized Code Scanner

PDF file name	Order number	Title
bpna100_v2r5.pdf	SC28-3123-50	IBM z/OS Authorized Code Scanner Guide

BDT

PDF file name	Order number	Title
bdta400_v2r5.pdf	SC14-7582-50	z/OS Bulk Data Transfer Installation
bdta500_v2r5.pdf	SC14-7583-50	z/OS Bulk Data Transfer File-to-File Transaction Guide
bdta600_v2r5.pdf	SC14-7584-50	z/OS Bulk Data Transfer Commands
bdta700_v2r5.pdf	SC14-7585-50	z/OS Bulk Data Transfer Messages and Codes
bdta800_v2r5.pdf	SC14-7586-50	z/OS Bulk Data Transfer Diagnosis Reference

Click on the hotlink and that opens the PDF for you

Thanks to **Geoff Smith** for creating this, and **Lionel Dyck** for informing everyone on IBM-Main!

IBM Hot Topics has a 'Forever Home'

- After a period of bouncing from one website to another, IBM Hot Topics articles finally have a home of their own - <https://zos-hot-topics.com/>.

The screenshot shows the IBM Z Hot Topics website interface. On the left, a navigation menu includes 'About', '2021 Hot Topics', '2020 Hot Topics', 'Hot Topics Archive', 'IBM Z Content Solutions', 'IBM Documentation', 'IBM Z Community Blogs', and 'IBM Redbooks'. A red box highlights the '2021 Hot Topics' and '2020 Hot Topics' items, with an arrow pointing to a callout box labeled 'Recent issues'. Another red box highlights the 'Hot Topics Archive' item, with an arrow pointing to a callout box labeled 'All the way back to 2001'. At the top right, a search bar is highlighted with a red box and labeled 'Search'. At the bottom right, a 'Subscribe' button is highlighted with a red box and labeled 'Subscribe'. Below the button, the text 'Be the first to hear from the people who know.' is visible, and a link 'Subscribe to our RSS feed now →' is circled in red with an arrow pointing to the 'Subscribe' button. The main content area features an illustration of people interacting with digital screens and a laptop displaying 'HOT TOPICS'.

IBM Developer Help

IBM Developer Support

- We've been followers of IBM-Main and really appreciate when a developer contributes.
- We were extremely disappointed to see the deletion of the DeveloperWorks blogs, which were replaced by '[IBM Communities](#)'.
- Cheryl monitors the '[IBM Z and LinuxOne Community](#)', especially the COBOL Language group, and appreciates the IBM developer support given to the COBOL community.
- There is also a [z/OS group](#) that is similar to IBM-Main, but also provides more consumable announcement materials, such as '[What's New in z/OS 2.5 ISPF](#)' and the quarterly z/OS Continuous Delivery updates.
- You can follow these communities between SHARE conferences to keep current. And if you have a question, ask it there.

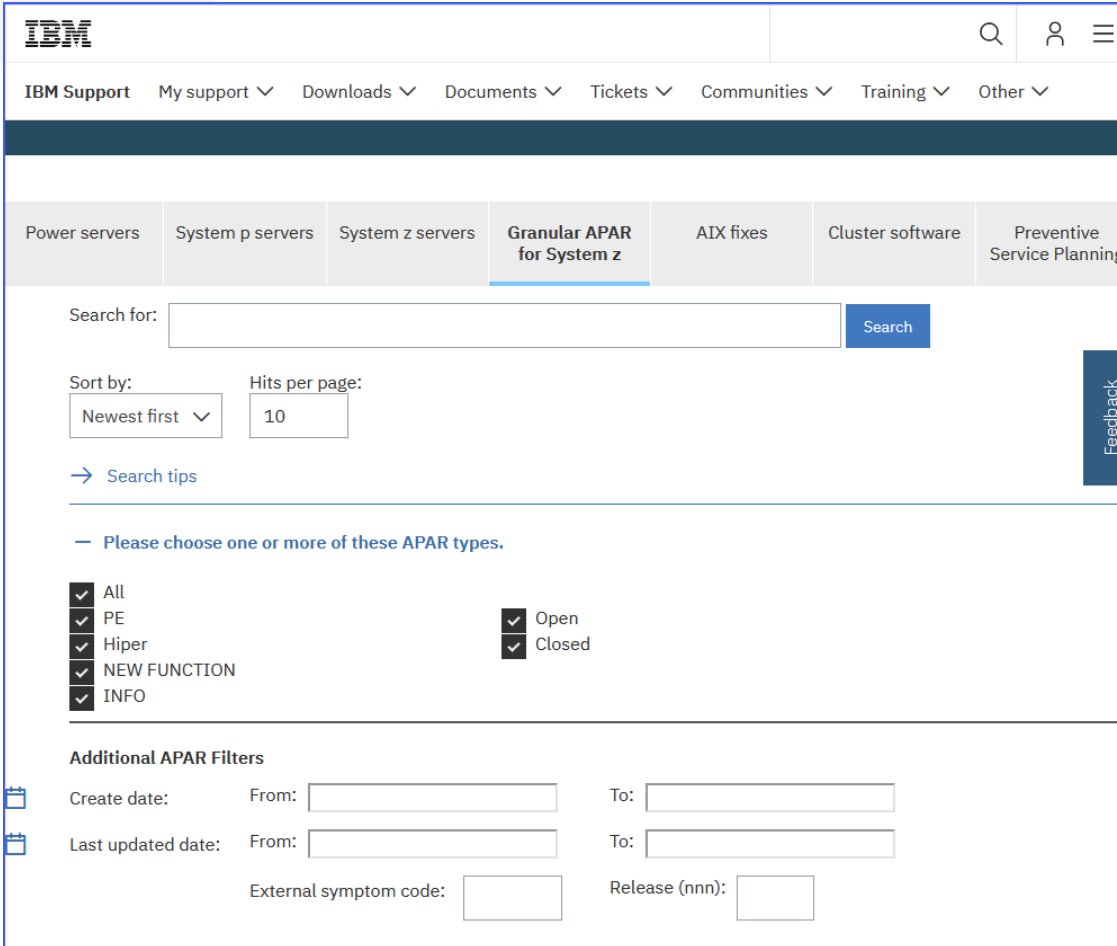
How To Find An APAR

Some people were recently discussing the best way to search for an IBM Z APAR.

Some use DuckDuckGo or Google search with 'apar xxnnnnn' when they know the number.

What do we use to review APARs?

- The most reliable site is the IBM Support Portal:
https://www.ibm.com/mysupport/s/?language=en_US. Apart from the occasional broken link, we find this to be the best site.
- As a fallback, or depending on what information we are looking for, we might use the IBM Granular APAR search site:
<https://www14.software.ibm.com/support/customercare/psearch/search?domain=gapar>



The screenshot shows the IBM Support Portal search interface. At the top, there's the IBM logo and navigation links like 'IBM Support', 'My support', 'Downloads', etc. Below that, there are tabs for different server types: 'Power servers', 'System p servers', 'System z servers', 'Granular APAR for System z' (which is selected), 'AIX fixes', 'Cluster software', and 'Preventive Service Planning'. A search bar is present with a 'Search' button. Below the search bar, there are options for 'Sort by' (set to 'Newest first') and 'Hits per page' (set to '10'). There's a 'Search tips' link. A section titled 'Please choose one or more of these APAR types.' contains several checkboxes: 'All', 'PE', 'Hiper', 'NEW FUNCTION', 'INFO', 'Open', and 'Closed'. Below this, there are 'Additional APAR Filters' including 'Create date' and 'Last updated date' with 'From' and 'To' date pickers, and 'External symptom code' and 'Release (nnn)' input fields.

IBM Technical Help Database

- On the <https://www14.software.ibm.com/webapp/set2/psearch/search?domain=sysz> webpage, you can do a generic search of all the documents listed (APARs, Flashes, Announcements, etc.).

The screenshot shows the search interface for the IBM Technical Help Database, specifically for System z servers. The interface includes a navigation bar with tabs for different server types: Power servers, System p servers, System z servers (selected), Granular APAR for System z, AIX fixes, Cluster software, and Preventive Service Plans. Below the navigation bar is a search bar with the text "Search for:*" and a "Search" button. There are also options for "Sort by:" (set to "Newest first") and "Hits per doc type:" (set to "10"). A "Search tips" link is provided. Under the "Document types" section, there are several checkboxes, all of which are checked: All, APARs, Flashes, FAQs, Preventive service planning, Product information, Redbook abstracts, Education, Solutions, Technotes, and White papers. A vertical "Feedback" button is located on the right side of the interface.

If you have z14 or z15, make sure you update CPENABLE to (5,15).

Check **John Shebey's** slides from [MVS Opening](#) for info about 4 HIPER APARs.

Get your network colleagues to check excellent [z/OS Communications Server Performance Update](#) session by **Mike Fitzpatrick** and **Dave Herr** for **important performance-related information** and APARs related to SMC V2, zCX, and AT-TLS. Some VERY impressive performance enhancements and helpful tips.

If you have a large common storage requirement, but are hampered by 31-bit Private area users, review JES2 APAR [OA61229](#) – this moves some JES2 control blocks (maybe 60MB?) to above-the-bar Private. See **Tom Wasik's** [JES2 Component Enhancements for z/OS V2R5](#) session for more info.

- And while there, check out the Spool Compression slides.

A little something for our **IMS** friends – APARs [OA60658](#), [PH33197](#), and [PH33306](#)
ADD support for IMS exploitation of Bypass Extent Checking – allows concurrent reads and writes within the same OLDS extent.

Db2 HIPER performance APAR [PH35088](#) - *High CPU Due to Long Prepares and Low Hit Ratio in Dynamic Statement Cache **When Dynamic SQL Plan Stability is Enabled.***

And another (**Db2 HIPER**) APAR [PH34200](#) - *High CPU in DIST Address Space After Long Running DRDA Threads Hit the System Profile Threshold Limit – **for systems that use the ‘Db2 System Profiles’ function** to monitor threads for remote applications accessing Db2 via TCP.*

Just so the WAS folks don't feel left out – here is a performance APAR for **Java/WAS**: [IJ29465](#) - Long GC Pause During Scavenge.



“Software Drag Racing”

- You might have noticed some discussion on IBM-Main recently about comparing the speed of various programming languages (initially not on Z).
 - As mainframers, we always take it for granted that Java is ‘slow’ compared to traditional programming languages.
 - **Andrew Rowley** from [Blackhills Software](#) ran a series of measurements comparing C, Java, and COBOL on z/OS on a z15, and kindly took the time to [document his results](#) – they make very interesting reading, especially when you look at the impact of relatively minor changes to the application *or* if you are using a sub-capacity CPC.

Tivoli Decision Support (or IZPCA) APAR [PH33033](#) zIIP enables the batch collector. Does NOT happen automatically after you apply the PTF – see APAR text for required manual actions.

z15 Upgrade Paths

We have been working with a number of customers, helping with evaluating and planning for upgrades to z15 CPCs.

1. The norm is to move to a z15 with the same number of CPs as your current CPC.
2. Some customers moving to sub-caps and getting more CPs as a result.
3. Others taking the opportunity to REDUCE the number of CPs.

Case 1 is BAU.

Case 2 requires careful checking to ensure that slower general purpose CPs will not have a negative impact on critical workloads.

Case 3 requires extra care with LPAR planning. Logical to physical ratio is likely to increase. Depending on your workloads and LPAR configuration, this *could* result in a higher LSPR MI Workload categorization.

- We *strongly* recommend using **Alain Maneville's** [LPAR Design tool](#) to optimize LPAR configuration for any upgrade, but *especially* for upgrades like this.

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.

The answer? Check out **Frank McCune's** [Refresh Your Catalog Maintenance Jobs](#) session **recording**. There have been a lot of great sessions this week, *but this one is a must-watch* if you are responsible for catalogs.

In most presentations and publications since 2010, we have been **STRONGLY** encouraging our customers to use z/OSMF. That's 11 years!

- And today's z/OSMF has many more features, has greatly improved performance, and has been redesigned for easy use.
- See Xiao Zhen Zhu's ("Joey") session this week on "[What's New in z/OSMF V2R5](#)".
- If you're having problems convincing others on your team about z/OSMF, be sure to look at the additional slides at the end of the deck that enumerate the benefits of z/OSMF. They're impressive.


Most sites report using it now, but the exploitation of all its features has been slow.

This survey comes from IBM's **Gary Puchkoff**, and two things are apparent:

- Slightly over half the clients utilized Workflows.
- Only 18% utilized or tried Software Management.

V2R5 upgrades will require training for most clients.

Which of the following z/OSMF services are you and your colleagues using?



- Workflows increased from 33.8% to 52.5%
- SCA, Desktop are v2.4 new features, won good adoption rate(5th and 7th)

Traditional CustomPac installs are going away. You can delay the inevitable a little if you rush in and order a z/OS 2.5 ServerPac right away), but would it not be better to participate and help influence its future direction?

Early feedback we have received:

- This new method does *not* eliminate the need for an experienced sysprog.
- There are functions missing that were in CustomPac.
- “New tenure” sysprogs feel more at home with z/OSMF than CustomPac.
- Apply all available z/OSMF service before you start.
- DON'T wait until you *have* to do this. Get experience with a CICS or Db2 or Sample PSI first.


Check out **Kurt Quackenbush's** [Installing z/OS V2.5 Using z/OSMF Software Management](#) session, then download the sample and try it for yourself.

And check out **Marna Walle's** “[Upgrade to z/OS 2.5 Technical Actions](#)” sessions for a treasure trove of z/OS 2.5 information (and stay until the end for the Q&A).

COBOL Migration – Why the Rush?

COBOL V4 produces code that will run on a 9672 (ARCHLVL 2?) Z15 is ARCHLVL 13!

- There are two driving forces
 - COBOL V4 compiler goes EOS on April 30, 2022. The compiled programs can run fine if they are LE-enabled. But if hit a problem trying to recompile, you can report problems only on COBOL V6+ - this means the migration of that program would need to be done *at that time*.
 - Companies want to reduce CPU time in order to reduce software costs. COBOL V6 allows them to do that with ARCHLVL compiler options (our clients have seen *up to* a 20% CPU and elapsed time reduction). The chart at the right shows that over 40% of all z/OS instructions can be used only on higher ARCHLVLs. Many of them are designed to reduce the CPU time for decimal calculations (i.e., a large part of COBOL apps).



* Thanks to **David Hutton** for this info

COBOL Update

One of the big challenges our customers are facing is the length of time it is taking to migrate from COBOL V4 (going End-of-Service on April 30, 2022) to COBOL V6.

Compiling with COBOL V6 takes more CPU time than COBOL V4, but the benefit should be lower CPU time when the program runs. Some teams are finding a little savings just from recompiling with COBOL V6, but the primary savings are if you use the ARCHLVL option.

However, there are tens of thousands of COBOL programs in some sites, and that is *before* you consider the complexity of dynamically-called programs that are called by many other programs.

These discussions led to the creation of our latest product, [AP4Z](#).

But in using our tool, we've found many things that can be done to save CPU time, even without using ARCHLVL.

COBOL Update

There are other ways to save CPU time.

- The program that consumed the most total CPU time of *all* application programs is a dynamically called one, so that would never show up in any SMF record.
- It was called over 4 Billion times! Would it be more efficient to move that code in-line in the calling programs? Need to find out who calls it.
- Its average elapsed time was about 8 times more than its CPU time. Why? Is it doing I/Os? Talking to Db2? If that was happening inline, would that processing be more efficient? Maybe easier on the L1/L2 caches?

modulename	language	Loads	Calls	Total ElapsedTime	Total CPUTime	Avg ElapsedTime	Avg CPUTime
XXXXXXXX	COBOL v5 / v6	18	4,290,642,435	61,200,526,547	8,074,390,394	14.26	1.88
XXXXXXXX	COBOL v5 / v6	6	1,444,955,475	2,494,565,302	518,911,776	1.73	0.36
XXXXXXXX	COBOL v5 / v6	58	263,158,716	419,381,778	97,499,302	1.59	0.37
XXXXXXXX	COBOL v5 / v6	11	257,997,146	169,830,451	82,525,681	0.66	0.32

Powered by Metabase Rows 1-19 of first 2000

COBOL Update

There are other ways to save CPU time.

- Identify executing programs that haven't been optimized and optimize them.
- Find old programs and modernize them, such as changing subscripted tables to indexed tables. Below is an example of two or more programs being loaded that haven't been recompiled in over 20 years. Do you even have the source?

optimizeOption	loadCount
N	83,707
2	10,561
0	4,362

Year	loadCount
2021	5,356
2020	600
2019	578
2018	6,446
2017	4,515
2016	517
2015	35,822
2014	44,641
2009	1
2000	151
1998	3

COBOL Update

While COBOL issues for a *Tuning Letter* article, Cheryl discovered that **Mike Chase** from the IBM COBOL Compiler Development group gives two fantastic **FREE** monthly webinars. One is on COBOL Migration, and the other is on COBOL Performance Tuning. You can register for upcoming webinars at <https://ibm.biz/COBOLWorkshopRegistration>. The same page provides links to previous recordings.

And don't miss **Tom Ross**' (Captain COBOL) [What's New in COBOL V6.3 Since GA: Continuous Delivery Items!](#) session tomorrow.



AND NOW – OVER TO MARIO....

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

Topics of the day

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 – [“z/OS Communications Server Performance Update”](#) by Mike Fitzpatrick and Dave Herr

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

1 Overview	33 Technical information
4 Key requirements	34 Ordering information
4 Planned availability date	44 Terms and conditions
5 Description	44 Prices
30 Statement of direction	51 Announcement countries
33 Program number	

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

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 SHARETM 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.

Artificial Intelligence 2021 Predictions - AIOps

By 2024

AIOps will become the **New Normal for IT Operations**, with at least **50%** of large enterprises adopting **AIOps Solutions** for automating major IT System and Service Management processes *



* Source:

IDC Future Scope Worldwide Artificial Intelligence 2021 Predictions

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



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!



WIKIPEDIA
The Free Encyclopedia

- Main page
- Contents
- Current events
- Random article
- About Wikipedia
- Contact us

Article [Talk](#) [Read](#) [Edit](#) [View hist](#) Not logged

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]



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](#)



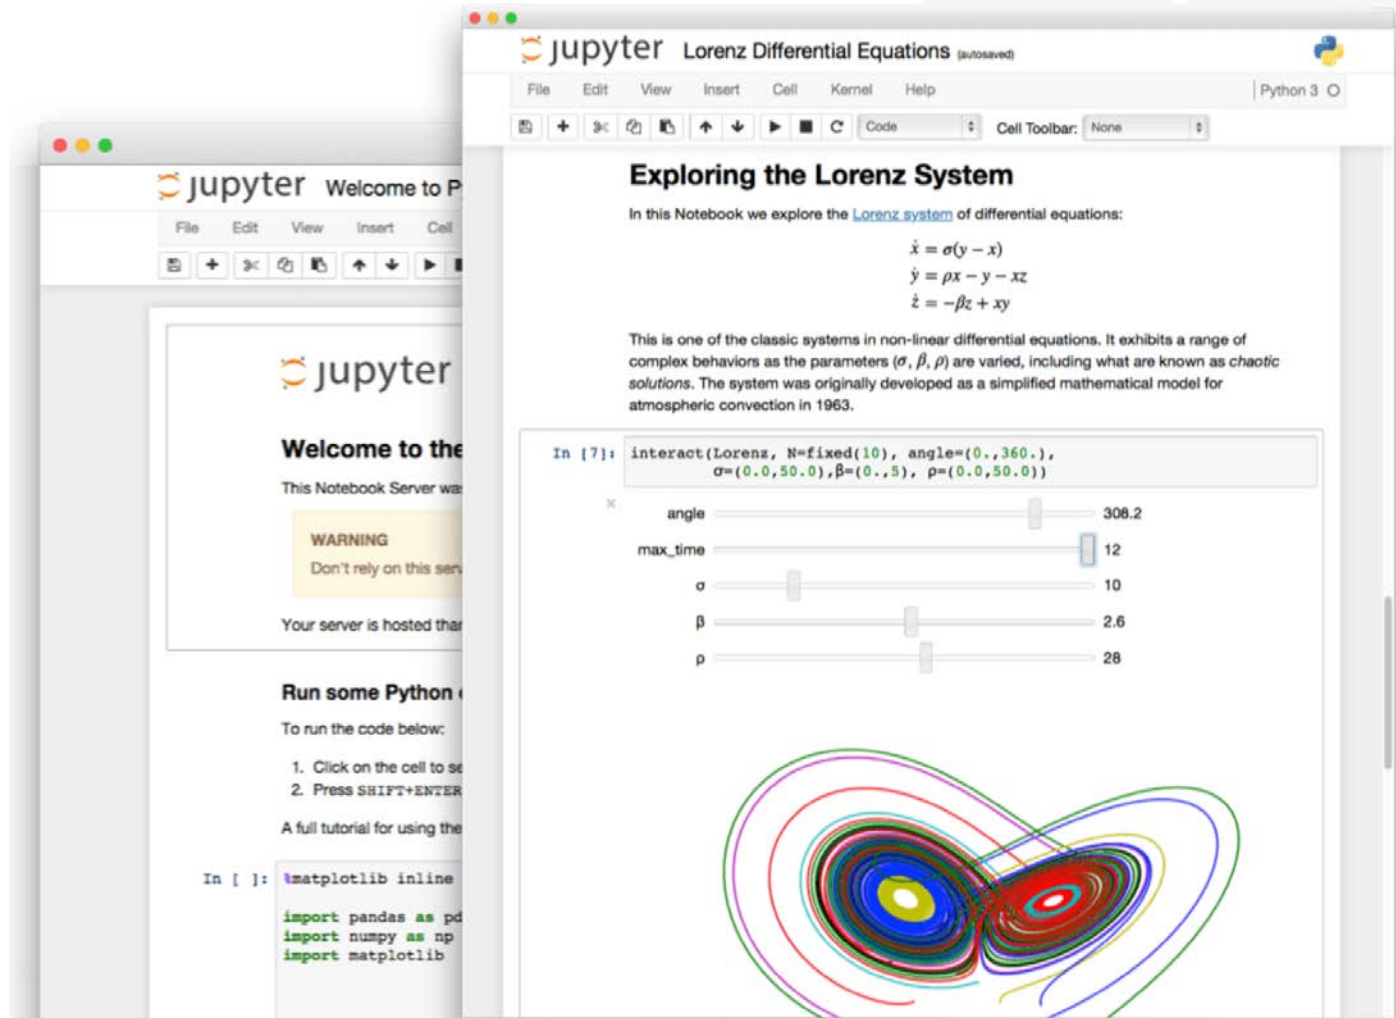
Credit: Jake vanderPlas, "The Unexpected Effectiveness of Python in Science", PyCon 2017

A sample Python “program”

*IcalParser

```
1
2 import icalendar
3
4 from pytz import timezone
5
6 if __name__ == "__main__":
7
8     icalfile = open("/home/mario/Downloads/mario.ics", "rb")
9
10    vcal = icalendar.Calendar.from_ical(icalfile.read())
11
12    localzone = timezone("Europe/Rome")
13
14    for component in vcal.walk():
15        if component.name == "VEVENT":
16            summary = component.get("summary")
17            description = component.get("description")
18            location = component.get("location")
19            startdate = component.get("dtstart").dt.astimezone(localzone)
20            enddate = component.get("dtend").dt.astimezone(localzone)
21
22    icalfile.close()
```

Jupyter 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



© Copyright IBM Corporation 2021

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)
```



System Utilisation



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.

180+ 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 ==> [red box]                      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 IEF989I SLIP TRAP ID=X33E MATCHED. JOBNAME=*UNAVAIL, ASID=0020.
00000090 1 *IGW048A IMF LRU Stalled SMSPDSE 876
00000090 last completed:08/07/2019 11:15:57.029629
00000090 BUFFER SPACE USED:16MB
00000090 2 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 be 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.

120+ 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%
```

Where to get our free tools



Free Tools

We occasionally create or come across small utilities or tools which we think everyone, not just subscribers, should have an opportunity to use.

[Home](#) > [Software](#) > [Free Tools](#)

We'll occasionally place various small utilities and tools on this page, free for the taking. We provide no warranty for these tools, nor do we provide any support. (If you have a problem or suggestion, please let us know and we may pursue it if we have time.) For our conditions of use, please see <https://watsonwalker.com/disclaimer/>. The utilities and tools have no copyright, and we request that you not attempt to sell them for profit. Click on the heading of each item to read more about the tool.

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

Thank you!

- For more information:
 - See our website at www.watsonwalker.com.
 - Send an email to technical@watsonwalker.com.
 - View our On-Demand session of Application Profiling.
 - 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 after z/OS 2.5 goes GA.
 - Access Out-of-Print IBM Systems Center Books at www.watsonwalker.com/publications/ibm-systems-center-books/. If you're willing to share any of your old ones, we'll do the scanning.

And don't forget to watch Kurt's [Installing z/OS V2.5 Using z/OSMF Software Management](#) session!

Complete your session evaluations for a chance at daily prizes!

To complete, visit
www.share.org/evaluation
and see your progress on the
leaderboard!

PERSON

PROGRAM

PROJECT