A PRACTICAL JOURNAL OF z/OS TUNING AND MEASUREMENT ADVICE

Cheryl Watson's Tuning Letter

Hi All,

Here's some exciting news: We told you in the previous issue that we hadn't been able to find a buyer for Cheryl Watson's Tuning Letter and that the final issue would be released in January 2025. This is NO LONGER TRUE!! It seems our announcement sparked a lot of interest, and we were inundated with conversations at SHARE Orlando. We are thrilled to tell you that we have found a buyer, and the Tuning Letter will continue! In fact, we're extremely confident we have found the *right* people to carry on what we started. That was VERY important to us!

We still have some of the details to work out so we aren't ready just yet to indicate who the buyer is, but Tom and I are certain that they will provide the same high quality exclusive content with the unbiased independence that you've appreciated until now. There are few companies we would entrust with our legacy - this is one of them. Their commitment and passion for the mainframe industry and its community is unparalleled. They are the perfect future keepers for the Tuning Letter. We're sure you'll agree. More news soon.

As Frank points out on page 10, my week at SHARE in Orlando was absolutely incredible! I am humbled by the sendoff I received from everyone involved. I didn't realize retirement would be so bittersweet. I haven't been able to send my thanks out to well-wishers because I was sidelined for six weeks right after SHARE - as a friend of mine says, “Getting older ain’t for sissies ☺”. But now I'm fully recovered and rarin’ to go. Now, where did I leave that paraglider brochure?

Life is good and I can't wait to share more details as we go forward!

Cheryl

Cheryl Watson's

Tuning Letter

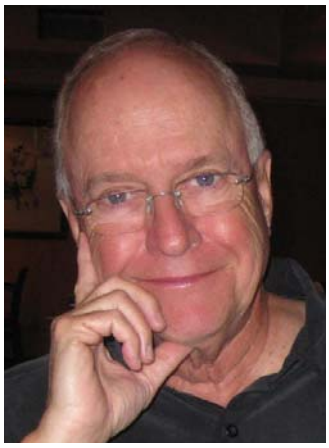

Tom Walker

2024 Number 1 (Vol. 34, No. 1)

Publishers: Tom Walker and Cheryl Watson

Editor: Frank Kyne

2024 Rates: Single-site (electronic & online archive), $1,780.
Multi-site discounts are available. Full rate schedule available at
http://watsonwalker.com/publications/tuningletter/rate-sheet/.

Payment may be made by major credit cards, by a check drawn on a U.S. bank, or by wire transfer. Send correspondence to:

Watson & Walker, Inc.,
1661 Ringling Blvd, PMB 3336,
Sarasota, Florida 34230-8036,
USA

Tel: 941-924-6565.

For customer service, send your emails to admin@watsonwalker.com. For technical questions, send emails to technical@watsonwalker.com.

Note: Implementation of any suggestions contained in this newsletter should be preceded by a controlled test and is the responsibility of the reader.

Watson Walker

Table of Contents

Management Summary

Welcome to our *Tuning Letter 2024 No. 1* issue. When we *started* this issue, we believed that this was going to be the last year of the Tuning Letter. However, Cheryl and Tom were approached by a number of companies during the SHARE conference, all interested in keeping the Tuning Letter up and running. Discussions are still ongoing, but now we are confident the Tuning Letter *will* continue to provide valued and independent information and guidance to our subscribers for many years to come. As you can imagine, it will take a little time to get everything agreed, but we hope to have details for you in the next Tuning Letter issue. As Cheryl would say - “Stay tuned” ☺.

Cheryl’s Last SHARE

Cheryl has been an institution at SHARE for as long as most SHARE attendees have been taking part. But all good things must come to an end, and as part of Cheryl and Tom’s retirement plans, they decided that the recent SHARE conference in Orlando would be Cheryl’s last one.

Given all that Cheryl has done for the SHARE organization and its members over the years, it was only fitting that Cheryl would get a special sendoff at her last SHARE conference. However, Cheryl isn’t just *any* SHARE member, so she couldn’t get just an average sendoff. Accordingly, Cheryl’s friends and admirers had organized *at least one* recognition event every day during the conference. Cheryl thoroughly enjoyed the week - it was truly a week to remember. For a summary of Cheryl’s week and some great photos, check out ‘Cheryl’s Last SHARE Conference’ on page 10.

Ransomware? On z/OS? *Surely* not!

Up until recently, my gut reaction to a question about Ransomware on z/OS would have been “Thankfully, that is one problem we *don’t* have to worry about.” Then I made the mistake of doing a little research on the topic. That led to a bit more research. And then more. And all that research led to an article that prompted Cheryl to opine “This is frightening AND enlightening”.

As you know, we tend to steer clear of security-related topics in the Tuning Letter. However, ransomware on z/OS is not just a security issue, it touches on *many* disciplines and is a logical follow-on to the topics we covered in ‘Recovery at Scale – Part 1’ in *Tuning Letter 2023 No. 2* and ‘Recovery at Scale – Part 2’ in *Tuning Letter 2023 No. 3*.

Ransomware is a particularly sensitive topic. The last thing we want to do is provide information that would help someone break into a z/OS system. On the other hand, we have

a duty to make our readers treat this topic with the urgency it warrants. We were very fortunate to have the assistance and guidance of **Mark Wilson**, Technical Director of Vertali. Mark is famous in both the traditional z/OS space and also in the scary world of IT security, so we couldn't have wished for a better collaborator.

With Mark's help, I believe we have managed to convey the reality of the risks facing the mainframe, but without providing hackers (or anyone else) with a cookbook of how to hack a z/OS system. If you don't read any other article in this issue of the Tuning Letter, *please* read 'Recovery at Scale - Part 3 (Ransomware)' on page 18. There is an old saying that "an ounce of prevention is worth a pound of cure" - well, you are going to need a LOT more than an ounce of *prevention* to protect yourself from a ransomware attach, but the cost of the *cure* might drive your company out of business.

User Experiences and Tips

The 'User Experiences and Tips' on page 34 article is always one of our most popular. Most of us have more work than time, so anything that helps us save a little time or do things more efficiently is always appreciated.

Our first tip in this article is related to getting the maximum value from the information in SMF type 23 records. By default, the type 23 records are created every 60 minutes, starting an hour after the system is IPLed. Our tip shows how to synchronize them with your CMF/RMF records instead, allowing you to align the information in the type 23 records with the detailed information about system activity in the CMF/SMF and other SMF records.

We also have *two* tips from our friend **Andrew Rowley** from Black Hill Software. The first one discusses measurements he took comparing processing times for CICS SMF records with no compression, CICS compression, CICS and zEDC compression, and zEDC compression on its own. He also provides information about using Java to drive the reading of the SMF records. His other tip describes his sample programs for reading SMF records in real time (using Java and *without* SMF exits or authorized code) and sending SMS texts for jobs or job steps based on their return code.

The final tip in this issue relates to IBM's zPCR tool, and the considerations for how it handles parked Vertical Low CPs. zPCR's default behavior might not be what you would expect, so if you are using zPCR for any CPC sizing activities, please make sure to check out this tip.

Introduction to Altair SLC

Many mainframe sites use SAS language programs, either in bought-in products such as MXG, or in home grown applications. In my experience, people that have learned to use SAS

really love it and become ardent advocates. Unfortunately, SAS System licenses are not inexpensive, making them a target for accountants trying to cut costs.

If you are under pressure to eliminate a SAS System license off your mainframe, one option is to move the application to another platform. However, it is frequently the case that the application is using data resident on z/OS. By the time you factor in the cost of remote access, or constantly transmitting the latest data to that other platform, you end up not really saving anything.

Another option, and the focus of our "Introduction to Altair SLC" article, is to replace the SAS System with Altair's SLC product. SLC, previously known as WPS (World Programming System) allows you to run programs written in the SAS language without having to license the SAS System on that z/OS image. Altair SLC is not guaranteed to be a 100% compatible replacement for SAS, however customers report savings of more than 50% compared to the cost of the SAS System, and both Altair and Merrill Associates support MXG running with SLC.

If you have the SAS System licensed on one or more z/OS images and are under pressure to cut costs by eliminating those licenses, we definitely recommend checking out Altair SLC. At a minimum, this will enable you to make a fully-informed decision. And possibly it will be a good fit for your environment, allowing you to replace the SAS System and earning some brownie points with your friends in your Finance department.

SMF's Many I/O-related Fields

While helping a client review their capacity planning and capacity evaluation processes recently, we noticed that some fields in SMF records that initially appear to be reporting the same thing actually have quite different values. A good example is "system I/O rate" - a metric frequently used to compare CPC capacities. When we looked a little more closely, we noticed that some fields differ by a factor of 2 or more. When we discussed this with our friend **Todd Havekost**, we agreed this would make an interesting topic for his "Making Sense of the Many I/O Count Fields in SMF" article in this Tuning Letter.

Todd has the luxury of sitting on a ocean of real customer SMF data. While two SMF fields might have similar-*sounding* descriptions, it is only when you look in actual customer data that you can determine if they are *really* reporting the same data. As an example, I would have assumed that the number of I/O interrupts handled by a system would be roughly the same as the number of I/Os started by that system. But when Todd looked in his data, we found that the number of interrupts was 7% greater than the number of I/Os in one system, but then ranged to as high as nearly 50 *times* greater in some other systems. This certainly isn't the first time I was wrong, but I never would have *known* that I was wrong without Todd's analysis of real world SMF data.

The numbers *are* all valid, however they are reporting on different aspects of I/O activity. Based on much analysis of his data, and invaluable help from various IBM teams, Todd was able to create a fascinating article that shows the relationships between the most popular I/O-related fields. With this understanding in hand, you can now be sure that any analysis you perform is based on the view of I/O activity that is the most appropriate for what you are investigating.

But that isn't all. To *me*, this article, and Todd's methodical approach to analyzing the data, provides an excellent example of the value of *really* diving into SMF data. Rather than just accepting the numbers in SMF fields at face value, Todd's article illustrates how you can develop a much deeper understanding of how z/OS works by considering why different fields are, or are not, reporting similar values. If you have a desire to become the go-to person in your site for performance questions, this article will, I believe, be a real eye-opener.

News You Can Use

Your teams simply don't have the time to find and review the many hundreds of APARs IBM opens and closes every month, especially if only a small subset are relevant to your environment. To save your colleagues' time, while also making sure you are aware of important new functions and performance-related APARs, *we* analyze APARs from the last quarter and present the results in the News article in every Tuning Letter issue. *'News' on page 62* also provides information about important hardware alerts, new IBM Techdocs, and upcoming technical conferences.

"IntelliMagic, an IBM Company"

Regular Tuning Letter readers will be familiar with the close working relationship we have developed with IntelliMagic over the last eight years. It truly has been a pleasure and an honor for us in Watson & Walker. The team in IntelliMagic are true z/OS experts - having 'colleagues' of their caliber to discuss questions with, and to fill gaps in our knowledge, has been invaluable to us. We also have benefited greatly from being able to use their IntelliMagic Vision product to quickly create powerful reports to support our Tuning Letter articles and conference presentations. In return, we are delighted to help out the IntelliMagic team whenever the opportunity arises. The founders of both of our companies have similar ethics and principles, making our 'partnership' very natural and frictionless.

As you might have seen, IntelliMagic was acquired by IBM on February 29, 2024, just before the SHARE conference. It is always slightly unnerving to see a smaller company being acquired by a much larger one. However, based on everything we've seen and heard, on both official and unofficial channels, IBM seems very serious about expanding both the customer base for IntelliMagic Vision and Vision's existing powerful capabilities.

Nearly every week we hear of another experienced z/OS performance person taking the retirement exit. The number of performance experts is constantly decreasing, while the performance and availability challenges faced by mainframe customers continues to grow. Realistically, the only way to bridge the void is through the use of intelligent performance products such as Vision. These products will *not* replace the people that have left, but they *can* make the remaining performance analysts much more productive. If you read back through Todd Havekost's Tuning Letter articles, you will see the common theme of using Vision's powerful capabilities to identify and investigate performance issues in significantly less time - that is more important now than ever before.

It is still early days in this new relationship. However, given the strengths of both parties, it holds the promise of becoming more than the sum of the two companies. If that comes to pass, I think we have some very interesting and exciting times ahead of us.

If you can think of a topic that you would like us to address in a Tuning Letter article, please let us know. We can't promise to address *all* suggestions, but the more interest we receive in a given topic, the more attention we will pay to it. Equally, if you have an experience that you would like to share with your peers, anonymously or otherwise, we would be delighted to work with you.

Thank you, and take care.

Frank and all the team at Watson & Walker

Cheryl's Last SHARE Conference

Cheryl's announcement of her upcoming retirement set the mainframe world alight, especially coming shortly before the Winter 2024 SHARE conference in Orlando. Cheryl has been a very active member of SHARE since her first conference in 1977. Over the ensuing nearly 50 years, Cheryl has an enviable attendance record and has become a SHARE institution. I still remember an attendee opining that "SHARE just isn't SHARE without Cheryl!" when a pending knee replacement operation caused Cheryl to miss her first SHARE in 40 years.

Cheryl is famous for her '*Cheryl's Hot Flashes*' Friday sessions at SHARE. I remember the first time I saw Cheryl in person - her SHARE Hot Flashes session was *completely* full, with people standing along the walls of the room and sitting on any available floor space. I remember thinking that I would have a heart attack if I ever had to stand up in front of so many people, but Cheryl just took it in her stride, chatting enthusiastically with this large audience in the same manner that she would chat one-to-one with an old friend. Cheryl was an *icon* - it never occurred to me that one day I would be working for her and her husband, Tom.

However, her very popular sessions were only one aspect of Cheryl's SHARE involvement. Cheryl was a SHARE volunteer starting with her very first conference. SHARE's volunteers are the unsung heroes that ensure that every aspect of each SHARE conference delivers the maximum value to attendees. Cheryl's commitment and dedication to her work as a volunteer typify her work ethic and are key ingredients in Cheryl becoming a trusted and admired figure, known to just about everyone in the mainframe industry.

An aspect of Cheryl's involvement with SHARE that many members might *not* be aware of is the time she put in behind the scenes to work with vendors. It might have been other tiny companies like Watson & Walker, or industry titans like IBM - Cheryl treated all with the same respect. She was also fearless in holding any vendor's feet to the fire when she felt their actions were not in the long term interest of their customers and the mainframe platform.I remember being at the SHARE Expo one time, and asking one of the attendees if he had seen Cheryl. He glanced around the area, saying "she was here just a few minutes ago, terrorizing one of the vendors." Despite Cheryl's diminutive stature, the idea of her haranguing some towering industry executive seemed perfectly normal.

Related to her quiet lobbying for customers' interests was Cheryl's role as the coordinator for SHARE's User Requirements process. IBM always encouraged customers to submit RFEs (Request for Enhancement) for improvements they would like to see in some IBM product. But unofficially, many IBM developers would advise customers "If you *really* want this to get

some traction, submit it as a *SHARE* requirement and lobby to get multiple SHARE members to support it". Cheryl put many long and frequently thankless hours into managing, organizing, and driving the SHARE requirements database and coordinating with the relevant IBM product owners. Knowing Cheryl, I'm sure she could tell you exactly how many days every SHARE requirement has been open (and I have IBM product managers that can attest to Cheryl's memory and tenacity for the SHARE requirements under her stewardship).

Given all this history with SHARE, and Cheryl's contributions to our industry, it was fitting that the SHARE board and the SHARE attendees made Cheryl's last SHARE conference a very special event.

The honors started at the keynote session on Monday morning, where SHARE President **Scott Fagen** presented Cheryl with the SHARE Distinguished Service Award.

Figure 1 - Cheryl Receiving SHARE Distinguished Service Award

This award recognizes outstanding service and support from someone who has worked for and with SHARE management over a period of years. As someone wryly commented, Cheryl had already been awarded just about every other SHARE award, so the

Distinguished Service Award was the only one left for SHARE to mark Cheryl's incredible contribution to SHARE over the last 47 years.

Speaking of Scott, he was one of the key group of people that worked to make the conference extra special for Cheryl. The others that put huge amounts of time, energy, inspiration, and love into organizing the tributes were **Art Gutowski**, **Marna Walle**, and **Steve Warren**. I would like to offer my personal thanks to these folks for making this a SHARE that Cheryl will never forget.

Immediately following the keynote session, Art (the fresh-faced young lad in Figure 2) introduced Cheryl for one of her trademark 'fireside chats' about her career at the *z/OS System Programming & Hardware KickOff* session.

Figure 2 - Art Gutowski Introducing Cheryl at the z/OS System Programming Kickoff

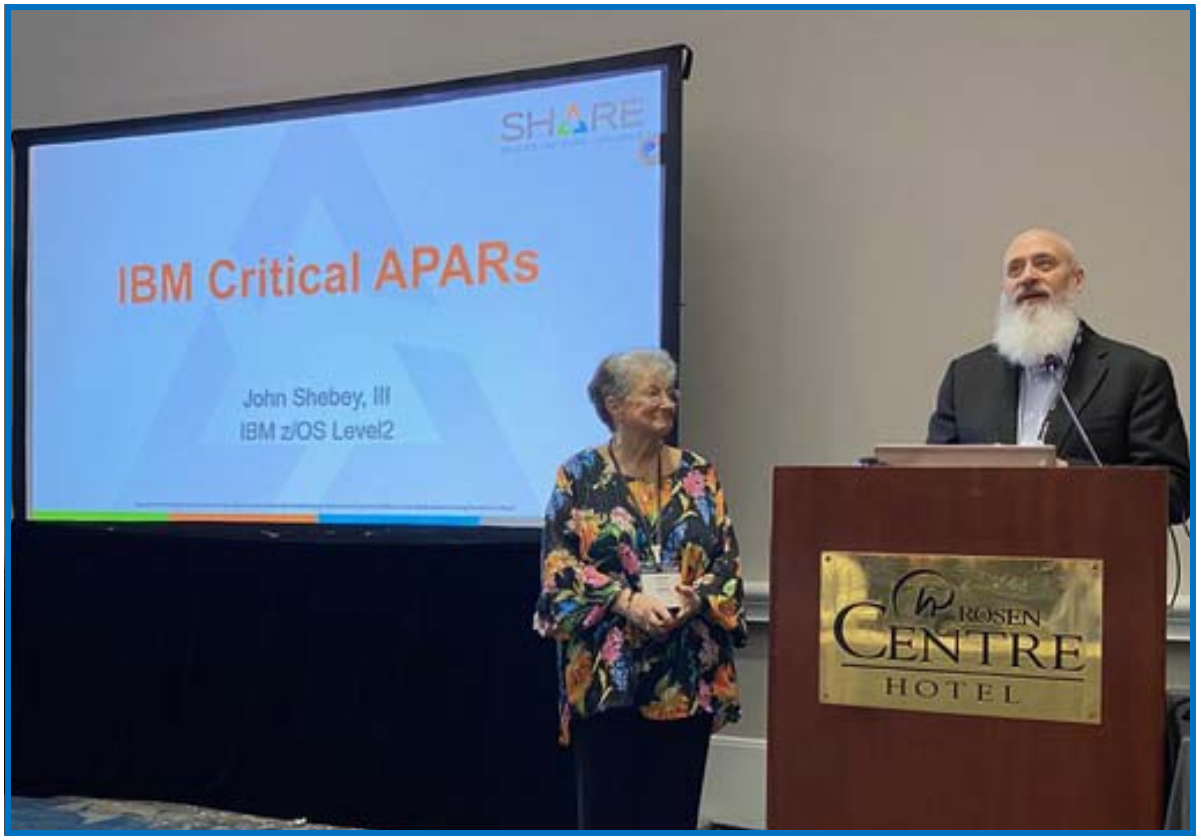

Cheryl's career, and her unique position as everyone's top trusted advisor (see Planet Mainframe article '*Cheryl Watson: Mainframe Influencer*') are real inspirations to any young mainframer that wants to make a difference, but also wants to stay in the technical, rather than the management, career track. Cheryl reached her position because of her work ethic, her always putting her customers' interests ahead of her own, her passion for research and learning, her love of the mainframe platform, her objectivity and fairness, her respect for everyone, regardless of their job title or position, and by assiduously avoiding industry politics - put simply, Cheryl became the industry's trusted advisor because she *earned* people's absolute trust.

In the words of two of the attendees at her session:

- "*To have gotten the opportunity to hear the reflections of Cheryl Watson, who has had such influence in my own career, was both bittersweet and most moving. Thanks for this opportunity*".
- "*Mrs Watson is amazing!*"

To cap a perfect session, Marna presented Cheryl with her favorite flowers, orchids. Cheryl and Marna are two of the best known girls in a largely-male mainframe world, so Marna's gift was a little like passing the baton to the next generation.

At a closed door session on Tuesday evening, **Jeff Cherrington** (representing Rocket Software), **Nicole Fagen** (representing Broadcom), **Paul Spicer** (representing BMC), and **Matt Whitbourne** (representing IBM) each gave a lovely tribute to the role Cheryl had played in their careers and how she worked tirelessly to promote the needs of customers and the mainframe platform in general (they were too polite to mention the 'terrorizing vendors' part ;)). Despite all the competition between these firms, it was really nice to see them join together to offer their gratitude and respect to Cheryl - that really meant a lot to her. On top all that, Steve Warren presented Cheryl with her very own z16. Well, OK, it was a Lego *model* of a z16, but it had the advantage that Cheryl was able to bring it home in the glove box of her car - how many other z16 owners can say that? ☺

Figure 3 - Cheryl With Her Portable z16

On Wednesday, the SHARE Planning team presented Cheryl and Tom with two picture frames loaded with bank notes from the many MANY countries they have visited over the years, along with notes from some destinations that are still on their bucket list (see Figure 4). Art and Marna did an amazing job of reaching out to other SHARE members and attendees to get the over 50 notes that made up this really thoughtful and heartfelt gift that now hangs in Cheryl and Tom's home in Sarasota.

Figure 4 - Cheryl Says "Learn the mainframe and see the world!"

Cheryl was also presented with her favorite cake (thanks for the inside line on that, Tom!) to celebrate her retirement.

Figure 5 - Cheryl's Favorite Health Food

All-in-all, it was a very enjoyable and memorable lunch and a nice chance in a hectic week for Cheryl and Tom to meet with many of their long-time SHARE friends.

On Wednesday evening, Cheryl, Tom, Mario, and I had the privilege of joining our friends in IntelliMagic and our common customers for the traditional IntelliMagic Customer Dinner. It is always a pleasure to have an opportunity to spend an evening with such nice people, but this customer dinner was particularly special. We had the all-too-rare opportunity to spend time with the brilliant and gifted (and very humble) co-founder of IntelliMagic, **Dr Gilbert Houtekamer** (shown with Cheryl in Figure 6 on page 16). It was really a very special event - in the company of *two* of the world's most outstanding mainframe performance experts, both making their last appearance at SHARE.

Figure 6 - Two Mainframe Performance Giants - Cheryl and Gilbert

It was also poignant because both Cheryl and Gilbert were celebrating their retirement. From a work perspective, it was sad to see our industry losing two of our leading lights, especially at a time when we so badly need *more* mainframe experts, not fewer. But from a personal perspective, both Cheryl and Gilbert have worked long and hard, and both richly deserve the opportunity to spend time on the many other rich experiences that await them out in the 'real' world. The IntelliMagic team had very kind words, and very thoughtful and special gifts, for Cheryl, Gilbert, and yours truly. That was a night that I will not soon forget, and I hope that I will have the opportunity in my retirement to spend more time with the many friends I have made through our relationship with the IntelliMagic team.

Cheryl's last SHARE conference closed fittingly, with her making a short speech at our zRoadshow session on Thursday. The zRoadshow session is the descendant of Cheryl's Hot Flashes presentation (which she kindly renamed when I started co-presenting that session with her!), so the timing of that session at the end of the SHARE conference seemed very apt.

Figure 7 - Cheryl and Frank

Of course, in typical Cheryl fashion, she used her last SHARE presentation to thank me for my work for Tom and her. This is by far THE BEST job I've ever had (and I've been exceptionally fortunate to have some really great jobs during my career), so it really should be me thanking them, rather than the other way around. But that is our Cheryl - always downplaying her own contributions and making sure to thank everyone else.

Finally, for the many people at SHARE and since then that have asked me "Is this *really* Cheryl's last SHARE?" - all I can say is, with Cheryl and Tom, every day is an adventure, so who knows what the future may hold. But I wouldn't be the least bit surprised to see Cheryl arriving by para glider at a future conference. As Cheryl would say, "Stay tuned." ☺

Recovery at Scale - Part 3 (Ransomware)

This series of 'Recovery at Scale' articles is intended to help you think about how you would recover from events that could occur on your z/OS systems. The first article, 'Recovery at Scale - Part 1', in *Tuning Letter 2023 No. 2*, discussed the loss of a sample critical system data set - a VTOC on one or more volumes. The next article, 'Recovery at Scale - Part 2', in *Tuning Letter 2023 No. 3*, discussed the more insidious problem of logical data corruption. In *this* installment, we will address Ransomware on z/OS.

The US government Cyber Security and Infrastructure Agency defines ransomware as:

> "*Ransomware is an ever-evolving form of malware designed to encrypt files on a device, rendering any files and the systems that rely on them unusable. Malicious actors then demand ransom in exchange for decryption*."

IBM's '*Definitive guide to ransomware 2023*' paper describes ransomware as:

> "*Ransomware is an online attack perpetrated by cybercriminals or nation state-sponsored groups who demand a monetary ransom to release a hold on encrypted or stolen data*."

There are many similar definitions on the Internet, but they all boil down to the use of encryption to deny the owning organization access to its data. The underlying common thread is that access to your data can be resumed if you are willing to make some form of payment. This is fundamentally different to an attack that *destroys or corrupts* your data, leaving you no choice but to recover from backups.

Not long ago, we would have said that a successful ransomware attack on a z/OS system was not in our top 10 (or top 100) list of things to worry about. However, this is an area that is recently getting a *lot* of attention from people that are much more plugged in to the security and risk side of our industry than I am. If *they* are worried about it, we felt it is time our readers also started worrying about it.

Target Audience

This series of articles is intended for *anyone* that works in a mainframe environment. While ransomware is a topic that is naturally associated with the security team, it is really something that should be at the forefront of *all* our thinking in these dangerous times. Within the *work* context, ransomware and security are business problems, not IT ones - the policy and investment decisions must be made at the *corporate* level and *implemented* by the IT staff.

We don't have the expertise or the time to provide a comprehensive guide for protecting yourself from ransomware - and even if we did, it would be out of date a week after we sent it to the publisher. Rather, our objective is to help you get a conversation going with your colleagues about what things *you* need to consider. Ideally, you will book a meeting room with a very large whiteboard for the day, provide lots of strong coffee, some nice treats, and maybe a defibrillator or two.

Before we go on, we want to thank the loyal Tuning Letter readers who provided invaluable knowledge, experiences, and assistance during the creation of this article. We especially want to thank Vertali's **Mark Wilson** for all his help, insights, and suggestions.

Introduction

We expect that you might be thinking "*This is z/OS. We don't have viruses or malware. Why do we need to worry about ransomware?*". To be honest, I used to feel the same way myself. After all, if it takes *years* for data and security experts to implement pervasive encryption in their shops, it would be impossible for some school kid to break into your system and encrypt everything before you noticed, right?

Unfortunately, it is not that simple. For a start, if someone is seriously trying to get into your mainframe, it is very likely *not* a school kid (and if a school kid IS able to get into your z/OS system, then you *deserve* to be hacked). More likely, it is some professional hacking organization or a state-sponsored hacking group. If you think the idea of state-sponsored hacking is a fantasy, you might want to take a look at a recently issued a report[1] by the US Federal Bureau of Investigation (FBI) and the US Cybersecurity and Infrastructure Security Agency (CISA) about such a group hacking into:

- U.S. Department of Energy (DOE)
- U.S. Environmental Protection Agency (EPA)
- U.S. Transportation Security Administration (TSA)
- Australian Signals Directorate's (ASD's) Australian Cyber Security Centre (ACSC)
- Canadian Centre for Cyber Security (CCCS), a part of the Communications Security Establishment (CSE)
- United Kingdom National Cyber Security Centre (NCSC-UK)
- New Zealand National Cyber Security Centre (NCSC-NZ)

These are very serious organizations - if hackers were able to get into *their* systems, do you honestly believe that they would not be able to get into *yours*?

[1] '*PRC State-Sponsored Actors Compromise and Maintain Persistent Access to U.S. Critical Infrastructure*'

Consider how much modern society depends on mainframes. How long would you be able to go without being able to use your credit or debit cards? How about if you couldn't access your bank account to pay your bills or get cash from an ATM? For all those 'cord-cutters' that got rid of your traditional phone service and cable TV, how would your life be impacted if the Internet and/or cellular phone service were down for an extended period? Imagine if your geographic area had no electricity for an extended period - you wouldn't be able to buy anything (even with cash) because the cash registers wouldn't work, you would have no light, you wouldn't be able to put gasoline in your car (because the gasoline pumps wouldn't work), you mightn't even be able to cook or have heating or cooling. If you look at the types of organizations that are being targeted by 'professional' ransomware groups, it is frequently the ones that keep society going, so don't think you are off the hook if you don't work for a mega-bank.

The point of all this is that taking out a single z/OS sysplex could have a far larger impact than taking down a few UNIX servers. Breaking into a z/OS system is hopefully a lot harder than breaking into a distributed system, but the 'rewards' are much larger. This makes mainframes a prime target for someone wishing you ill.

If you are still thinking "*That's all well and fine, but I've never heard of a ransomware attack on a mainframe*", I have a question for you. If *your* company or organization was the victim of a ransomware attack on your mainframe, would *you* announce it to the world? Just because you haven't heard of it happening somewhere doesn't mean that it hasn't happened. And, to be honest, I believe *very little* in the sanitized, lawyer-approved, statements issued by companies after they have been hacked. So just because I'm not reading stories in the general media about mainframes being victims of a ransomware attack, that doesn't give me confidence that it isn't happening.

Prozac[a] time

ALL of the information provided in this article is easily accessible to anyone with access to Google or other search engines. So if you are thinking, "*They really shouldn't be publishing this information*", I'm sad to inform you that there is other material in the public domain that contains *far more* detailed and dangerous information than anything we provide in this article.

In order to fully comprehend the challenge you are facing, I *beseech* you to spend just half an hour Googling 'z/OS' together with terms such as 'hacking', 'ransomware', or 'exfiltrate'. Then have a look through the material listed in 'References' on page 31. If you don't need medication for your nerves after that, you are a much stronger person than me.

a. Prozac is a medication that is used to treat certain mental/mood disorders such as depression or panic attacks.

What *Exactly* Is a Ransomware Attack?

For the purposes of this article we will stick to the description of ransomware described above; specifically:

- It involves an unauthorized person encrypting some or all of your data.
- It differs from the scenarios in the first two articles in this series because it *might* be possible to make your data/applications/systems available again with no loss of data, no recovery, and no recovery from the recovery. Assuming, of course, that you are willing and/or able to pay the ransom.

However, a ransomware attack, particularly by a group that is sophisticated enough to get into a z/OS system, is unlikely to be as 'simple' as just encrypting your data. More than likely, the attackers will first take a copy of your most sensitive data. This results in them being able to not only demand a ransom to regain access to your data, but also threaten to make your data freely available to the public and/or sell that data on the Dark Web. Of course, like all blackmailers, if they find a victim that is willing to pay them once, they are very likely to hit that victim a second or more times - I recently saw a paper that reported that nearly 80% of companies that paid a ransom were hit again within 6 months.

A Sample Real World Experience

An American healthcare company called Change Healthcare was the victim of a ransomware attack in February of this year. You can find out more about the impact of the attack in this American Medical Association report. Given the impact, it is understandable that they decided to pay the $22,000,000 ransom.

However, surprise, surprise, it transpires that there really *is* no honor amongst thieves. The group that actually hacked the system were not given their share of the spoils by their client. They obviously concluded that bringing their complaint to the Better Business Bureau would not achieve much, so they decided that they *also* would look for a ransom from Change Healthcare in return for not exposing the terabytes of data they had stolen before they encrypted it.

The cost of the attack on Change Healthcare *up to the time of writing* has been $872,000,000, and it is expected to be *double* that by the time the dust settles. Wow, that's quite a bit of money. They probably could have afforded to hire a few more security administrators for that sort of money. Despite the age-old adage that "Prevention is better than cure", we humans seem to have a boundless ability to find funding *after* a disaster occurs, but nearly zero ability to find funding to *avoid* the disaster in the first place.

If you want to know more about this event, just Google "Change Healthcare" - just about all the hits will be about the cyber-attack.

In any case, it is obvious that protecting yourself from, and recovering from, a ransomware attack is not something that can be achieved by any single IT department or discipline. Encryption is generally viewed as the bailiwick of the IT security team, while resiliency is typically owned by the infrastructure teams. Ransomware straddles both disciplines:

- It starts with someone gaining unauthorized access to your system and escalates to that individual using encryption or manipulation of encryption keys as the means to impact your organization's business.
- Some or all applications in your system will be unavailable, and recovery from the attack (if you refuse to pay the ransom) will likely involve some level of data/file recovery. Even if your organization agrees to pay the ransom, there will be a huge amount of work for the infrastructure staff to investigate exactly what happened, *how* it happened, ensure everything has been returned to normal, and then address the weaknesses in your processes that allowed the hacker to be successful.

Because this article is part of our Recovery at Scale series, we will focus primarily on the availability aspects of a ransomware attack. An attack that results in unauthorized access to your company's data but no impact to availability, while obviously critically important, is somewhat outside the scope of this article.

I say 'somewhat' because there is a very fine line between someone gaining unauthorized access to your z/OS system, and that individual then using encryption to lock out your applications. If someone is skilled enough to gain unauthorized access to your z/OS systems with the intent of exfiltrating your data, the logical next step would be to encrypt your data to divert attention from their theft of your data and also potentially generate some additional revenue for themselves. As you can see, ransomware is a risk that must be tackled in a holistic manner.

“As A Service”

Don’t you just love the way marketing departments seem to have an uncontrollable need to jump on the latest ‘trending terms’? For example, we now have Software-As-A-Service, Platform-As-A-Service, Infrastructure-As-A-Service, Function-As-A-Service, Authentication-As-A-Service, XaaS (‘Anything-As-A-Service’), and countless others. Not wishing to feel left out of things at the next marketing convention, the hacker community now offer Ransomware-As-A-Service.

This is such a sad reflection on the state of our society that I don’t even know where to start. However, imagine if someone were to concentrate all available expertise (and throw in some AI for good measure) on creating an RaaS offering specifically targeting mainframe systems - that is a challenge *I* would not want to face if I was the executive responsible for delivery of z/OS-based services to my organization. Perhaps it would be better to start preparing *now*, rather than waiting and having to recover later?

Stages of a Ransomware Attack

The number and names of the stages of a ransomware attack vary depending on which organization you refer to. For example, IBM uses five stages:

- Initial access
- Post-exploitation
- Understand and expand
- Data collection and exfiltration
- Ransomware deployment

If you Google “stages of a ransomware attack”, you will find many organizations, each using their own name for between three and ten stages. However, the underlying actions are pretty common regardless of how you label them. Because all our readers are users of one or more IBM operating systems, we will use the IBM structure and naming to describe a ransomware attack in a little more detail.

Initial Access

The first step, not surprisingly, is for the hacker to gain access to your network.

This might be achieved by exploiting a vulnerability in your network. Hardly a week goes by without some new exposure being identified. Given the plethora of network hardware devices (routers, switches, firewalls, and so on), firmware levels, management and security software, and interactions between all these (likely from multiple vendors), it is hardly surprising that network management teams are in a constant battle to try to keep up with the hackers, never mind staying one step ahead of them. Ironically, the extremely high availability demanded by customers and your own organization make it difficult to schedule planned outages required to keep everything up to date.

Even if you are diligent and fortunate enough to avoid any vulnerabilities in your network configuration, you still have to deal with the risk of malware such as a Remote Access Trojan which can provide access to your network to an unauthorized external party.

There are *very* few z/OS systems today that are not connected to a network that is accessible from the Internet. There are even fewer where *every* device that can access that system is connected to *only* that system and *nothing* else (no Internet access, no email access, no USB connections, no ability to FTP to or from a device that does not have the same restrictions, etc.).

Networking is *far* outside our area of expertise, and we are sure your network colleagues can tell you stories that will give you many sleepless nights. You can also find a little more information specifically aimed at a mainframe audience in the TechChannel article '*Mainframe Security: Are There More Gaps in Your Defenses Than You Thought?*' by **Trevor Eddolls**. The point here is that you should be mindful of how likely it is that a determined and knowledgeable hacker will be able to reach a point where they can attempt to logon to at least one of your z/OS systems.

Post-Exploitation

Having gained access to your network, the next step is for the hacker to obtain a userid and password on at least one of your systems. There are seemingly endless ways for them to do that, many of which are the same mechanisms that hackers use to obtain a victim's credentials so they can access their bank account or health information or other sensitive data - for example, phishing, keystroke loggers, network sniffing, Remote Desktop Protocol (RDP), unaddressed security exposures, or the use of default passwords, just to name a few.

But that is only the start. Apparently there are sites on the Dark Web where one can purchase thousands of userids and passwords, including, possibly, userids and passwords from *your* systems. And while I was working on this article, I heard that OpenAI have announced a new capability called Voice Engine, that can recreate a person's voice from a 15-second recording. The technology is so powerful that it has only been released (so far) to a small set of businesses because they (OpenAI) are still trying to understand its potential dangers. There was also a story about some poor financial institution employee who participated in a video call with their colleagues and followed the directions of their CFO to wire $25,600,000 to someone that turned out to be a fraudster. But it wasn't only the recipient that was a fraud - apparently every 'person' on the call, other than the one that transferred the money, was a deepfake. Then today I saw a story in Wired magazine about the US Air Force experimenting with AI-controlled fighter jets in dogfights with human-controlled ones, and another story titled "Microsoft teases deepfake AI that's too powerful to release"[2]. Ah yes, every day brings a new bounty of reasons to relax, secure in the knowledge that my mainframe is immune from all this madness - or maybe not.

Combine all this 'super' new technology with the ability for just about anyone to run their own AI, and I think it is nearly guaranteed that some unauthorized party is eventually going to try to get on to your system unless you start taking defensive actions *right now*.

The use of MFA (Multi-Factor Authentication) massively reduces the risk of a system being compromised. MFA should be used for *all* users of the system, not just the so-called

[2] If it is too powerful to release, why are they talking about it? Are they *deliberately* trying to entice hackers to steal it?

privileged users. Deploying MFA reduces the risks poised by shoulder surfing, keyboard logging software, users writing their password down, and so on.

I'm sure (I hope!) your security colleagues are already heavily engaged in preparing for these threats. For those of us that are *not* security gurus, some documents that might help you better understand the steps that can be taken are:

- The US government's CISA '*#StopRansomware Guide*'.
- The UK's National Cyber Security Center '*A guide to ransomware*'.
- IBM's '*Definitive Guide to Ransomware 2023*' paper.

Most of the activities in the first two stages involve an unauthorized person gaining access to your network and getting a ID and password to logon to your system. On that basis, most of the actions to protect your organization from this stage would lie within your security and network groups. However, *everyone* in your organization must play their part by being aware of how benign actions on their behalf (such as opening a PDF that appears to have been sent by a colleague) could open the door for an attack.

The Sneak Attack

I have it on good authority that there is more than one company in the world where production jobs (with production levels of access) get submitted to the mainframe from distributed systems. I guess it is not beyond the realms of possibility for someone to hack into one of those distributed systems and modify the JCL that is automatically submitted to add in a few steps of their own, thereby effectively giving them access to your system without ever having to actually logon to your system. IF that were to happen in your organization, would anyone even notice before it was too late?

Supply Chain Attack

In a somewhat similar vein (someone hacking into something that is not under your control) is the idea of a supply chain attack. An example of this is the 'Solarwinds' hack in 2020. Solarwinds is a software company based in the US. One of its products is a performance monitoring system called Orion. As a performance monitor, Orion needs privileged access to IT systems to obtain log and system performance data. Orion is installed on distributed systems around the world.

So, we have a piece of software that has privileged system access, and is already installed on tens of thousands of systems - that sounds like a magnet for would-be hackers. And that is what happened - hackers managed to insert malware into the Orion software in the Solarwinds systems and then sit back and wait for Solarwinds to unknowingly distribute that to their customers. When the corrupted Orion software was installed and activated by a customer, the malware then spread all over the customers network, opening back doors that the hackers could then use to break into the customers' systems.

This leads me back to the same question - IF something like this were to happen in *your* organization, would anyone even notice before it was too late? As uncomfortable as it is to ask this question, how confident are you that *every* vendor that provides software for your z/OS systems has bulletproof processes to ensure something like this could *never* happen to them?

Understand and Expand

Having obtained credentials to access your system, the next step is to start gathering information about your system and establishing a mechanism that will reliably allow the hacker(s) to continue accessing your systems. How long are they likely to be on your system? According to IBM's '*Cost of a Data Breach Report 2023*', hackers spend an average of 204 days in a network before the breach is detected, followed by another 73 days on average to contain the breach. 277 days - just think how much information you could gather about a system in just *one day*, never mind more than *200 days*, not to mention all the little

changes you could make to the system to allow yourself to log back on in the future or if your use of your hacked userid is discovered.

If you think that you are safe because you have a limited number of userids with elevated access rights (such as the ability to update APF data sets or grant SAF accesses), you must consider that there have been multiple bugs and vulnerabilities in IBM and ISV products over the years that would allow an ordinary TSO user to turn themselves into a super user, with the ability to do anything they want. As **Mark Wilson** reminded me, “If I can update an APF data set, the only limitation on what I can do is how good of a programmer I am, and my imagination.” And trust me when I tell you that Mark has a pretty scary imagination!

A customer told me recently that they didn’t need to worry about ransomware impacting their z/OS Db2 data because their applications would start crashing the instant someone encrypted the data they were accessing. However, based on **Robert Catterall’s** ‘*The Various Aspects of Db2 for z/OS Data Encryption*’ blog post, Db2 data set-level encryption is transparent to the application as long as the Db2 database services and system services address spaces (i.e., the Db2 DBM1 and MSTR address spaces) have the required SAF authorities. If someone has the knowledge and the tools to break into your system, it doesn’t seem impossible that they would also have a way to grant the required SAF authorities to Db2, potentially allowing applications to continue accessing encrypted Db2 data sets, transparently to the applications - until they (the hacker) decided it was time to change the key and remove that access.

Someone told me that to *really* protect your system, you have to put yourself in the position of the hacker. Let me give you an example. As part of renewing our home insurance recently, I checked around my house to make sure that it was burglar-proof (short of someone breaking down a door or some other obvious, destructive, means to breaking in). Everything seemed secure, so I was a happy homeowner. Then I accidentally locked myself out of the house on a very cold winter day. It took me a whole 5 minutes to (non-destructively) break back into my own house. Nothing had changed; however, when I was locked out (and freezing!), I looked at my house from the perspective of someone desperately trying to get back in.

To translate that experience into how you would protect your systems - consider that you somehow managed to lock yourself and your colleagues out of your system, and you now have your execs breathing down your neck. Using the collective experience and knowledge of your team, would you have a way to make changes to your system so you could regain access? Shared disks maybe? Or FTPing a job from your PC? Maybe you have a special started task that can issue SAF commands?

Another common trick for hackers (I am told) is to modify the backup process so that your backups are useless. For example, how difficult would it be to modify your full volume dump jobs so they are all backing up the same volume? The jobs would probably take roughly the same time as before, they still produce output files, they still end with RC=0 - everything

seems great until you try to use them to restore your system to some point in the past and you discover that you now have 2743 copies of your JES2 spool volume.

For those that feel that any malevolent actions would be immediately spotted, what would happen if someone damaged a critical data set that is required to IPL your system, but that is not touched again until the next IPL (6 months later)?

Now you are starting to think like a hacker - this is what you could be up against.

Anyway, let's move on to the next stage of the ransomware attack.

Data Collection and Exfiltration

You now have a hacker wandering around your system(s), potentially with the ability to gain access to any resources that look interesting.

They say that 'data is the new oil'. And just as there is a ready market for black market oil (or just about anything else), there is also a very active market for stolen data, meaning the hackers could sell your data to the highest bidder.

Or they might use the fact that they have a copy of your data as additional leverage to get you to pay the ransom, by threatening to make your most sensitive data available to the public. Having your customers' bank transactions available for all to see is probably not the best way to build customer loyalty. It is difficult to know which would be more damaging; the stampede of customers taking their business elsewhere, or the lawsuits for not adequately protecting your clients' data.

I bet you are thinking "if that happened on *my* system, we would immediately notice the additional network traffic". Hopefully; but it is not unusual to see tens of thousands of FTPs running every day. On top of standard FTP you have products like MQ File Transfer, plus all the products that trap database updates and transfer them to your distributed systems. In my experience, the infrastructure staff are aware that this activity is going on, but they rarely are familiar enough with it to detect anomalous behavior.

Ransomware Deployment

The final stage in the attack is where the attacker removes client access to their own data. This can be achieved in many ways, for *example*:

- If the installation has been diligent and careful and has already encrypted all their data, all I need to do as a hacker is encrypt or compromise the keys that you have encrypted the data with. Whoever owns the keys owns the data. Mark and his team spend a lot of time talking to their clients (both those who are *considering* Pervasive Encryption *and*

those who have already done it) about their key management processes and ensuring they are robust. For a hacker, weak key management represents relatively low hanging fruit.

- Another option is to encrypt a data set in place, being careful to give the database manager access to the required SAF resources so it is transparent to the applications using that data (until the hacker changes the key and removes access).
- Another option is that a program could read a record, encrypt it, and write the encrypted record back to the data set. Depending on how heavily used that data set is, the encryption activity might be noticed immediately, or not at all (until the data set is needed for recovery processing).

Obviously there are many other ways that a hacker could remove access to your data, and in a manner that holds out the possibility that access can be regained in return for some form of payment.

If you read the available documentation, there is broad agreement that every installation should have a ransomware policy that determines how you would react in the event of a successful (for the hacker) ransomware attack - do you pay, or do you start trying to recover from your backups?

The topic of ‘immutable backups’ could easily take up an entire article all on its own. If there is sufficient interest and demand, we will cover that in a future Tuning Letter issue.

But for now, we have one last topic before we close this article.

What To Do Now?

The overwhelming majority of material about ransomware is aimed at distributed systems. Nevertheless, there is certainly helpful information out there that also applies to mainframe environments. Apart from anything else, your mainframe almost certainly is not 100% cut off from the outside world, and the first two stages of a ransomware attack are targeted at components outside your z/OS systems. As an individual that is able to logon to your z/OS systems, you should be aware of the threats and the mechanisms that hackers use to try to access customer systems. We have no doubt that our readers are very vigilant and savvy, but we have heard stories about people that are just as smart and careful and responsible as you are being used to allow hackers into their system. And don’t forget that story above, about that poor bank employee that was duped into giving away $25M - how many of us can honestly say that we would not have made the same mistake given the same scenario?

However, coming closer to home, what can you do to protect your beloved z/OS systems? First, deploy MFA for all users, as mentioned above. Next, consider the use of network

activity monitoring and File Integrity Monitoring (FIM), combined with products to detect anomalous behavior.

I think a fantastic idea would be to form two teams of your *most experienced* infrastructure specialists, with one team (the 'Red' team) representing a group of hackers, and the other team representing your department. Grab a meeting room for a day, and see how long it takes for the 'hacker' team to achieve certain objectives, and how long it takes the 'home' team to identify a recovery process. Even better would be if you could use a test sysplex. Give the home team a week to ensure everything is secure (making sure to note anything that had to be changed). Then give the Red team a week to:

- Break into the sysplex using all your z/OS experience, your knowledge of all the backdoors and methods you have used to get back in when you locked yourself out in the past, your programming expertise, and any z/OS hacking tools you can download from the Internet (and no, we are not going to provide a list of those tools!),
- Exfilterate some data,
- Make some changes to SAF accesses,
- Create some backdoors to let themselves back in if they are discovered,
- Encrypt random files, and
- (The most fun part) do *anything* else you can think of that is reversible and would make your systems unusable.

Whoever loses buys dinner for the victors. ☺

I think this would be an enjoyable exercise. More importantly, it would get people *really* thinking about ALL the routes into your system, and all the types of damage someone could do - and hopefully avoid the mistake I made in thinking my home was secure when it really wasn't.

Having gone through that exercise, and hopefully addressing any gaps or exposures that you found, you might consider bringing in an expert penetration tester to see if they can find anything you missed.

Parting Suggestion: Order a copy of a vanilla z/OS Customized Offerings Driver (5751-COD), test and document it, and then lock it away somewhere that there is *zero* possibility of it being altered. In the worst case, where your systems are compromised and you need to stop them immediately before any more damage can be done, this will provide you with a system you can trust to start investigating exactly where you stand and how you can start to recover.

Parting Suggestion 2: If you are ever unfortunate enough to be the victim of a ransomware attack, after the dust settles and everything has returned to normal, you are guaranteed to have a visit from someone that wants to pick your hacked system apart to figure out what the hackers did, and ensure they haven't left any 'gifts' in your system that have not been detected yet. This means that you should not *touch* the compromised environment (SMF data, logs, system volumes, user volumes, *everything*). And that means that you need the ability to restore your backups to a separate set of disks - you do not restore them on top of the (possibly) compromised version of those volumes.

If I was an auditor or a forensic sysprog and I found the compromised volumes had been overwritten, my first two questions would be 1) who authorized that, and 2) what are they trying to hide? I don't think any of us would want to find ourselves in that position.

References

For some Tuning Letter articles we really have to scrape the bottom of the barrel to find relevant and helpful material to list in the References section. This is not one of those articles! We Googled "ransomware" and got 103,000,000 hits. Hmm - apparently this ransomware stuff is not a secret..☺. So we narrowed it down by adding "mainframe" - that reduced the number of hits to 'only' 184,000. Still a lot of reading, even for the most effective speed readers. So, the following are the documents or websites that we found to be particularly informative:

- Reddit discussion about whether there have been any cases of successful ransomware attacks on a mainframe. Interestingly, this was the very top of Google's list of 184,000 hits on ransomware and mainframe. This discussion has links to many other Internet sites that discuss ransomware specifically in relation to mainframes. It makes for quite scary reading, but I would consider this to be required reading for whoever is responsible for worrying about ransomware attacks on your z/OS systems.
- US CISA '*Ransomware Guide*'.
- List of US government websites that provide various useful information about ransomware: https://www.cisa.gov/stopransomware/general-information
- SHARE in San Jose 2017 presentation '*Ransomware on the Mainframe Checkmate!',* by **Chad Rikansrud**, hosted by NewEra Software. Plus, here is a link to a YouTube recording of Chad's session: https://www.youtube.com/watch?app=desktop&v=i-DbTy3bEj8. And you can find Mark Wilson's SHARE 2024 version of the presentation *here*.
- Techchannel article, '*The Complete Beginners' Guide to Hacking a Mainframe*', by **Trevor Eddolls**.

- IBM Paper '*Definitive Guide to ransomware 2023*'.
- IBM Redbook, *Getting Started with z/OS Data Set Encryption*, SG24-8410.
- IBM '*Ransomware protection solutions*' website.
- IBM Ransomware *hub website*.
- Short article by **Paul F. Renda** (from December 2021) about why ransomware could be a real risk in the near future.
- GSE UK Conference 2020 Session '*Mainframe Security Combating Ransomware With File Integrity Monitoring'* by ***Al Saurette***.

Summary

We would love to be able to dedicate every Tuning Letter issue to performance tips and information, with maybe some relevant new release or new function information thrown in for some variety. However, we have a responsibility to our readers to draw their attention to topics that are critical to the mainframe platform or their organizations. Ransomware is one of those topics.

Based on conversations with other experienced mainframers, the ones that have not looked closely at ransomware or even given it much thought are content in their belief that their systems are secure, and ransomware is something they don't need to concern themselves with. I was one of that group.

Then I started researching this article and speaking to people like Mark that know infinitely more about this topic than I do. And now I am firmly in the *other* group. I have children and family that I love, so the little knowledge on this topic that I now have terrifies me. The understandable, but false, sense of security that many of my mainframe friends and peers have also really scares me. I really wish that Ransomware *was* a topic that we could categorize as 'a distributed system issue' and then ignore, but everything I've seen leads me to believe that is not the case.

This article only scratches the surface of this important topic. Protecting yourself from a ransomware attack, and certainly from stages 1 and 2 of that attack, is far outside our field of expertise. Preparing yourself for recovery from a ransomware attack is a little closer to home, and is a topic we will likely address in a future Tuning Letter issue.

For now, we hope this article is effective in encouraging you to do some research of your own, to enable you to make a more informed decision about how concerned you should be about a ransomware attack on *your* z/OS systems. At a minimum, we hope this article will allow you to actively participate in a ransomware discussion with your peers from other platforms.

Finally, I want to thank **Mark Wilson** for his invaluable help and patience in the creation of this article. If you find it helpful, you have Mark to thank for that.

User Experiences and Tips

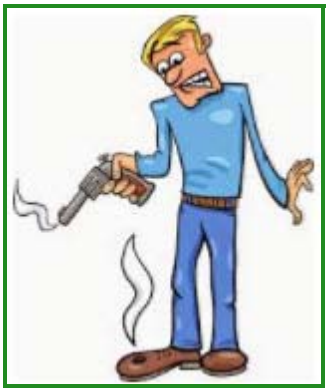

The User Experiences section has always been one of the most popular sections of the Tuning Letter. After all, who doesn't like to learn from other people's experiences, be they good or bad? Also, from time to time, we come across useful tidbits of information that we believe would be of value to our readers, so we like to pass those on as well. In this issue, we have two interesting and valuable items for you:

Synchronizing Your Type 23 and Type 7X SMF Records

Following on from our Planning for an Upgrade to z16 article in Tuning Letter 2022 No. 4, we've been looking at the metrics some of our readers use to compare the CPC capacity before and after some event.

Most sites, and IBM's ITRR calculation, use a variety of the RMF/CMF-produced type 7x SMF records, with some also including information from the type 113 CPU MF records. But one site told us that they also use information from the type 23 (SMF Status) records. I had looked at specific fields in the type 23 records before, but never at the full set of fields. Apart from the obvious information about SMF buffer use and activity to the SMF logstreams, it also contains information about system-wide I/O rates, the uncompressed and zEDC-compressed number of bytes written to each SMF logstream, the total number of TCB and SRB dispatches, plus more. The customer we spoke to uses the number of TCB and SRB dispatches as an indicator of the volume of work being done on the system, which seems eminently sensible.

Based on their input, we did a little more work on this concept to determine how strong is the correlation between the number of dispatches and the total CPU consumption. However, we immediately hit an inhibitor - the default interval for the type 23 records is one hour, and they rarely align with the type 7x records because by default they are aligned with the IPL time. Looking at data from multiple readers, it appears that most sites use the default settings. However, all that needs to be done to get the creation and duration of the type 23 records to align with the RMF/CMF records is to specify the following in your SMFPRMxx member:

```
STATUS(SMF,SYNC)
```

As it happens, while working with Todd on the "Making Sense of the Many I/O Count Fields in SMF" article I noticed that the type 70 records also contain the number of TCB and SRB dispatches (SMF70TCB and SMF70SRB accordingly). As a result it is not necessary to process the type 23 records to get this information. But there is still other useful information in the type 23 records, and it helps to have that aligned with the corresponding type 7x data, so we still recommend including that STATUS(SMF,SYNC) statement in your SMFPRMxx.

Speaking of SMF Records and zEDC...

There has been some discussion on IBM-MAIN recently about the use of zEDC. Our first in-depth article about zEDC was titled 'zEDC', back in *Tuning Letter 2014 No. 2* - nearly 10 years ago. At that time, the zEDC function was provided by PCIE-attached zEDC cards. On z15 and z16, the zEDC function is provided by zEDC Accelerators on each chip.

We had some measurements back in that original zEDC article, comparing disk space, CPU time, and elapsed time for no compression, traditional SMS compression, and zEDC. However, we have not had a complete set of zEDC measurements since the zEDC card was replaced by the zEDC Accelerator in the z15 generation of CPCs. Given the comments on IBM-MAIN, we thought a brief reminder would be valuable.

There are many different types of data sets on a z/OS system, some of which support zEDC and some that don't. Unfortunately, zEDC is not supported by every data set type, nor does it support the exact same set of data set types as the older SMS-provided compression. On the other hand, it *is* supported for SMF data in a log stream, it is supported for use with JES2 spool files, and it is supported by PDSE members. The full list of IBM zEDC exploiters (including information about which ones require the chargeable zEDC feature on z/OS, and which do not) is shown in Figure 8 on page 36, courtesy of **Marna Walle** (you can also see this slide on Marna's LinkedIn page).

Figure 8 - zEDC Exploiters (© IBM Corporation)

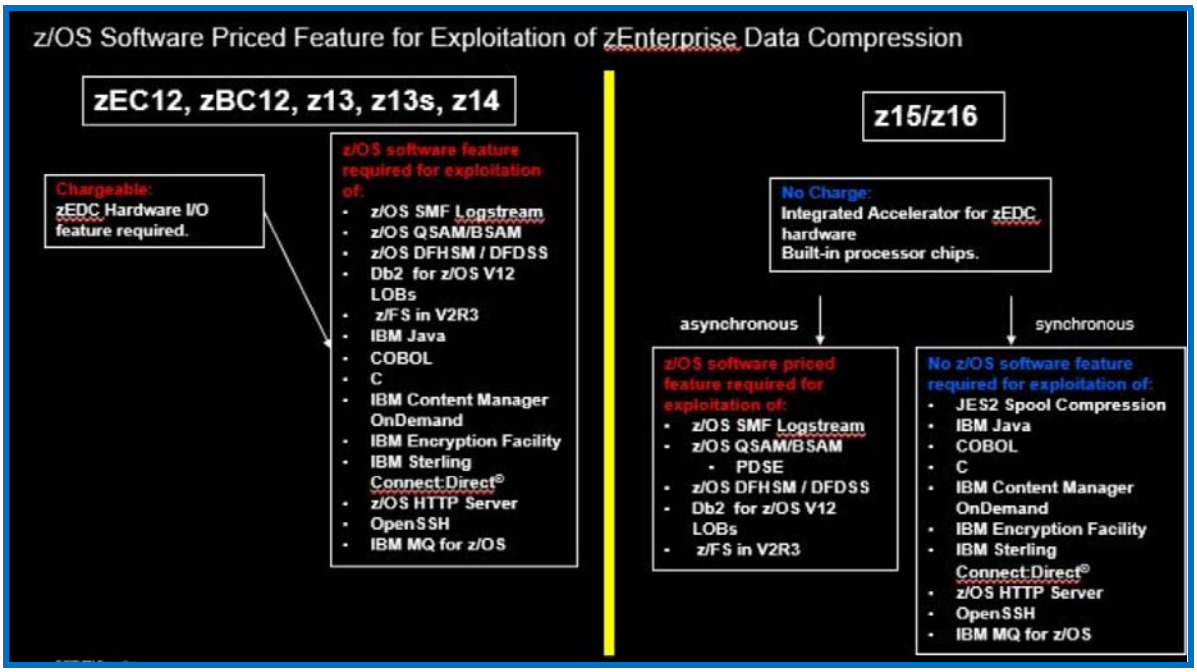

For now, we just want to concentrate on the use of zEDC for SMF data. If you were following the IBM-MAIN discussion about zEDC, you should have seen a post by **Andrew Rowley** from Black Hill Software in Australia. He ran a number of measurements using:

- ◆ No compression
- ◆ CICS software compression
- ◆ CICS and zEDC compression
- ◆ zEDC compression

He also included the very interesting results of measurements using Java to read the data (rather than IFASMFDP). I don't want to steal Andrew's thunder, so I will simply recommend that you have a read of his findings for yourself.

The tie-in between this and the type 23 records is that the SMF23BBC field reports the number of bytes written to the associated logstream *before* compression, and field SMF23BAC reports the number of bytes *after* they were compressed by zEDC. These two fields allow you to very quickly and easily see the compression ratio achieved for each of your SMF logstreams. Note that the compression ratio for the offloaded SMF data (in sequential disk data sets) is likely to be different because SMF and QSAM use zEDC in slightly different ways.

Many thanks to Andrew for sharing the very interesting results of his measurements.

Sending SMS Texts from z/OS

Are you fed up hearing your beloved z/OS platform being described as 'legacy', 'obsolete', or 'old tech'? Well, we are too. So you can imagine our elation when we read Andrew Rowley's article on his website about sending SMS texts directly from z/OS, based on SMF data obtained via the Real Time SMF interface (that was described in the 'SMF Streaming' article in *Tuning Letter 2016 No. 2*).

Andrew's article uses an example of gathering information in real time from SMF type 30 step end and job end records, summarizing the information for the job, and sending an SMS notification (using the Twilio service) in case of a non-zero return code. No SMF exits or authorized code are required (because the program uses the SMF Real Time Interface), and minimal GCP MIPS are consumed because the programs are written in Java and therefore run on zIIP engines.

Getting the programs up and running and customized for your environment would certainly require more Java programming skills than *I* possess. However, Andrew's EasySMF:JE product (available for a 30-day free trial) takes all the wizardry out of navigating your way to the fields you are interested in, and his working samples would get anyone with decent Java programming skills up and running and receiving SMS texts before you know it.

So there! Who says the mainframe isn't trendy? What's that you say? SMS messaging has been around since 1992? Well, that's *reasonably* new, in MVS terms, isn't it? What about Instagram support? And z/OS selfies? Sheesh, there's just no pleasing some people... ;)

zPCR Tip

We did some work recently with a customer that had quite a large number of Vertical Low logical CPs across their many non-production LPARs - certainly more in many cases than the two VL CPs/LPAR suggested by **Kathy Walsh** in her '*Number of Logical CPs Defined for an LPAR*' white paper. Their logical-to-physical ratio was quite high, resulting in LPAR overhead as high as 10% at times.

According to Kathy's paper, and subsequent discussions with several IBM z performance experts, the major reasons for trying to minimize the number of Vertical Low CPs are:

- The higher the ratio of logical-to-physical CPs, the more work PR/SM has to do to manage all those logical CPs ('LPAR Management' time).
- The more UNparked CPs/zIIPs an LPAR has, the higher is the MVS management cost – additional control blocks, longer control block chains to be scanned/managed, potentially more inter-CP locking and latching. Note this is MVS management time, not LPAR Management time.

- Unparked Vertical Low CPs/zIIPs share the processor share of the LPAR's Vertical Medium (VM) CPs/zIIPs, so unparking a VL reduces the share of the VM(s). And the more unparked VLs you have, the lower is the share (i.e. relative priority) of each one. (For more info, see this excellent '*WLM and Hiperdispatch*' presentation by **Robert Vaupel**).

- VL CPs/zIIPs have demonstrably poorer Relative Nest Intensity (RNI) and Cycles Per Instruction (CPI) values – just look in your own SMF 113 records if you want to verify that.

- PR/SM dispatches VL CPs/zIIPs on cores belonging to another LPAR's VH CPs/zIIPs, polluting the caches of those VH CPs/zIIPs. Sadly, there is no metric to easily quantify this.

Based on all these good reasons, one of the actions we suggested was to decrease the number of Vertical Low CPs to a number closer to that required to deliver the maximum capacity used by those (test) LPARs.

Naturally, before making widespread changes, they wanted to know what sort of savings could potentially be achieved. I didn't have a quick answer to that question, however I recalled noticing some recent changes in IBM's zPCR tool related to Parked engines.

My initial expectation (which turned out to be incorrect) was that enabling the "Remove Parked LCPs from the LCP count when copying partitions into zPCR" option (as highlighted by the lower arrow in Figure 9 on page 39) would get zPCR to model the effect of reducing the number of LCPs configured in the LPAR by the number of LCPs that were parked in the selected interval.

Figure 9 - zPCR Create LPAR Configuration from EDF screen (© IBM Corp.)

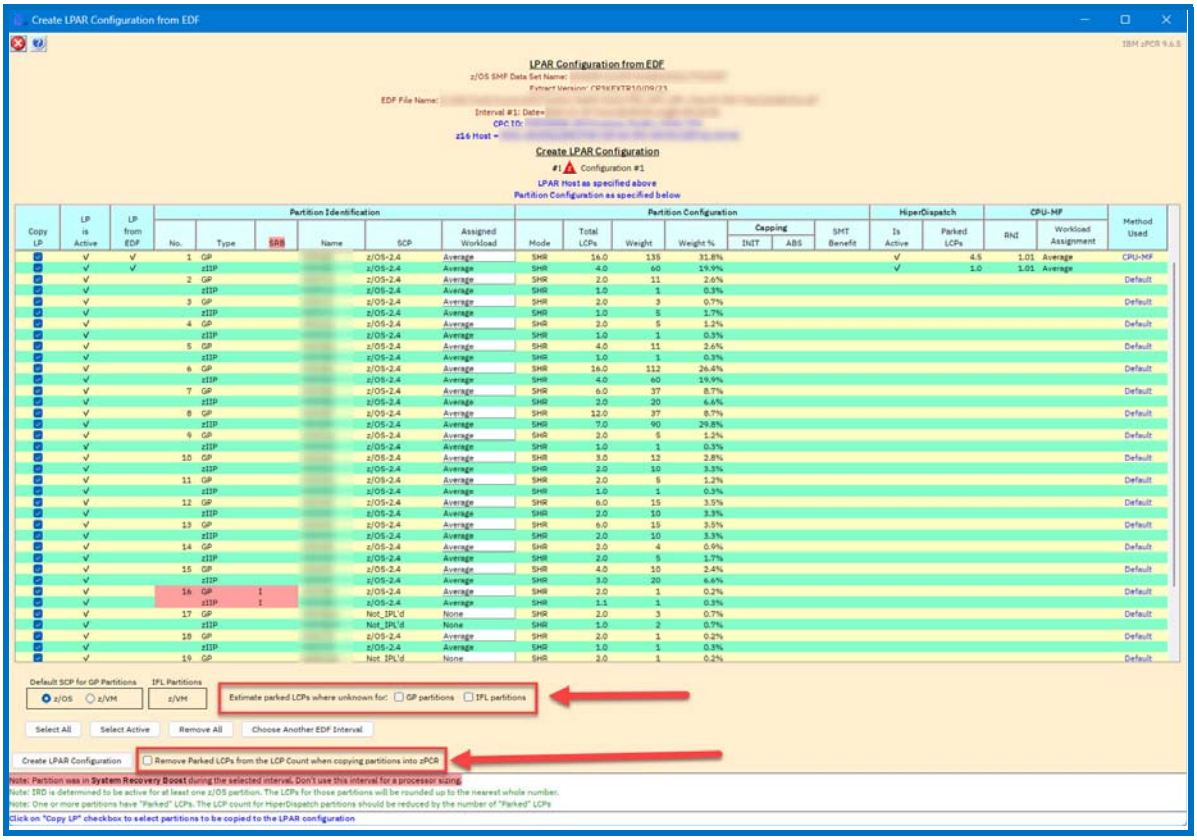

However, I was lucky enough to have a chat with zPCR experts **Brad Snyder** and **Shawn Lundvall** at SHARE. They explained to me that, *by default*, zPCR treats parked engines in the same way as every other LCP when calculating the CPC capacity (that is, it ignores the fact that they were parked). Selecting the “Remove Parked LCPs....” option tells zPCR that it should treat parked engines as parked engines. In other words, to get a more accurate estimate of the current CPC capacity, make sure that you select that option on the zPCR “Create LPAR Configuration from EDF” screen after you have selected the interval you wish to use.

One other important point while on the topic of zPCR and VL LCPs. zPCR only sees information about which engines are parked, and how long they are parked for, for a given LPAR if you include the SMF type 70 records from that system. For the other systems on that CPC, it can only guess at those numbers. If you do *not* select the check boxes on the “Estimate parked LCPs where unknown for ...” line (as highlighted by the upper arrow in Figure 9) it again assumes that *all* LCPs for those LPARs were unparked for the entire interval. If you would like it to attempt to treat the parked LCPs more appropriately, you must select those two check boxes.

If you think selecting these check boxes will have little or no impact in your configuration, try selecting them and see how large is the difference between the two estimates. If it was possible to specify a default behavior, based on my experience, I would want the default to be to always *select* those check boxes.

Introduction to Altair SLC

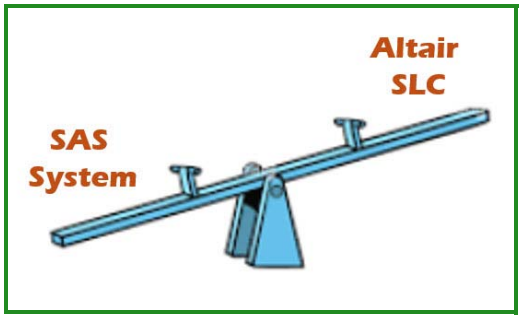

Many of our readers, both large and small, have a big investment in SAS language programs. Products such as Merrill Associate's MXG, Broadcom's MICS, and our own BoxScore product are written in SAS language, and many companies have their own, or bought-in, business applications written in SAS language. Unfortunately, the SAS System frequently appears high on the hit list of many of our cost-cutting colleagues in Finance. This places you in a difficult position - if some applications written in SAS language are moved to other platforms, the effective cost per user for the remaining users increases, making it difficult to justify keeping a SAS System license on z/OS. On the other hand, removing this much-loved programming language is not going to win you any popularity contests with your colleagues, especially if they have to rewrite their home-grown SAS language applications.

In 2003, a UK-based company called World Programming created a product called WPS ('World Programming System') that ran programs written in the SAS language, but did not require customers to have a SAS System license. In 2021, Altair, an established global company in the engineering software and services industry, bought World Programming and renamed WPS to 'Altair SLC'.

According to the Altair website, switching to Altair SLC typically results in cost savings of between 50% and 70%.

This article is not in any way going to get down into the gory details of the SAS language or the SAS System or Altair SLC. Rather, our objective is simply to increase the awareness of this alternative to the SAS System among our readers, and to point out that Altair SLC *might* offer you a cost-effective way to keep both your SAS language users *and* your accountants happy.

Target Audience

This article is potentially of interest to any site that either uses the SAS language (on any platform, but naturally we are focused on z/OS) or that considered it in the past but were put off by the cost.

This article is not going to discuss the fine details of migrating from the SAS System to Altair SLC. Rather, its purpose is simply to make you aware of the Altair SLC product and to relay the results of our initial research, hopefully saving you a little time.

I want to thank my colleague **Scott Barry** and Altair's **Oliver Robinson** for all their help and patience in handling my many questions about both SAS and Altair SLC. This article wouldn't exist without their generous support.

Introduction

Two popular SMF analysis tools, MXG and MICS, are written in the SAS language. Early in her illustrious career, Cheryl worked on MICS when it was owned by Morino Associates. A little later, Cheryl wrote BoxScore in ... you guessed it, the SAS language. With that history, you can understand why we have a keen interest in what is happening in the 'SAS world'.

Terminology

Because people often use the term 'SAS' to refer to different things depending on the context, we will try to be precise with our phrasing in this article to avoid any potential confusion.

SAS Language - The *SAS programming language* is just like any other programming language (COBOL, PL/I, Java, and so on) in that it has functions and verbs and syntax rules and so on. Anyone could use any simple editor to *write* a program in the SAS language, just as you could with COBOL, Java, and so on. However, like those other programming languages, you need something to *execute* your program.

In this article, we use the term "SAS language program" to mean a customer or vendor-created program written in the SAS language. We try to avoid the term "SAS program" because that could be taken to mean a program provided as part of the SAS System or a program provided by the SAS Institute.

SAS System - The *SAS System* provides an environment to interpret SAS language programs, provide services to those programs, manage SAS-format data sets, and so on. While anyone *could* write a SAS program, you need a license for the SAS System to be able to *execute* that program on z/OS.

SAS Institute - The company that owns the SAS System software and many other products that are based on the SAS System.

Questions and Answers

Most customers that are interested in using Altair SLC on z/OS are probably using SAS on z/OS today, and therefore will want to know how Altair SLC compares to the SAS System. On that basis, we felt that the most appropriate format for this article would be a Q&A one. We hope you find this useful.

Exactly what *is* Altair SLC?

Altair SLC provides a runtime environment that is conceptually similar to the SAS System. It is *not* a tool to convert SAS language programs into some other programming language - Altair SLC is able to read and execute programs written in the SAS language. Altair SLC can run SAS language programs with commonly used syntax, including DATA steps, procedures, macros, and so on.

Altair SLC also offers a user interface called Altair Analytics Workbench that includes coding tools and low-code/no-code tools for working with data. The Workbench is an IDE/GUI for use on Windows, Mac, or Linux desktop. Altair also offers a TSO interface for z/OS called Server Explorer.

There is a third component called SLC Hub. SLC Hub is a separate piece of software that works with SLC using the same Altair Units licensing mechanism. SLC Hub is typically useful for medium to large enterprise deployments. There is currently no z/OS or TSO version of SLC Hub. We don't believe that you *need* to use SLC Hub, however, it does provide a means to centralize management of a site's users, SLC/SAS language code, plus various Altair SLC/SAS-language (environment) objects, like databases, code-pieces, user-access, and site administration functions such as load balancing, governance, scheduling, provisioning, and security.

This means that no new programming skills are required - if you have been writing SAS language programs for the last 20 years, you can continue using your SAS skills and experience to maintain your existing programs and create new ones.

Altair SLC runs on z/OS on the mainframe, as well as under Windows, Linux, MacOS, and AIX on distributed systems.

Does Altair SLC support 100% of SAS language programs?

You might view a migration from the SAS System software to Altair SLC to be a little like migrating from an old version of COBOL to the latest release. Just as you would not recompile *all* your old COBOL programs and run them in production without testing them, similarly you would have a migration project to move your SAS applications over to Altair SLC.

There *is* some old SAS language syntax that Altair SLC does not support, and there is new syntax that SLC *does* support. Support for SAS language syntax support is based on customer requests and Altair SLC development priorities.

Altair provides professional services to assist with the evaluation and migration to Altair SLC, as well as tools that will scan your SAS language programs code (from any platform) and highlight any statements that might cause an issue during migration. While most common SAS language statements and constructs are supported, it is possible that you might have

some SAS language programs that would require work to get them to function in an equivalent manner with Altair SLC.

The code analyzer mostly performs syntax verification. It doesn't really exercise the Altair SLC vs SAS System operational processing, however, that testing would be part of the normal migration process, just as you would perform when moving to a new version of COBOL.

In Altair's experience, the percentage of programs with unsupported syntax is often zero or very small for mainframe users. The Altair team have years of experience with running SAS language programs on z/OS and have worked hard to support this set of syntax. Distributed (off-mainframe) workloads tend to be more likely to have unusual or advanced syntax that the Altair team may not have come across before.

SAS provides its own types of files - are all those file types supported by SLC?

Altair SLC can read SAS proprietary-format databases ("sas7bdat" - MXG PDBs, for example). It can also read, and in many cases write, many other SAS data file formats. There are some ancillary "object library" formats which are not really data-related, but instead are used with SAS operation/execution – these object libraries are generally not compatible with Altair SLC.

Altair SLC provides its own database format (for READ and WRITE). There are also SLC-provided robust built-in procedures for install, setup/configuration, backup, and recovery, equivalent to what the SAS System provides to its customers.

SLC stands for SAS Language Compiler. Does that mean that you must compile your SAS language programs rather than running them interpretatively?

Both the SAS System and Altair SLC support pre-compiling SAS language programs into an object library. In our experience, most customers run their SAS programs interpretively, meaning that the compile-and-execute all happens in one step, conceptually a little like Java. This is a very efficient process, meaning that there is generally little benefit to be had from pre-compiling. Both the SAS System and Altair SLC support running SAS language programs interpretively.

How transparent to users and developers would a switch to SLC be?

There are minor differences between the JCL statements used, but otherwise the two environments are very similar looking. From a SAS language program developer's

perspective, there are minor differences and unique characteristics with a SAS System to Altair SLC migration, but they are easily adapted to.

What about vendor products written in the SAS language?

There are many vendor applications written in the SAS language, some of which are provided by the SAS Institute, and some by other vendors. There is no comprehensive, up-to-date table showing exactly which releases of which SAS-based applications are supported with Altair SLC.

As you would expect, SAS Institute offer many products based on the SAS language, and Altair do not necessarily offer an alternative to every one of them. However, the Altair team *are* aware of customers replacing the following SAS Institute products with the corresponding Altair product:

- Base SAS
- SAS/STAT
- SAS/AF
- SAS/FSP
- SAS/GRAPH
- SAS/ETS
- SAS/OR
- SAS/QC
- SAS/CONNECT
- SAS/Access (for many databases)
- Enterprise Guide
- Enterprise Miner
- DI Studio
- Metadata server
- Stored processes

If you have a SAS-based application that is not in the list above and are interested in migrating to Altair SLC, contact Altair.

Does it work with MXG?

Yes, MXG works with Altair SLC, and such a configuration is supported by Merrill Associates. There is just one version of MXG that covers both SAS System and Altair SLC customers, however, the MXG code *does* detect when operating in an Altair SLC environment and adjusts its behavior, where needed, to adapt to subtle SAS System vs Altair SLC differences. Altair (and before that, WPL) have worked with my colleague Scott Barry, with Merrill Associates, and with MXG customers to ensure MXG functions correctly with Altair SLC. Many Altair SLC mainframe customers are using it with MXG.

Does it work with BoxScore?

Watson & Walker's BoxScore has been tested/exercised with Altair SLC. All supported functions were performed successfully and produced the expected output/outcome, so we (Watson & Walker) support BoxScore with both the SAS System and Altair SLC.

Does Altair work with MICS?

According to the Broadcom website:

> *"MICS has not been certified on the WPL/WPS platform and there are no plans in the product roadmap to do so."*

That statement is from 2022 and we are not aware of any changes to Broadcom's strategy in this regard since then.

Can the SAS System and Altair SLC co-exist on the same z/OS?

Yes, that would typically be the case during the period when you are migrating to Altair SLC.

Can you intermix Altair SLC on z/OS with SAS applications on other platforms?

Yes, to some extent. For example, it is possible to have a SAS application using the SAS System on a PC talking to SLC on z/OS using RSUBMIT, PROC UPLOAD/DOWNLOAD, and so on, with a few limitations. This interoperability can be useful in migration scenarios.

The SAS System does not exploit zIIPs. What about Altair SLC?

SAS language programs do not exploit zIIPs, regardless of whether the program is running under the SAS System or Altair SLC.

Are there any benchmarks to compare the performance of Altair SLC to SAS System?

The Altair team have not performed any benchmarks. However, they have put a huge amount of effort into performance optimization over the years. Feedback from customers suggests that performance is often comparable, although these are different pieces of software and performance can vary depending on workload and environment. Altair is happy to address performance issues on a case-by-case basis with customers.

Does Altair SLC provide metrics or performance analysis tools?

The SAS System optionally creates an SMF record to be written for every PROC or DATA step containing resource usage statistics for CPU time, memory, and EXCP count. Altair SLC currently does not produce SMF records, but we believe they would be open to considering such a request for enhancement.

Both SAS System and Altair SLC provide configuration settings that can help with increasing the performance measurement data points that are rendered in the SAS or Altair SLC log output – this could be used for a head-to-head performance comparison.

What about licensing models on z/OS?

Altair SLC supports full-capacity, sub-capacity, and usage-based pricing. SLC on z/OS does not track system size/MSU capacity, CPUID, and so on. Altair has a flexible approach to z/OS licensing and is happy to work with customers to select a sizing/scope mechanism that works for the customer.

Additionally, Altair offers free trials and Proofs of Concept for Altair SLC on z/OS, and a free community edition for Windows and MacOS workstations.

Relative Costs: According to the Altair website, switching to Altair SLC typically results in cost savings of between 50% and 70%.

Status of WPL/SAS Litigation?

Weren't there restrictions on WPS sales in the US?

Some years back there was a court case in the US between SAS Institute and World Programming. During that time, WPS was still available to, and used by, existing US-based customers, but WPS was not marketed to new customers in the US. All litigation between SAS Institute and World Programming or Altair is now fully resolved and there are no restrictions on licensing Altair SLC software anywhere in the world. For more information see this news release.

References

The following websites provide a little additional information that might be helpful.

- https://altair.com/resourcelibrary/?category=Customer%20Stories provides links to examples of many different users of Altair products. To get information about the use of Altair SLC on z/OS, go to this page then search on ‘mainframe’.
- Short (4-page) document describing the capabilities and advantages of Altair SLC and the migration process for customers moving from the SAS System: https://altair.com/docs/default-source/resource-library/da_print_technicaldocument_altairslc_letter.pdf
- Home page for information about using SLC to run SAS language programs, and links to related Altair web pages: https://altair.com/run-sas-programs-on-mainframes
- There are also many other documents (such as a Migration Guide for Altair SLC on z/OS) about Altair SLC on the Documentation link on the Altair Community website. To see them go to the Community website, then click on Documentation (up near the top of the page) and then filter on ‘Altair SLC’.

Summary

Anyone that has been around IT in general, and the mainframe in particular, will know that the right answer to any question is “It depends”. We believe that we have helped our readers save millions of dollars by encouraging them to *consider* sub-cap CPCs when looking at their upgrade options. That doesn’t mean that we believe that a sub-cap CPC is the right answer for every configuration. But we *do* believe that our readers should make fully-informed decisions. In the case of CPCs, that means that sites should have *evaluated* whether a sub-cap CPC would be a good fit for them, and factored that information into their final decision.

For readers that are under pressure to reduce or eliminate the costs of running SAS programs on z/OS, we feel that Altair SLC should be included in their list of options. For some sites, it might be a perfect fit. For other sites, it might not. But we want to help our readers make an informed decision about the best way to proceed - hence this article about Altair SLC.

We want to thank **Oliver Robinson** in Altair, and especially my colleague **Scott Barry**, for their invaluable assistance with this article. We hope their time and patience will pay off in helping you make the right decision for your organization.

Making Sense of the Many I/O Count Fields in SMF

Regular Tuning Letter readers eagerly await the latest article from our friend, IntelliMagic's **Todd Havekost**. Todd is an award-winning presenter with a wealth of real world experience and a boundless enthusiasm for anything related to z/OS performance. I've been working with Todd on and off since 2015, and he still amazes me with his knowledge, his organization, and his attention to detail. We couldn't ask for a better contributor and workmate.

In this Tuning Letter issue, Todd addresses a question that has long vexed me (and many of our readers) - why do different fields in SMF records that all sound like they are reporting the same thing (I/O rates) have different values? **Kathy Walsh's** excellent *How to Measure that New z15* presentation discusses a variety of methods for comparing CPC capacity, many of which use I/O rate as a key part of the calculation. However, with different I/O rate-related fields reporting different values, which one should you use? And *why* do they contain different values? If these questions resonate with you, this is the article you have been waiting for.

Target Audience

This article is intended for *anyone* that works with SMF data, whether you use it for performance reporting, capacity planning, understanding system behavior, or charging for use of system resources.

Additionally, if you are still new to SMF data, Todd's article provides an *excellent* example of how you can use SMF data to deepen your understanding of what is happening in z/OS - critical for anyone that hopes to one day become the go-to performance guru in your site.

I want to thank Todd for his outstanding work on this article. This article wouldn't have been possible without Todd's vast experience and his access to real customer SMF data. We also got great support from IBM's Data Gatherer, DFSMS, and IOS Development teams - thank you to all of them for their patience and enthusiastic assistance.

Introduction

One common way to quantify and characterize workloads is to use measurements of I/O activity. This is illustrated by the fact that one metric often relied upon when calculating

Internal Throughput Rates to measure the impact of processor upgrades is I/O interrupt rate divided by CPU consumption. In the methodology described in the **Kathy Walsh** "*How to Measure that New z15*" presentation, one of the two metrics in the "sweet spot" (combining higher accuracy and lower effort) on her chart about maximizing accuracy and minimizing data collection effort is IRATE/GCP USED, i.e., total rate of I/O interrupts divided by consumed CPU capacity expressed in terms of CPC cores.

However, there are a plethora of SMF metrics that report I/O activity across numerous record types including SMF type 70, 72, 74, 78, 42, 23, 30, 6x (VSAM), 92 (zFS), Db2 SMF records, and probably others. We don't have the space or time to cover them all in this article, so we will limit ourselves to the most widely-used ones (types 7x and 42) here.

Many of the fields we will be looking at have similar descriptions in the SMF manual, such as "total number of I/Os" or "start subchannel count", however when you look at those fields for the same interval they can contain significantly different values.

To assist with our understanding of the counts, we analyzed the values of these metrics from multiple systems in two sites and listed the fields in relative order of the values. This resulting order was consistent across the two sites, confirming that the apparent disparities were not an anomaly. This approach also helped us identify the fields that appeared to be reporting *consistent* values, and the ones reporting significantly *different* values (indicating that different categories of I/Os were being included or excluded from the different fields).

The categories of I/Os we have identified so far include:

- Traditional non-zHPF Disk I/Os.
- zHPF Disk I/Os.
- Synchronous I/Os (via zHyperLink).
- Paging Disk I/Os.
- zHyperWrite I/Os *to* a Secondary Device, and Consistent Reads *from* a Secondary. These may appear a little different to 'normal' I/Os because the device is not online to the reporting system. However, for the SMF fields we investigated for *this* article, these I/Os are treated the same as 'normal' I/Os.
- "Other" I/Os (Tape, Communication Devices, Graphics Devices, Unit Record Devices, and Character Reader Devices).
- 'I/Os' to PCIE-attached devices such as OSA adapters, zEDC cards (prior to z15), PCIe Cryptographic Co-processors, RoCE Express cards, and so on.

The values shown in Figure 10 on page 50 and throughout the remainder of this article represent cumulative rates across multiple CPCs from one of the sites over a one-day time interval.

Figure 10 - I/O Count Fields from SMF 42, 70, 72, 74, and 78 Records

Row	Record	Field	Value
1	SMF 70.1	SMF70SLH + TPI	394366
2	SMF 78.3	R783IPII	390971
3	SMF 74.1	SMF74SSC + SQR	338183
4	SMF 42.n	All I/Os w/sync	318113
5	SMF 42.n	Captured I/Os w/sync	318086
6	SMF 70.1	SMF70NIO	259662
7	SMF 78.3	R783IIFS	254958
8	SMF 74.1	SMF74SSC	235418
9	SMF 42.n	All I/Os less sync	215346
10	SMF 42.n	Capt I/Os less sync	215319
11	SMF 72.3	R723CIRC	192165
12	SMF 74.1	SMF74SQR	102765
13	SMF 74.1	SMF74SQW	0
14	SMF 42.n	S42SNROS	102767
15	SMF 42.n	S42VSION	27

As you can see, there are significant differences between the different metrics. This immediately raises the question of - when someone asks you for the “system-level I/O rate”, *which* of the more than 10 I/O-related fields are they referring to? In this example, the total I/O interrupt rate was nearly 52% more than the number of I/Os reported in the SMF70NIO field, and more than double the rate reported in the R723CIRC field in the SMF 72.3 records, so which of those fields should you use? The remainder of this article steps through those fields and explains, as best we are able to, why they contain significantly different values.

SMF 70 Fields

We will begin with the SMF 70.1 record, because that is a core record for many types of analysis, and because it contains the fields used in the IBM processor capacity evaluation methodology mentioned above. Most sections of SMF 70 records are generated at the z/OS system level, though some sections are at other levels (e.g., all LPARs on that CPC, logical CP).

The first fields mentioned in Figure 10 are SMF70SLH and SMF70TPI. SMF70SLH reports the number of times the Second Level Interrupt Handler (SLIH) was entered. This information is retrieved from field PCCASLIH in the PCCA (Physical Configuration Communication Area) control block of each logical processor[3]. If you are not familiar with

[3] The PCCA is documented in the *z/OS MVS Data Areas Volume 3 (ITK - RQE)* manual.

MVS interrupt handling and would like to know more, the *1.17.2 Types of interrupts* and *1.34 I/O Interrupt processing* sections of the IBM Redbook *ABCs of z/OS System Programming Volume 10*, SG24-6990, will tell you all you want to know.

The SMF70TPI field reports the number of Test Pending Interruption (TPI) instructions. The SLIH always issues a TPI instruction just before it completes, to determine if there are more interrupts waiting to be processed. If the TPI ends with completion code 1, that means there is at least one more interrupt queued, so rather than ending, the SLIH immediately retrieves that interrupt and starts processing it (*without* incrementing the SMF70SLH counter). This is illustrated nicely in the chart in Figure 11 from **Robert Vaupel's** outstanding *High Availability and Scalability of Mainframe Environments* textbook. The SMF70TPI value is retrieved from the PCCASTPI field in the PCCA.

Figure 11 - MVS Interrupt Processing (© Robert Vaupel)

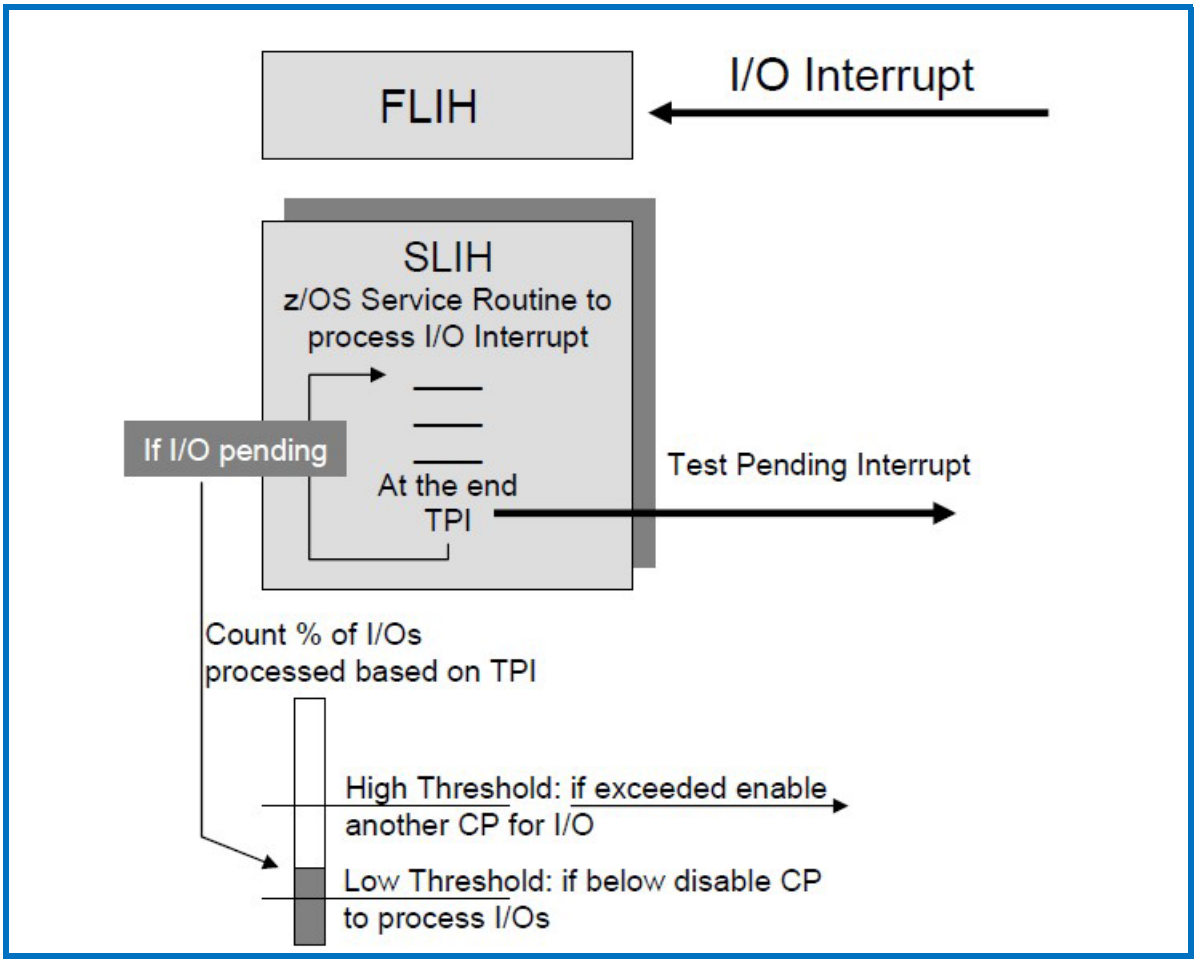

As a result, to determine the *total* number of I/O interrupts processed by a z/OS system, you need to sum the SMF70SLH and SMF70TPI fields. The sum of these provides the IRATE component used in the IBM IRATE/GCP USED formula referenced above.

IBM ITRR calculation uses SMF70SLH and SMF70TPI values.

Synchronous I/Os

In late 2017, IBM introduced zHyperLink technology that enables qualifying I/Os for supported (short distance point-to-point) configurations to be completed up to ten times faster than “standard” (zHPF, z High Performance FICON) I/Os (in under 30 microseconds, depending on the distance). These I/Os are completed synchronously, maintaining control of the CPU and not resulting in an I/O interrupt. Thus, synchronous I/Os are not reflected in any measurements of I/O interrupts such as those in SMF70SLH and SMF70TPI. Additionally, as you will see, synchronous I/Os are included in *some* other SMF fields that report I/O counts, but not in others.

In keeping with the best IT traditions, there are (at least) two completely unrelated uses of the term 'Synchronous I/O':

- There are the I/Os issued on a zHyperLink (those are the ones we mention in this article).
- Db2 also uses the term for I/Os that result when a Getpage cannot be satisfied from a buffer pool, resulting in an I/O to the database that is ‘synchronous’ to the unit of work.

When we talk about Synchronous I/Os *in this article*, we will always be referring to I/Os issued over a zHyperLink. From the perspective of this article, the database reads that Db2 calls 'Synchronous I/Os' are the same as any other disk I/O, and we treat them accordingly.

For more information about Synchronous I/Os and zHyperLink, refer to the 'Meet the Future of I/O - zHyperLink' article in *Tuning Letter 2018 No. 2*.

Another interesting field from the SMF type 70 records is SMF70NIO which is described in the SMF manual as “the number of I/Os for this CPU.” It is based on information from an Object Code Only (OCO) IOS control block.

Figure 12 - I/O Count Fields from SMF 70 Records

Row	Record	Field	Value	Level	Source	SMF manual text / Formula
	SMF 70.1	SMF70SLH		z/OS system	PCCASLIH in PCCA	# of entries to I/O SLIH
	SMF 70.1	SMF70TPI		z/OS system	PCCASTPI in PCCA	# of TPI with CC=1
1	SMF 70.1	SLH + TPI	394366	z/OS system		SMF70SLH + SMF70TPI
2	SMF 70.1	SMF70NIO	259662	z/OS system	Internal IOS control block	Number of I/Os for this CPU

The source for the SMF70NIO field is always updated for every Start Subchannel instruction (and *only* **Start** Subchannel instructions) the I/O Supervisor issues, without regard to what that device is. Therefore. it includes I/Os to channel-attached devices such as disk, tape, unit record equipment, and so on. Despite this, as you can see in Figure 12 on page 52, there is a significant difference between the value presented in the SMF70NIO field and the total interrupt rate as reported by the SMF70SLH + SMF70TPI fields. Some possible reasons for the difference are:

- The SMF70SLH+SMF70TPI fields include interrupts from PCIE devices (OSA adapters, for example), but the 'I/Os' to those devices do not use the Start Subchannel (SSCH) instruction, and therefore are *not* included in the SMF70NIO field.
- I/Os issued by programs that use Program Controlled Interrupts (PCI), such as the MVS Program Loader, paging I/Os, XCF CTC communication, and other essential system tasks[4] can generate multiple interrupts per Start Subchannel.
- Many types of activities performed by modern storage subsystems (e.g., flash copies, state changes) that generate interrupts that are not associated with I/Os from programs. Given the increasing amount of replicated data that exists in customer configurations (with Metro Mirror, Global Mirror, Metro Global Mirror, FlashCopy, SafeGuarded Copy, and so on), these types of interrupts are likely to grow as a percent of the overall number of interrupts as time goes on.
  - An example of a scenario where you might have much larger numbers of interrupts than there are started I/Os is a small Sandbox system. Such a system might not issue many I/Os, but it will still receive state change interrupts from volumes that are offline (but physically connected) to that system.

Note: Synchronous I/Os are not included in SMF70NIO, but because they do not generate interrupts, they are not included in SMF70SLH+SMF70TPI either. Accordingly, Synchronous I/O activity is unrelated to any potential difference between the number of I/Os and the number of interrupts.

Unfortunately, there are no detailed metrics about the types of interrupts that are being handled by the SLIH, so there is no easy way to definitively explain the difference between the number of interrupts and the number of started I/Os. The information above should be some help.

You also need to consider how important it is to know why the numbers are different. If you are looking for an indicator of the total volume of work being processed by a system, particularly for performing before-and-after comparisons, the number of interrupts (as used by IBM in their ITRR calculations) might be the better number to use. If you are looking for an

[4] Both JES2 and JES3 used to use PCI, however neither of them do anymore, except in the extremely rare case of a direct-attached (JES-mode) 3800 printer.

I/O metric from SMF 70 records that correlates to business workloads, SMF70NIO might be a better choice than SMF70SLH+SMF70TPI.

In any case, understanding that there *are* differences between the total number of 'I/O interrupts' and the number of started I/Os, and some of the reasons for those differences, might help you select the metric that is the most appropriate for the use you have in mind.

SMF Type 23 Fields

SMF record type 23 is titled “SMF Status” and as the name suggests it provides general information about SMF processing for both legacy and logstream approaches (e.g., records written, buffers used). But it also has some overall system information that some sites use to compare relative CPC capacities. Two of those fields are interesting in the context of this article - the SMF23NIO and SMF23NID fields[a]. *These fields get their data from the same source as SMF70NIO*.

If you compare the SMF23NID and SMF70NIO fields you might see differences in some intervals. In our experience, the differences we saw were due to the type 23 and type 70 records being produced at different times. By default, the type 23 records are produced every hour, based on the time of the IPL. Type 70 records are generally produced every INTVAL minutes, on the SYNCVAL time (where INTVAL and SYNCVAL are both defined in the SMFPRMxx member). However, you *can* tell SMF to create the type 23 records based on the INTVAL and SYNCVAL settings - you do this by specifying STATUS(SMF,SYNC) in your SMFPRMxx member. Setting up your SMFPRMxx in this way should result in nearly identical values in the corresponding fields in the type 23 and type 70 records.

a. Both of these fields are currently described in the SMF manual as “total number of I/Os for this interval” - in fact, the SMF23NI**O** field contains cumulative values, so the description of that field is a little inaccurate - hopefully that will be addressed in a future update to the SMF manual

SMF 78 Fields

SMF 78.3 records are described as reporting “I/O Queueing Activity”. The section with global I/O measurement counts is generated at the IOP (Input Output Processor) level.

Terminology

Depending on where you look in IBM documentation, you will see the term **IOP** (I/O Processor) or **SAP** (System Assist Processor). They are two names for the same thing. Because the term 'IOP' appears to be used more frequently in the RMF and SMF documentation, we will use that term in all cases in this article.

The IOP is a system Z Processor Unit (PU) mostly dedicated to handling the starting and completion of I/O requests[a]. In the 'old days', the functions carried out by IOPs today were collectively called the 'Channel Subsystem'. There is limited IBM documentation on exactly which 'I/O' requests are handled by the IOPs. From the perspective of this article, however, the important thing is that information about the total number of Start Subchannel requests and I/O interrupts handled by each IOP are reported in SMF type 78.3 records.

If you are really interested in the role of the IOP, refer to IBM Redbook *ABCs of z/OS System Programming, Volume 10*, SG24-6990, last updated in 2018, or an even older IBM paper from 2007, *Input / Output: A White Paper*, by **John Kettner**.

a. The IOP is also involved in Asynchronous zEDC requests, some Storage Class Memory (SCM) requests, and Server Time Protocol (STP) processing.

The R783IPII field (“Number of processed I/O interrupts”) is based on data retrieved directly from each IOP, and reflects the number of I/O interrupts handled by the IOP. As seen in Figure 13, the values it reports are very close to the sum of the SMF70SLH and SMF70TPI fields seen above. We've also seen a few examples where the R783IPII number was a little higher than the SMF70SLH + SMF70TPI values, but always very close.

Figure 13 - I/O Count Fields from SMF 70 and 78 Records

Row	Record	Field	Value	Level	Source	SMF manual text / Formula
1	SMF 70.1	SLH + TPI	394366	z/OS system		SMF70SLH + SMF70TPI
2	SMF 78.3	R783IPII	390971	IOP	Internally obtained from the firmware	Number of processed IOP interrupts
6	SMF 70.1	SMF70NIO	259662	z/OS system	Internal IOS control block	Number of I/Os for this CPU
7	SMF 78.3	R783IIFS	254958	IOP	Internally obtained from the firmware	Number of I/O functions initially started

Another interesting field in the 78.3 record, R783IIFS, is also based on data retrieved directly from each IOP. It is the number of start-subchannel (SSCH) requests that were started for the reporting system on that IOP. As you can see in Figure 13, the number of started I/Os as reported by the IOP(s) (field R783IIFS) is very close to the SMF70NIO field. The values in the SMF70NIO field come from an IOS control block while the R783* numbers come directly from the IOP, so some small difference is to be expected. They are also retrieved by the Data Gatherer at slightly different times - another possible explanation of small differences. It is also possible that there are some requests (a very small number, based on our data) that are included by either IOS in its SLH and TPI counters, or by the IOP, but not included by the other. We were unable to identify *exactly* which types of requests those might be.

Since synchronous I/Os are driven directly by the general purpose CP or the zIIP, and they do not result in I/O interrupts, synchronous I/Os are also not included in either of these IOP fields.

SMF 72 Field

The R723CIRC field in the SMF 72.3 record reports the total non-paging DASD I/O Start Subchannel count at the WLM service class level. The Data Gatherer uses a WLM service to gather this information, and WLM in turn gathers that information from the OUXB and ENCB control blocks, both of which are maintained by IOS.

As you can see in Figure 14, there is a significant difference between the sum of the R723CIRC fields for every service class and the value of the SMF70NIO field. We mentioned previously that the SMF70NIO field contains information for many I/O categories (for example, tape I/Os, channel-attached communications devices, paging I/Os, and so on), while the R723CIRC field does not include paging I/O or I/Os to non-DASD devices. While the SMF70NIO field is intended to provide a system-level view of activity, the R723CIRC field provides a view that is more application/workload centric. These are factors in explaining why the R723CIRC value is lower in Figure 14 than the SMF70NIO or R783IIFS fields.

Figure 14 - I/O Count Fields from SMF 70, 72, and 78 Records

Row	Record	Field	Value	Level	Source	SMF manual text / Formula
1	SMF 70.1	SLH + TPI	394366	z/OS system		SMF70SLH + SMF70TPI
2	SMF 78.3	R783IPII	390971	IOP	Internally obtained from the firmware	Number of processed IOP interrupts
6	SMF 70.1	SMF70NIO	259662	z/OS system	Internal IOS control block	Number of I/Os for this CPU
7	SMF 78.3	R783IIFS	254958	IOP	Internally obtained from the firmware	Number of I/O functions initially started
11	SMF 72.3	R723CIRC	192165	WLM service class	OUXBIOSC in OUXB	Total non-paging DASD I/O start subchannel count

SMF 74 Fields

SMF 74.1 records report a broad range of I/O metrics including counts at the device level. The SMF74SSC field reports the start subchannel count, including SSCH and RSCH (resume) instructions. It is populated from field ECMBSschRschCount in the device ECMB (Extended Channel Measurement Block). By default, ECMBs are only created for disk and tape devices. The CMB parameter in the IEASYSxx member of Parmlib can be used to extend this to include other (mainly older) types of devices such as Unit Record devices, Character Reader devices, and so on. SMF74SSC only captures I/Os for devices that have associated ECMBs, which may help explain why its value is less than SMF70NIO.

Traditional SMF 74 fields like SMF74SSC only include metrics for (legacy) asynchronous I/Os. To cater for the 'new' synchronous I/Os, an entire set of SMF74S* fields (including SMF74SQR and SMF74SQW) that report exclusively on synchronous I/Os has been added to the end of the Device data section of the SMF 74.1 records. SMF74SQR reports the

number of successfully completed synchronous I/O read requests. “Successfully” differentiates from reads initiated as synchronous I/Os but that had to be re-driven as standard zHPF asynchronous I/Os because the requested data was not found in the disk subsystem cache (as required for a successful synchronous I/O).

SMF74SQR (reads) and its companion SMF74SQW (for writes) are populated from fields ECMXSynReadCnt and ECMXSynWriteCnt in the ECMX (Extended Channel Measurement Block extension). The ECMX control block is a control block extension introduced for synchronous I/Os.

RMF only includes metrics in the SMF 74.1 records for the types of devices specified in the ERBRMF00 member (see the sample member in Figure 15).

Figure 15 - Excerpt from Sample ERBRMF00 Parmlib Member

```
/*****************************************************************/
/* PART 1: MEASUREMENTS                                          */
/*****************************************************************/
  CACHE                    /* CACHE STATISTICS                  */
  CHAN                     /* CHANNEL STATISTICS                */
  CPU                      /* CPU STATISTICS                    */
  CRYPTO                   /* CRYPTO HARDWARE                   */
  DEVICE(DASD)             /* DIRECT ACCESS DEVICES MEASURED    */
  DEVICE(NOTAPE)           /* NO TAPE DEVICES MEASURED          */
  DEVICE(NOCHRDR)          /* NO CHARACTER READER DEVICES MEASURED */
  DEVICE(NOUNITR)          /* NO UNIT RECORD DEVICES MEASURED   */
  DEVICE(NOCOMM)           /* NO COMMUNICATION DEVICES MEASURED */
  DEVICE(NOGRAPH)          /* NO GRAPHICS DEVICES MEASURED      */
  DEVICE(NONMBR)           /* NO SELECTION BY DEVICE NUMBERS    */
  DEVICE(NOSG)             /* NO SELECTION BY STORAGE GROUPS    */
```

In addition to adding these SMF 74.1 fields, Figure 16 now also includes a column indicating whether metrics include synchronous I/Os. (The site from which I captured this sample data has not yet implemented synchronous write I/Os.)

Figure 16 - I/O Count Fields from SMF 70, 72, 74, and 78 Records

Row	Record	Field	Value	Disk I/Os	Paging I/Os	Sync I/Os	Tape I/Os	'Other' I/Os	Level	Source	SMF manual text / Formula
1	SMF 70.1	SLH + TPI	394366	Yes	Yes	No	Yes	Yes	z/OS system		SMF70SLH + SMF70TPI
2	SMF 78.3	R783IPII	390971	Yes	Yes	No	Yes	Yes	IOP	Internally obtained from the firmware	Number of processed IOP interrupts
3	SMF 74.1	SMF74SSC+SQR	338183	Yes	Yes	Yes	Yes	Yes	Device	See component fields below	SSCH & RSCH + successful sync I/Os
6	SMF 70.1	SMF70NIO	259662	Yes	Yes	No	Yes	Yes	z/OS system	Internal IOS control block	Number of I/Os for this CPU
7	SMF 78.3	R783IIFS	254958	Yes	Yes	No	Yes	Yes	IOP	Internally obtained from the firmware	Number of I/O functions initially started
8	SMF 74.1	SMF74SSC	235418	Yes	Yes	No	Yes	Parm (Note 1)	Device	Device ECMB (Extended Channel Measurement Block), based on field ECMBSschRschCount	Start subchannel count (SSCH & RSCH)
11	SMF 72.3	R723CIRC	192165	Yes	No	No	No	No	WLM service class	OUXBIOSC in OUXB	Total non-paging DASD I/O start subchannel count
12	SMF 74.1	SMF74SQR	102765	No	No	Reads	No	No	Device	ECMX (Extended Channel Measurement Block extension) field ECMXSynReadCnt	# of successfully completed sync I/O read requests
13	SMF 74.1	SMF74SQW	0	No	No	Writes (Note 2)	No	No	Device	ECMX (Extended Channel Measurement Block extension) field ECMXSynWriteCnt	# of successfully completed sync I/O write requests

Note 1: I/O reporting of device types specified in RMF ERBRMF00 Parmlib member
Note 2: Selected site is only doing sync reads, no sync writes

SMF 42 Fields

As you might expect, the DFSMS-owned type 42 records are more focused on data set, storage class, and volume-level activity than the Data Gatherer-owned type 7x records which are generally more focused on system-level activity. Also, while most of the type 7x fields described above include I/Os to *multiple* types of devices, the type 42.5 and 42.6 records contain information about *only disk* activity. This helps explain the reason SMF74SSC (which includes at least tape and potentially other non-disk devices) reports higher values than the SMF 42 metrics, as you can see in Figure 17.

Figure 17 - I/O Count Fields from SMF 42, 70, 72, 74, and 78 Records

Row	Record	Field	Value	Disk I/Os	Paging I/Os	Sync I/Os	Tape I/Os	'Other' I/Os	Level	Source	SMF manual text / Formula
1	SMF 70.1	SLH + TPI	394366	Yes	Yes	No	Yes	Yes	z/OS system		SMF70SLH + SMF70TPI
2	SMF 78.3	R783IPII	390971	Yes	Yes	No	Yes	Yes	IOP	Internally obtained from the firmware	Number of processed IOP interrupts
3	SMF 74.1	SMF74SSC+SQR	338183	Yes	Yes	Yes	Yes	Yes	Device	See component fields below	SSCH & RSCH + successful sync I/Os
4	SMF 42.n	All I/Os w/sync	318113	Yes	Yes	Yes	No	No	Data set		S42DSION + S42VDION + S42VXION + S42VVION + S42VSION
5	SMF 42.n	Captured I/Os w/sync	318086	Yes	Yes	Yes	No	No	Data set		S42DSION + S42VDION + S42VXION + S42VVION
6	SMF 70.1	SMF70NIO	259662	Yes	Yes	No	Yes	Yes	z/OS system	Internal IOS control block	Number of I/Os for this CPU
7	SMF 78.3	R783IIFS	254958	Yes	Yes	No	Yes	Yes	IOP	Internally obtained from the firmware	Number of I/O functions initially started
8	SMF 74.1	SMF74SSC	235418	Yes	Yes	No	Yes	Parm (Note 1)	Device	Device ECMB (Extended Channel Measurement Block), based on field ECMBSschRschCount	Start subchannel count (SSCH & RSCH)
9	SMF 42.n	All I/Os less sync	215346	Yes	Yes	No	No	No	Data set		DSION + VDION + VXION + VVION + VSION - SNROS - SNWOS
10	SMF 42.n	Capt I/Os less sync	215319	Yes	Yes	No	No	No	Data set		DSION + VDION + VXION + VVION - SNROS - SNWOS
11	SMF 72.3	R723CIRC	192165	Yes	No	No	No	No	WLM service class	OUXBIOSC in OUXB	Total non-paging DASD I/O start subchannel count
12	SMF 74.1	SMF74SQR	102765	Yes	No	Reads	No	No	Device	ECMX (Extended Channel Measurement Block extension) field ECMXSynReadCnt	# of successfully completed sync I/O read requests
13	SMF 74.1	SMF74SQW	0	Yes	No	Writes (Note 2)	No	No	Device	ECMX (Extended Channel Measurement Block extension) field ECMXSynWriteCnt	# of successfully completed sync I/O write requests
14	SMF 42.n	S42SNROS	102767	Yes	No	Reads	No	No	Storage class		Number of synchronous I/O read successes
15	SMF 42.n	S42VSION	27	Note 3	No	No	No	No	Storage class		SSCH count from "System I/O statistics" section

Note 1: I/O reporting of device types specified in RMF ERBRMF00 Parmlib member
Note 2: Selected site is only doing sync reads, no sync writes
Note 3: Only I/Os that DFSMS cannot properly categorize are reported in S42VSION as "System I/Os"

The SMF 42.5 records contain sections with I/O metrics at the storage class and volume levels and are created at the end of each SMF interval. The SMF 42.6 records contain similar information at the data set level. For purposes of this discussion, we are combining data from both 42.5 and 42.6 sources to present a more complete picture of total disk I/Os. These metrics are derived when IOS (I/O Supervisor) in combination with the device driver (which, for synchronous I/Os, is Media Manager) requests an SSCH I/O operation to a device.

These SMF 42 metrics differ from the “legacy” SMF 74 metrics described above in that synchronous I/Os are included in many of the metrics. Additionally, the SMF 42.6 records include zHyperWrite I/Os to, and Consistent Reads from, Metro Mirror secondary volumes.

The primary field of interest to us is S42DSION (“total number of I/Os”) for each data set. To get a complete count of I/Os, this value is added to S42VDION (VTOC), S42VXION (VTOC Index), and S42VVION (VVDS).

S42VSION - System (or, 'Uncaptured') I/Os

One other metric that can be included to derive a complete total is S42VSION, "system I/Os". This field and a related section in the 42.5 record were created by APARs OA55709 and OA55710 (in 2019). These can also be thought of as "uncaptured" I/Os, as they reflect atypical situations where the control blocks DFSMS uses to categorize the I/Os do not exist. Examples include RMF requests to retrieve DSS or FICON director metrics, the standalone "device release" to end the hardware reserve for a device, and copy services commands to control advanced function operations that are unrelated to data stored on the device. z/OS systems staff will recognize the parallels with CPU capture ratios, which quantifies CPU that cannot be assigned to a specific service class.

Synchronous I/Os are included the S42__ION metrics described above and are also tracked separately in the Synchronous I/O section of the SMF 42.5 and 42.6 records. S42SNRDT and S42SNWTT capture the Synchronous I/O Read and Write Attempts. As explained above, for a synchronous I/O to be "successful" the data must be present in the cache of the storage controller ("cache hit"). S42SNROS and S42SNWOS capture the numbers of successful synchronous I/O reads and writes, respectively. The differences between those values and S42SNRDT and S42SNWTT ("attempts") quantify the I/Os that had to be re-driven as standard zHPF asynchronous I/Os. In my sample data, 86% of synchronous I/O attempts were successful.

In the midst of all the challenges with accounting for differences in the various I/O counts, it was very encouraging to observe that the counts of synchronous reads reported by SMF74SQR and S42SNROS were effectively identical!

References

As you have probably concluded by now, we wish there was more documentation to help you understand exactly what is or is not included in the SMF fields discussed in this article. We found the following documents to be helpful:

- IBM Manual, *z/OS MVS System Management Facilities*, SA38-0667.
- IBM Redbook, *ABCs of z/OS System Programming Volume 10*, SG24-6990.
- IBM Paper, *Input Output: A White Paper*, by **John Kettner**.

There is also helpful information in the text of some related APARs. And the SMF mapping macros in the SYS1.MACLIB data set sometimes contain a little more information than is provided in the SMF manual.

Editor's Note: We are all used to seeing fields in SMF records that are specifically for IBM's own use - in fact, there are even entire subtypes that are described as “This information is for IBM internal use only” in the SMF manual. However, for SMF fields that are labeled something other than Reserved or IBM internal use, I would encourage IBM and *all* vendors to provide the information customers need to be able to use that data correctly. Information that is obvious to the developer that is responsible for a particular SMF record type isn't always as clear for those of us that are limited to the description in the SMF manual or the mapping macro.

Summary

I (Frank) want to start by thanking Todd for his really outstanding work on this article. Every time I sent Todd a rhetorical question, he immediately came back with an answer or a suggestion based on his analysis of real customer data. I have encountered situations in the past where the values of two fields didn't make sense, but I've never seen so many I/O-related fields laid out in such a structured manner, and based on real world SMF data.

Both Todd and I would have loved to be able to explain, down to the single I/O, every difference between every pair of fields. We didn't achieve that, but we did learn a great deal from the data and the insights provided by IBM's IOS and DFSMS experts. Even if we can't definitively explain every last difference, just knowing that there are differences is very valuable. As someone who loves playing with SMF data, and helping readers understand the values they see, I believe this article will be invaluable.

Apart from the intended purpose of helping readers better understand these related fields, I believe this article holds another very important lesson for us. Regardless of whether you have 4 years or 40 years of experience, actually looking into your SMF data can deliver invaluable insights and lessons. It is all well and fine to know the description in the SMF manual, but to really understand what is happening under the covers, you need to look at the data and consider what it is telling you.

At the start of this exercise, I just assumed that the number of interrupts seen by a system would be roughly the same as the number of I/Os that were started by the system. But when Todd showed me data from different systems, where the number of interrupts ranged from being 7% more than the number of I/Os to being nearly 50 times more, I was forced to realize that my 'understanding' of interrupt processing was far too simplistic. I am still far from being any sort of expert on this aspect of z/OS performance, but at least now I know what I don't know, and I'm on a mission to address that shortcoming. And that wouldn't have happened without Todd's enthusiasm and hard work (and patience!) and the great support we received from the Development folks in IBM.

If you want to become the next Todd or the next Cheryl, this article holds a very valuable lesson for you - never stop digging and asking 'why' until you have an answer that makes sense. I hope you enjoyed it and found it valuable.

News

Are you fed up with the 24x7 news cycle, the feeling of information overload, the constant jockeying for attention, and the lingering doubt about the veracity of the information you are seeing? Well, you've come to the right place. This section of the Tuning Letter provides a *quarterly* summary of hardware-related flashes, important APARs, useful IBM Techdocs, and information about upcoming conferences - in short, a distillation of information we encountered over the last quarter that we feel you should know about, all nicely sorted, filtered, and packaged in one place.

Target Audience

We believe that *every* reader should find something of interest in this section - if you don't, then please let us know. Our objective is to make you aware of information that you might otherwise miss *and* give you back some time in your busy days. If there is information that we are missing that would be valuable to you or your colleagues, let us know and we will add that if possible.

IBM Z Hardware Alerts

As you know, IBM delivers fixes and enhancements to CPC microcode from time to time. Given the rate at which IBM has been delivering new service bundles for the z16 CPCs, it would be prudent to speak to your service representative about finding the 'Goldilocks' service level for your z16 - you don't want to apply new Bundles *too* quickly, but you also don't want to get too far behind current service levels.

Anyone that has a Resource Link ID can use the Machine Alerts section to check the various Driver Bundles and the problems they addressed, or new functions they delivered, as shown in Figure 18 on page 63[5].

[5] You can check the latest hardware alerts by logging on to Resource Link, then clicking on 'Problem Solving' in the left panel, then 'Machine Alerts' in the Alerts section.

Figure 18 - IBM Resource Link Machine Alerts Screen

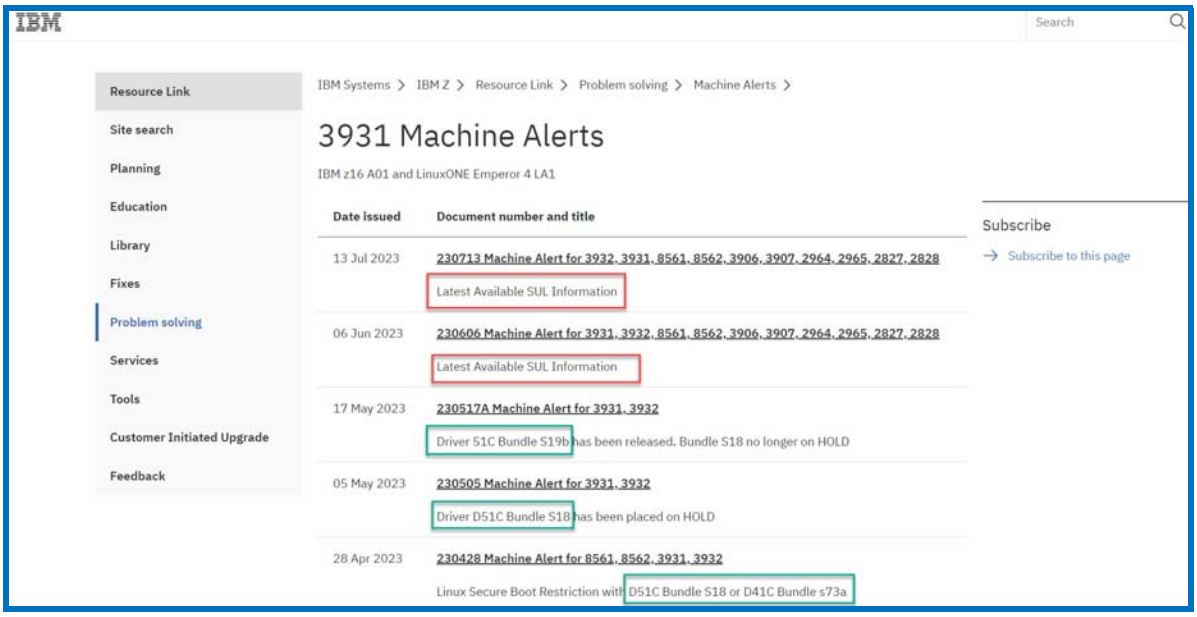

The most interesting alert over the last quarter is one about an issue that may result in a complete system outage for IBM z16 servers after error recovery occurs on a Coupling Express2 Long Reach adapter. That Machine Alert was issued on March 19, 2024. The fix for the problem was announced in a different part of the site (Hiper Alerts - see below) on April 17, 2024.

Note: There have been *seven* Alerts for z15 and z16 (and a subset of those also apply to z14, z13, and even zEC12) since December 31, 2023. We recommend reviewing the purpose for those alerts with your service representative to ensure you are not missing out on performance enhancements, or exposing yourself to availability issues.

z CPC customers can get MCL information for *their* CPCs by logging on to Resource Link, going to the Tools section and then selecting > Machine Information > Machine List, then select your CPC, then click on the EC/MCL option, and finally see the Missing MCL Details link over on the right in the Details section.

There is also a “Hiper alerts” link in the Problem Solving panel - this refers specifically to *Hardware* HIPER alerts, not the normal software ones. There were two new Hiper Alerts during the last quarter, one for the fix for the CE LR problem mentioned above, and one for a communication problem on the internal service network on z16 systems. While there generally is not a high level of activity on that page (the last one was nearly a year ago, in May, 2023), it is wise to monitor this page to be on the safe side.

ResourceLink Apps moving to new sites

Note that IBM is in the process of moving much of the material from the ResourceLink website to other sites. You can find more information here.

Hopefully you have already seen this, but there was also a Machine Alert titled "*Preparing customer firewalls and proxies for the upcoming infrastructure changes – Call Home, Electronic Fix Distribution*". This is an update to an announcement made in October 2023, and includes additional detail.

Red Alerts

There was one 'Red Alert' issued during the last quarter.

Red alert '*SMP/E SOURCEID RSU2312 is incomplete*' was released by IBM on January 29, 2024. It applies to anyone that obtained PTFs from IBM between January 4, 2024 16:00 GMT and January 26, 2024 01:00 GMT. During that time the list of PTFs assigned to SMP/E SOURCEID RSU2312 was incomplete. The list was missing a number of PTFs closed COR in November 2023. As a result, some recommended PTFs may not have been delivered, and the SMP/E ++ASSIGN statements for SOURCEID RSU2312 delivered with PTF orders were incomplete. The recommended action is to submit a new order for recommended PTFs - that should result in your getting the missing PTFs as well as the corrected ++ASSIGN statements.

To be sure that you are aware of new Red Alerts as soon as they are issued, use the IBM subscription service to subscribe to Red Alerts. Ensure that at least *two* people in your site are getting automatically alerted by IBM any time a new Red Alert is issued. Hopefully, Red Alerts will be few and far between, but when a Red Alert *is* issued, you want to be sure that you don't miss it.

For an up-to-date list of Red Alerts, see: https://www.ibm.com/support/pages/node/959863.

Preparing for z16

Even if you are not on the verge of installing a shiny new z16, we all know how long it can take to roll service out to all your systems. Therefore, it would be prudent to check the IBM FIXCATs for the z16 if you are planning on installing a round of service on your systems. There are separate sets of FIXCATs for the z16 A01 (**`IBM.Device.Server.z16-3931.*)`** and A02 (**`IBM.Device.Server.z16A02-3932.*`**) models, and separate FIXCATs for z16 Exploitation, Recommended Service Updates from the PSP bucket, and fixes that are required for z/OS to run on a z16.

We expect that all our readers are familiar with using FIXCATs. If you are not, download the latest HOLDDATA (make sure to retrieve the Full 2 Year file) from the HOLDDATA website and check the IBM FIXCAT website (https://www.ibm.com/support/pages/ibm-fix-category-values-and-descriptions) for a brief

description and some sample SMP/E statements that will use the FIXCAT information to generate a report showing the fixes that are missing from your systems. If you're new to FIXCATs, start with our introduction in our 'Fixing Your Cats the Easy Way' article in *Tuning Letter 2015 No. 1*.

Upcoming Withdrawal from Marketing Dates

We wanted to remind readers that own z15s of the upcoming deadlines for ordering upgrades:

June 30, 2024	Withdrawal of new builds and *hardware* MES changes for z15 T02. This is also the last date to order upgrades to a z15 T02 from a z13 or z14.
December 31, 2024	Withdrawal of LIC upgrades (that is, upgrades that do *not* require the installation of new hardware - an example of a LIC upgrade might be purchasing memory that is already installed but not previously available to you) for z15 T01 CPCs.
June 30, 2025	Withdrawal of LIC upgrades for z15 T02s.

Just to draw this to your attention - IBM previously announced that IMS 15.4 is the last IMS release to support zEC12/zBC12 CPCs. In support of this, IMS APAR PH58758 adds preconditioning code to enforce the requirement that IMS releases after 15.4 will require a minimum CPC level of *z13*. Hopefully all our readers are on newer technology than z13, but just in case you are not, we wanted to bring this to your attention.

We hope that all our readers that have any responsibility for hardware planning and management are familiar with **Don Clarke's** outstanding *Mainframe Lifecycle History* paper. I really can't say enough good things about this document. It presents all the information you need about IBM CPC key dates in an easy-to-consume format. It shows the possible upgrade paths for both business class and enterprise class models.

Unfortunately it is not possible to subscribe specifically to Don's document, but if you go to https://www.ibm.com/mysupport/s/article/Subscribing-to-Products?language=en_US you can find instructions for subscribing to topics via IBM's My Notifications service. If you set up a subscription for a CPC ('z14', or 'z15', or 'z16' for example), you will be notified whenever he updates it.

Preparing for z/OS 3.1

z/OS 3.1 went Generally Available on September 29, 2023. You can find more information about the enhancements in this new version of z/OS in the following Tuning Letter articles:

- 'What's New in z/OS – 3.1 Preview', *Tuning Letter 2023 No. 1*. This is based on the Preview announcement and covers supported hardware levels plus the enhancements to the z/OS Base, JES2, and SDSF.
- 'What's New in z/OS – 3.1 Preview Part 2', *Tuning Letter 2023 No. 2*. This is also based on the Preview announcement and covers DFSMS and its components, Communication Server, z/OSMF, RMF, zACS, zWIC, USS, RACF, Change Tracker, and others.
- 'What's New in z/OS - 3.1 Enhancements Part 3' in Tuning Letter 2023 No. 3. It is based on the August 8, 2023 formal announcement of z/OS 3.1 (announcement letter number AD23-0308) and covers all the enhancements to z/OS 3.1 that we could find *that were not already mentioned in the Preview announcement*.

Now that z/OS 3.1 is available, it is time to start preparing your current systems for the migration to 3.1. The best single source of information to help you plan for that project is **Marna Walle's** two *Upgrade to z/OS 3.1* sessions from SHARE in Orlando:

- Session 10510: '*Upgrade to z/OS 3.1, Part 1: Planning*'.
- Session 10511: '*Upgrade to z/OS 3.1, Part 2: Technical Actions'.*

Reminder: Because SHARE material is only available to SHARE members now, you must be logged on to the SHARE website before you can access these presentations.

There are two IBM FIXCATs related to z/OS 3.1:

IBM.Coexistence.z/OS.3.1	Fixes that allow z/OS V2.4 and z/OS V2.5 to coexist with, and fallback from, z/OS 3.1.
IBM.TargetSystem-RequiredService.z/OS.3.1	Fixes required on other IBM products to allow them to run on z/OS 3.1.

Important: Java 17 Requirement: Refer to the *Software requirements for running z/OS 3.1* document for information about z/OS products that use Java as well as other product requirements for z/OS 3.1.

Java 17 on z/OS

It really doesn't seem that long since IBM announced Java 11 support on z/OS, packaged in a product called *IBM Semeru Runtime Certified Edition for z/OS Version 11* (Product ID 5655-DGJ). So I was quite surprised to find on the Java 11 Lifecycle web page that Java 11 reached End of Marketing on April 1, 2024 and the announced End of Support date is just over six months from now - November 30, 2024.

The follow-on version of Java, Java 17, is delivered by *IBM Semeru Runtime Certified Edition for z/OS Version 17* (Product ID 5655-UA1) which became generally available on

August 25, 2023. You can find more information in the IBM Community post titled '*Java IBM Semeru Runtime Certified Edition for z/OS, Version 17 is now available!*' by **James Tang**.

Recommendation: Note that the End of Support date for Java 8 on z/OS is unchanged at September 30, 2026. However, as per IBM Withdrawal notice AD22-0559, Java 8 on z/OS was withdrawn from *marketing* on January 31, 2024. Hopefully you *already* have Java 8 and IBM Semeru V11 (*and* IBM Semeru V17) installed on your system.

Recent APARs

This section contains information about APARs that were closed over the previous quarter that we feel are important or will be of interest to our readers.

z/OS Quarterly Continuous Delivery Announcements

IBM normally releases a z/OS Enhancements announcement once every quarter, listing recent enhancements that were shipped in the Continuous Delivery service stream. The announcements are not necessarily 90 days apart, but there *will* be one in each quarter.

The 1Q 2024 Continuous Delivery announcement for z/OS 3.1 is available at https://community.ibm.com/community/user/ibmz-and-linuxone/blogs/fiona-king1/2024/02/27/zos-31-1q-2024-enhancements?CommunityKey=2eba32a4-9426-4331-81e7-8ea120e4fd99

You can find a list of all the z/OS 3.1 announcements at https://community.ibm.com/community/user/ibmz-and-linuxone/blogs/fiona-king1/2024/02/12/zos-announcements

New Function APARs Website

We have mentioned IBM's New Function APARs website in the past. The intent of that site is to publicize APARs that deliver new functions for any IBM product that runs under z/OS - so it is not only enhancements to z/OS itself, but also CICS, Db2, and the other major subsystems, as well as tools like the Db2 Admin Tool and CICS Performance Analyzer - basically any IBM product that supports SMP/E. We have found that it isn't always 100% complete, but it is certainly a valuable tool for any system programmer. The biggest problem we have is closed APARs whose titles were never updated from just "New Function" - however, that is not the fault of the people that maintain the New Function APARs site.

If you want to access the site using HTTP, you must have an "IBM Support Secure File Transfer ID". If you don't already have such an ID, you can easily get one by going to

https://www.secure.ecurep.ibm.com/transferids/, and entering your IBM ID and password. Make sure to keep a note of the transfer ID and password that is generated. When you have a transfer ID, go to https://testcase.boulder.ibm.com/ to logon and navigate to /fromibm, then /mvs, and look for the mvsstore* files.

Using the Information in This Section

IBM creates LOTS of APARs, many of which will not apply to you. Of the ones that do, many of those are for trivial problems that will not have any significant impact on your business. However, there are a small number that *are* important to your business, possibly because they deliver valuable new functions, or maybe they impact performance or CPU consumption. In this section we do our best to highlight APARs that we believe will be of interest to large numbers of our readers.

New Function APARs

The major new z/OS functions delivered in the last quarter were enhancements to ABO, Capacity Provisioning Manager, Comm Server, CICS, Db2, DFSMS, IOS, MQ, Omegamon, RMF, SCRT, USS, WLM, zFS, and z/OSMF. Some of them roll z/OS 3.1 enhancements back to previous releases.

z/OS-Related New Function APARs

Note: The component ID for z/OS 3.1 is R7E0.

- OA61972 (OMVS 2.5-3.1) - **New Function** - *z/OS UNIX Support for Namespaces*. These system calls will allows z/OS UNIX-based applications to run in containers natively on z/OS and follows the Open Container Initiative (OCI) specifications. This is in support of the new z/OS Container Platform.

z/OS Container Platform (z/OS CP): You can find much more information about this new capability on the z/OS Container Platform website.

- OA62281 (OMVS 2.5-3.1) - **New Function** - *z/OS UNIX Support for Containers*. This PTF adds z/OS UNIX functionality to enable WLM support for IBM z/OS Container Platform (zOSCP) technology. The z/OS UNIX exec callable service will now support setting the _BPX_CONTAINER_POD_ID, _BPX_CONTAINER_ID, and _BPX_CONTAINER_QUAL environment variables. These variables will allow z/OS Container Platform workloads to be properly classified for WLM.

- OA62734 (OMVS 2.5-3.1) - **New Function** - *z/OS UNIX Support for Mount Namespaces*. More support for z/OS CP. This PTF adds z/OS UNIX functionality to enable support for mount namespaces, the pivot_root callable service, and additional flags on the mount and unmount shell commands.

- OA63911 (zFS 2.5-3.1) - **New Function** - *Providing Enhancements to z/OS File System Checks*. This APAR provides two new health checkers that can be used to assist in monitoring and improving cache performance (ZFS_CACHE_PERFORMANCE) and also with identifying zFS file systems with exception states (ZFS_EXCEPTIONS).

  New performance related health check for zFS.

  Note: If you are looking for documentation on these new health checks, use the z/OS 3.1 Health Checker User's Guide because the z/OS 2.5 version of that manual is no longer being updated.

- OA64530 / OA64552 (DFSMS 2.4-3.1) - **New Function** - *Cascaded Incremental Resynchronization*. Refer to this section in the z/OS DFSMS Copy Services manual for information about this new capability. Note that this support appears to largely be replacing existing functions, but more efficiently or in a simpler way.

- OA64559 / OA64569 / OA64570 / OA64571 / OA64907 / OA65219 / OA65488 (DFSMS 2.5-3.1) - **New Function**. All but one of these APARs contain the same description - "DFSMS Infrastructure for future new function." Whatever the new function is, APAR OA64559 appears to be the base APAR (all the other APARs are 'sysrouted to' from that APAR). Interestingly, the one APAR in that list that differs from this pattern is one for the z Workload Interaction Correlator. It is interesting the insights that you can glean from browsing through APARs, even ones that appear to say very little.

- OA65253 (XMLSS 3.1) - **New Function** - *XML System Services Support for Java 11 and 17 Semeru Migration*. The Apache Xerces XML Parser is no longer shipped with Java 11 or Java 17 (isn't Open Source wonderful?). Therefore, the jars associated with the parser are now packaged with XML System Services.

- OA65319 (DFSMShsm 2.5-3.1) - **New Function** - *Improve Scalability of the DFSMShsm Java hbackup Utility*. This APAR enhances the Java hbackup command to support a new -i option. The new option allows users to specify the host id of the DFSMShsm host to which the backup request should be directed. This allows users to distribute large UNIX file backup requests across multiple DFSMShsm hosts in a single LPAR. It would be *even nicer* if hbackup supported a common queue (like the HSM Common Recall Queue) to automatically distribute the tasks across the available hosts, but we'll take what we can get, and hope for better in the future.

- OA65675 (CPM 3.1) - **New Function** - *Capacity Provisioning Manager Support for 64-bit*. Prior to this APAR, CPM ran in 31-bit mode with Java 8. This APAR delivers support for Java 11 and 64-bit. APAR OA65482 delivers the BCPii support for this change to CPM.

- OA65797 (RMF 2.4-3.1) - *Increased Processor Usage in DDS(GPMSERVE)*. This APAR addresses a problem with high CPU consumption in the GPMSERVE address space by reverting to using the XL C/C++ Runtime Library _LARGE_TIME_API. Unfortunately, the APAR doesn't say what API it was using previously that was causing the increased CPU consumption.

- OA65801 (OM 550) - **New Function** - *Enhancements to CPU Loop Index That Extend Its Functionality*. This APAR adds the ability for a site to assign relative impacts of an address space's looping on other system address spaces, based on counts in the loop calculation.

Recommendation: Given the impact a looping address space can have on CPU consumption, any site using Tailored Fit Pricing and Omegamon for z/OS should investigate this new function.

- OA66012 (ADG 3.1) - **New Function** - *z/OS Data Gatherer REST Services and SMF Explorer Updates and Corrections*. In addition to some minor corrections, this APAR adds support of client certificates as an authentication method for the z/OS Data Gatherer REST services.

- OA66013 (ADG 3.1) - **New Function** - *z/OS Data Gatherer REST Services Java 17 Support*. This APAR delivers support for using Data Gatherer REST services together with Java 17.

- OA66064 (WLM 2.5-3.1) - *Updates to WLM Execution Delay Monitoring and Buffer Pool Mgmt*. This APAR addresses three problems related to WLM delay monitoring. The most perplexing one might be that the delay states for SESS_NETWORK and SESS_SYSPLEX are reversed, resulting in network delays being reported as sysplex delays and vice versa.

- OA66079 (DFSMSrmm 3.1) - **New Function** - *Semeru 17 Migration*. This APAR is required for RMM users wishing to use the HLL API on z/OS 3.1.

- OA66145 (RMF 3.1) - **New Function** - *Deliver Support for z/OS 3.1 WLM Enhancements and Grafana Plugin*. This APAR delivers RMF support for two WLM enhancements in z/OS 3.1: WLM AI Initiator and Implicit CPU protection. It also delivers RMF support for the Grafana plugin, a new function we hope to cover in more detail in a future Tuning Letter issue.

- PH39613 (CS 2.5-3.1) - **New Function** - *Communications Server Support for z/OS Container Platform*. This APAR delivers Comm Server support for zOSCP workloads by introducing a new type of VIPARANGE dynamic VIPA (DVIPA) called ZCONTainer in addition to other enhancements.

- PH53743 (IZBR 1.2) - **New Function** - *The Logical Corruption Protection (LCP) Function Within IZBR Does Not Correctly Support CyberVault Topology*. Prior to this APAR, the IZBR LCP User Interface used the Primary system to connect to a CyberVault. However, a CyberVault solution only allows for outgoing connections from the CyberVault. This APAR alters IZBR so the IZBR LCP Interface on CyberVault will now initiate all communications to the Primary site using a secure TCP/IP connection.

- PH57444 (Chg Tracker 1.0) - *Collection of Fixes*. This APAR is not flagged as a New Function one, however it addresses 21 issues in Change Tracker so we wanted to bring it to your attention.

- PH59284 / PH59285 / PH59286 (CSM 6.3.10) - **New Function** - *This PTF Will Upgrade Copy Services Manager to the 6.3.10 Dependency Level of Maintenance*. You can find information about the enhancements in CSM 6.3.10 in this document: https://www.ibm.com/support/pages/system/files/inline-files/csm_release_notes_6_3_10.pdf

- PH60156 (ABO 2.2) - **New Function** - *Improved Scalability at ABO Optimization Time*. This APAR adds improvements to reduce CPU time consumption and memory use by ABO while optimizing large input modules.

CICS-Related New Function or Performance APARs

- PH59256 (CICS 6.1) - **New Function** - *Provide a Script to Report PTF Level of CICS Jars*. This APAR delivers a new script 'getserviceinfo', that can be used under direction from IBM Service personnel. APAR PH58123 provides this new script for CICS 5.5 and 5.6.

- PH59350 (CICS 6.1) - **OPEN** - *CICS is not Processing Inbound Chunked Data Efficiently*. This APAR addresses a problem where CICS tasks show very high CPU consumption when using a chunked transfer-coding in order to send data from a client to CICS. The APAR contains a workaround until the PTF is available.

- PH59376 (CICS 6.1) - **OPEN** - *DFHDS0102 Shows QR TCB Dispatch / Interval Ratio Above 100%*. The DFHDS0102 message reports the ratio (expressed as a percentage) of the accumulated CPU time to the accumulated dispatch time for the QR TCB and the dispatch / interval ratio of the QR TCB. The frequency of this message is controlled by the INITPARM=(DFHQRCPU='nn') SIT parm. This APAR addresses an error that results in an incorrect percentage being reported.

The following web pages provide information about all the fixes for the following CICS releases:

- Fix list for CICS Transaction Server for z/OS 5.5.

- Fix list for CICS Transaction Server for z/OS 5.6.

- Fix list for CICS Transaction Server for z/OS 6.1.

And this web page provides information about the End of Service dates for CICS products:

- End of Service dates for CICS products.

In particular, please note that the End of Service for CICS TS 5.5 has been announced - September 30, 2025. If the z/OS folks stick with their traditional cadence, z/OS 3.2 will become generally available on that same date. I'm sure you will all be rushing to install z/OS 3.2 as soon as you can get your hands on it, so take the opportunity to move your CICS 5.5 regions to a new release in the meantime ;).

And speaking of new CICS releases - IBM announced the Open Beta for CICS TS 6.2 - see the announcement for more information. For a list of *all* CICS product announcements, and links to those announcements, refer to: https://www.ibm.com/support/pages/announcements-cics-products.

Db2-Related New Function/Performance APARs

- PH59381 (Db2 12-13) - **HIPER** - *QUIESCE Utility With WRITE(YES) Does Not Write All Updated Pages From GBP.* This HIPER APAR addresses a problem whereby a QUIESCE with WRITE(YES) does not necessarily result in all updated GBP pages being castout to disk. The updated pages are still in the GBP, so there is no data integrity issue. But if you have a process that expects the disk copy to reflect the latest position, that process might fail or it might copy out-of-date pages. The problem is caused by PTFs UI92363 or UI92364, so either apply the PTFs for this APAR, or SMP/E RESTORE those PTFs.

- PH59534 (Db2 13) - **New Function** - *Db2 New Function Level*. For information about Function Level 505 for Db2 13, along with activation instructions and application considerations, refer to: https://www.ibm.com/docs/en/db2-for-zos/13?topic=levels-function-level-505.

- PH59631 (Db2 13) - *Db2 V13 Excessive Notifies After a Write Gets a GBP Structure Full Error*. This APAR addresses delays in data sharing and a large amount of ECSA storage allocated by IRLM due to excessive notifies sent for castouts after group buffer pool structure full errors. We have seen customer service impacts as a result of floods of IRLM notifies, so, even though this APAR is *not* marked HIPER, if it was *my* system, I would apply the PTF sooner rather than later.

MQ-Related New Function APARs

- PH57431 / PH59260 / PH59270 / PH59271 / PH59272 / PH59273 / PH59274 / PH59275 / PH59276 (MQ 9.3.5) - **New Function** - *IBM MQ V9.3.5 Continuous Delivery (CD)*

Release. For a full description of the enhancements in MQ 9.3.5, refer to https://www.ibm.com/docs/en/ibm-mq/9.3?topic=delivery-whats-new-changed-in-mq-935

z/OSMF New Functions

z/OSMF is now unquestionably a core component of z/OS, delivering multiple functions that are not available elsewhere in z/OS. We truly hope that you have it started on at least one of your systems and that you have started to get familiar with it. As time goes by, more and more products and functions will require that z/OSMF is up and running.

A valuable source of information is https://ibm.github.io/z/ for LOTS of useful z/OSMF info AND an attractive user interface - in Git! Who knew...

If you have not already joined the z/OSMF Guild (see the Guild website) we *highly* recommend that you do so - it is an invaluable source of useful information about z/OSMF.

- PH56207 (z/OSMF 3.1) - **New Function** - *z/OSMF Parmlib Syntax Validation REST API*. This APAR delivers a new z/OSMF component called Parmlib Management. The first iteration of this new function provides a Parmlib Syntax Validation REST API which can be used to drive automatic syntax validation of 38 types of z/OS Parmlib members. Note that there is NO IZUPRMxx parm for the new plugin - if you want to use it, the plugin must be enabled using the z/OSMF desktop General Settings task.

  This first delivery does not provide a GUI - hopefully that will come in future releases/APARs. However, a very interesting aspect of this new capability is that the z/OSMF team developed a parser that can read the available documentation to determine the valid syntax and keywords for a given Parmlib member. This will hopefully make is easier for the z/OSMF team to roll out support for additional members. There could also be a knock-on benefit for all z/OS customers - if the z/OSMF team discover errors in the documentation, IBM will hopefully correct the documentation, thereby benefiting everyone.

Recommendation: You can find much more information about this enhancement, as well as a really excellent Demo, in the *New REST API for Parmlib Management and Desktop Enhancements* z/OSMF Guild session (you can download the slides, but you must watch the recording to see the Demo).

Even if you are not particularly interested in the new Parmlib plugin, **Joey Zhu's** demo is *still* worth watching because of his demonstration of how you would test and play with the z/OSMF API functionality.

- PH56233 (z/OSMF 3.1) - **New Function** - *z/OSMF Configuration Support for PARMLIB Management Plugin*. This APAR provides z/OSMF Configuration support for the new PARMLIB Management plugin.

- PH56708 (z/OSMF 3.1) - **New Function** - *z/OSMF Desktop Supports Data Set and File Search on a Remote System*. This APAR enhances the z/OSMF desktop user interface to:
  - Search data sets and files on a remote system with the Data Set and File Search tool and perform the same actions which can be performed on the local system.
  - View the jobs on a remote system with the Job Output task, and perform the same actions which can be performed on the local system.
  - Compare two sequential data sets, data set members or files from two different systems.
  - Support syntax highlight of Python programming language.
  - Add an option to enable Parmlib Management in the General Settings task.

  Note that, in relation to this APAR, a 'remote' system is any system that is not in the same sysplex as the z/OSMF you are logged on to. The remote system could be in another sysplex in the same site, or in a distant site; the only requirement is that there is a z/OSMF instance in the same sysplex as your remote system, and that there is network connectivity between the two sites.

  I would *strongly* encourage you to view the z/OSMF Guild video referenced in APAR PH56207 above. There are some REALLY nice capabilities in here - for example, being able to easily compares files or PDS members in two different sysplexes or two different sites, being able to easily copy files or members between sysplexes or sites, and much more. We have been saying for a long time that z/OSMF is increasingly becoming the place where system programmers will spend much of their day, with functions being available in z/OSMF that are simply not available in ISPF - the functions delivered by this APAR reinforce that statement.

- PH56716 (z/OSMF 3.1) - **New Function** - *z/OSMF Feedback Collection Enhancement*. This APAR:
  - Enhanced z/OSMF to gather certain product usage information, including system operational data, and send it back to IBM on a monthly basis[6]. The monthly feedback data sent to IBM is stored in the /global/zosmf/data/feedback directory in JSON format.and can be used for reporting purposes by the customer. Those JSON files can be utilized to generate reports using other third-party tools.
  - Added an option to enable or disable the monthly feedback collection in the General Settings task. It is enabled by default.

Note: You can control whether the operational data is sent to IBM. However, if you do *not* want the data to be sent, you must explicitly disable it.

  - Added an option to collect system operational data in the Provide IBM Feedback task.

[6] At this time, the "system operational data" is information about IBM products from the z/OS Function Registry.

- PH57412 (z/OSMF 3.1) - **New Function** - *Network Configuration Assistant Support for zERT Support for OpenSSH Upgrade*. This APAR adds Network Configuration Assistant Support for new SSH key exchange algorithms in the zERT technology in line with z/OS OpenSSH upgrades. Note that APAR PH58110 (see below) must also be applied for these new algorithms to be recognized by 3.1 z/OS Communications Server

- PH58105 / PH58110 (z/OSMF 3.1) - **New Function** - *zERT Network Analyzer Support for zERT Support for OpenSSH Upgrade*. This APAR updates the z/OSMF zERT Network Analyzer to import, query, and report from SMF 119, subtype 12 records containing new key exchange methods and key types supported in z/OS 3.1 OpenSSH version 8.4p1.

- PH59317 (z/OSMF 3.1) - **New Function** - *z/OSMF RM Plugin Grafana Support Updates*. This APAR corrects error handling for Grafana servers and delivers general maintenance. You can find information about z/OSMF support for Grafana here: *IBM RMF for z/OS Grafana Plugin*.

- PH59439 (z/OSMF 3.1) - **New Function** - *z/OSMF Semeru 17 Support*. This APAR changes the default Java level for z/OSMF 3.1 from Java 11 to Java 17 ("IBM Semeru Runtime Certified Edition for z/OS Version 11" to "IBM Semeru Runtime Certified Edition for z/OS Version 17"). The APAR text states "Although z/OSMF can continue to run on IBM Semeru Runtime Certified Edition for z/OS Version 11 for existing users, IBM recommends you to upgrade to IBM Semeru Runtime Certified Edition for z/OS Version 17 as early as possible, because IBM Semeru Runtime Certified Edition for z/OS Version 11 will be planned to be end of service by November 30, 2024".

  What about those customers that will still be on z/OS 2.5 after that date? z/OSMF in z/OS 2.5 uses Java 8 (Java 11 is not supported for use with z/OSMF in pre-z/OS 3.1 releases). The announced end of service date for both Java 8 and z/OS 2.5 is September 30, 2026, meaning that z/OSMF customers on z/OS 2.5 *will* be supported until z/OS 2.5 reaches it end of service date.

Serviceability APARs

We've seen a number of customer situations over the last year where IBM was unable to positively identify the cause of a problem, either because the dump that was captured was only a Partial Dump, or because the dump didn't contain all the required diagnostic information and the customer was unable to recreate the problem.

Providing best practices information for capturing diagnostic information would take up an entire article. However, while researching APARs for this Tuning Letter issue we noticed a fair few that are described as "Serviceability" improvements, so we though it would be valuable to include the list here.

- PH58992 (Db2 12-13) - *Db2 12 and Db2 13 for z/OS Serviceability Fix*.

- PH59020 (Db2 13) - **OPEN** - *Db2 New Function to Enhance Db2 TRACE Facility*.

- PH59477 (CICS 6.1) - *Improve CIST and CISD Diagnostics*.
- PH59853 (Db2 13) - **OPEN** - *COPY and COPYTOCOPY Utilities Serviceability*.
- PH60073 (Db2 (12-13) - *Diagnostic APAR That Provides -DISPLAY BUFFERPOOL SERVICE(8) Command*.

If you have any of these products, it might be prudent to apply the fixes for these APARs sooner rather than later.

Sub-Capacity Reporting Tool (SCRT) APARs

On December 13, 2023 IBM shipped APAR OA65794 which introduces a new version of SCRT, Version 30. This is an important event because IBM stipulates that you must use the latest SCRT *version* to create sub-capacity reports to be submitted to IBM. IBM will only accept reports for periods after January 2024 if they are created using V30.

You can always find information about the latest SCRT release, and the enhancements/fixes they deliver, in this document: https://www.ibm.com/downloads/cas/P9RDWLAX.

If you don't want to use the SMP/E-managed version, you can download a copy from http://ibm.biz/SCRTDownload.

- OA65794 (SCRT V30) - **New Function** - *SCRT 30.1.0*.

  This APAR delivers the following changes:

  – Support for Container Interval Rate Consumption reporting in new report section I7.

    Section I7 is optional and is enabled via a new SPECIAL control statement.

  – Support for zIIP usage reporting in new report section I8.

  – Support for variable capacity reporting in existing section V9.

  – The sample JCL has removed "-Xgcpolicy:optthruput" from IBM_JAVA_OPTIONS

  Support has been added for:

  – IBM Open Enterprise SDK for Node.js 20 – 5655-NOS

- OA66078 (SCRT V30) - *SCRT V30.1.0 Fails With JAVA Exception Unable to Convert CPUT - SCRT 30.1.1*. In large environments that result in "broken" SMF type 70.1 records, SCRT may miscalculate MSU consumption and/or MSU rate or may terminate with an error. This APAR addresses that error.
- OA66109 (SCRT V30) - *SCRTTOOL603 When SCRT is Run for VSE System and OS Product is Active Outside Reporting Period But Not Within the Period Itself*. This APAR only applies to SCRT users running multiple versions of the z/VSE operating system. See the APAR text for exact details of the configuration that must exist in order to cause this problem.

News from IBM Techdocs

IBM's Technical Sales Library, also known as simply 'Techdocs', can be found at http://www.ibm.com/support/techdocs. On that site you can search across all document types, or you can select to search just one type. However, we have found that over time, there has been a blurring of what is a Technote and what is a White Paper. It might still make sense within the framework of IBM's internal structure, but for an external consumer of that material, it is largely irrelevant.

If you are a regular visitor to the Techdocs site, you will be aware that many documents *appear* to be new or recently modified. But are they really? Many of them are simply 'updated' to get around IBM's rules about archiving 'obsolete' information. Despite multiple pleas from us (and customers) over the years, most documents do not have a change history. Many, if not most, don't even use change bars, so it is next to impossible to know if a document actually contains any new information or not.

When you combine that with the plummeting numbers of genuinely new documents, it simply is no longer sensible for us to continue to provide a list of all the Techdocs that have been 'updated' over the previous quarter. We feel that we can provide much more value to our readers if we invest the time we spend on researching the Techdocs in other areas of the Tuning Letter.

Nevertheless, we will keep this section and use it to inform our readers about any valuable, genuinely new Techdocs that we come across.

However, there *is* one set of documents that always do an outstanding job of clearly documenting all changes - we're referring to the CPS family of tools - zPCR, zBNA, CP3KEXTR, and so on. We will continue to include a link to those every time there is a change.

Here is this quarter's crop of updates to the CPS tools:

Getting Started with zPCR (IBM's Processor Capacity Reference)

This page provides access to zPCR Version 9.6.5 and to the related educational materials.

This update to zPCR 9.6.5 delivers the following changes:

- Packaging and Installation: IBM zPCR 9.6.5 is delivered as a single install package for Windows, identical for IBM Clients, IBM Business Partners, and IBM Employees. This change simplifies the process of building and testing install packages.

  IBM Java: A newer version of the IBM Java is included with the 9.6.5 install package. Therefore, you will be required to uninstall the previous IBM zPCR version before installing IBM zPCR 9.6.5.

- LSPR Table view: On the LSPR Table Control window, the Selected Families group box is now limited to only include recent processor families. To view older families, select All Families under Processors Displayed.

  The Selected Family check box default settings may be assigned on the Preferences window.

- LPAR Configuration Capacity Planning function:
  - EDF and RMF input: When multiple files are read, if the CEC IDs are different, the resulting dialog will now identify the file considered as the Master.
  - Measured SMT Benefit: Using EDF or RMF input, when a zIIP or IFL partition's utilization is less than 20% or when a partition's measured SMT benefit is greater than 90%, the SMT benefit is considered as being unreliable. While transferring the LPAR configuration to zPCR, a dialog will identify such cases and the default Estimated SMT value will be assigned instead. While viewing the SMT Benefit window, the ability to assign the Measured SMT benefit is provided.

- Corrections and Changes:

  Control Panel window: Corrected problem when the LPAR host was modified, the 2nd part of the LPAR configuration name was not being updated accordingly.
  - LPAR Host and Partition Configuration window: Corrected abend problem when copying partitions from a previous zPCR multi-configuration study file.
  - Partition Detail Report window: Corrected Capacity columns lineup on CSV output.
  - zAAP/zIIP Loading window: Corrected minor inconsistency of HTML output with actual window.

- Current EDF support:
  - z/OS - CP3KEXTR version 4.34 available on 11/28/2023.
  - z/VM - CP3KVMXT version 2.9g, available on 09/26/2023.

- Documentation:

  The IBM zPCR User's Guide and Online Help have been updated.

IBM Z Batch Network Analyzer (zBNA) Tool

This page provides access to zBNA and related documentation. The latest version of zBNA is V2.5.5 (Updated on March 19, 2024). The following enhancements have been added to zBNA since the last Tuning Letter issue:

- V2.5.5:
  - Starting with this release, for the Windows platform, there is a single zBNA installation file, zBNAInstallwithJava.exe, which is available to Clients, Business Partners, and IBM employees via the regular distribution locations.

- Additionally, the IBM Java Runtime version has been updated in this release, which requires an uninstall of the previous version. With this said, the V2.5.5 installation will not complete until the previous version is uninstalled
- Added “Proc Step” column to DFSORT IBM Z Sort table.
- Additional Classic Batch charts are supported in the Graph Aggregator application.
  - BATCH1008 – Total Job Initiator Time in MIPS
  - BATCH1017 – Total IIPCP Time in MIPS
- Miscellaneous other enhancements and bug fixes

zBNA has grown tremendously in scope and capabilities over the years - it is quite likely that it contains functions that you’ve never used. To get the most value out of it, have a look at the zBNA Enablement website - a single source for educational videos about zBNA.

Other Interesting Techdocs Documents

- *IBM z16 Server Time Protocol (STP): PTP/NTP Time Direct to CEC Accuracy Performance.* IBM white paper about Timer enhancements on z16. If you are responsible for STP in your organization, get yourself a large pot of very strong coffee and prepare to learn more than you wanted to know about FINRA regulations, MiFID II, STP, NTP, PPS and ... that was where I dozed off. But, seriously, if you are Ms. or Mr. STP in your organization, this is a very important paper that you really should read and understand.

- *z/OS Management Facility 3.1 Resource Requirements*. This excellent document, created by **Tian Na** in the z/OSMF development group, introduces the resource requirements for z/OSMF 3.1 and shows z/OSMF resources usage during startup and in idle time. It also provides some guidance on tuning z/OSMF startup. This is a must-read by anyone responsible for z/OSMF.

- *z/OS 3.1 Communications Server New Function APAR Summary*. This is a useful, and short, guide to significant new functions delivered by Continuous Delivery APARs. Please make sure your network colleagues are aware of this guide.

- *IBM Semeru Runtime Certified Edition for z/OS_17.0 - IBM Lifecycle Page*. There has been some confusion over the marketing and support dates for Java 17 for z/OS, so this page should be in your bookmark list - it should reflect IBM’s latest announcements on key dates for this product.

- *IBM Z Performance Report: Optimizing zIIPs with Banking Services from SAP on z16*. This very interesting paper describes a number of measurements comparing different mixes of general purpose CPs and zIIPs, and the impact of each configuration on CPU consumption, throughput, and response times. Definitely worth a read if you have workloads that can consume a lot of zIIP capacity. Unfortunately the test configuration did

not use SMT2, and there are no comparisons of enabling SMT2 vs not using it - an unfortunate gap in an otherwise excellent paper.

- *Switches and Directors qualified for IBM zSystems FICON and FCP channels*. This web page that lists the qualified Brocade and Cisco DWDMs has been updated to add new firmware levels for both vendors.

- *IBM Best Practice: Collect Correlator SMF Records for Performance Problem Diagnosis*. This white paper describes the reasons *why* IBM recommends enabling and collecting the z Workload Interaction Correlator SMF records, provides information about the products that exploit the Correlator, and explains how to enable the Correlator SMF records.

- *z/OS TCP/IP Routing*. This excellent white paper explains TCP/IP routing for those with no networking experience - even *I* could understand it. With so many performance enhancements in the z/OS Communication Server, the information in this white paper will help z/OS performance specialists identify the options that are appropriate for their environments.

Conferences

The 2024 conference season is underway now. In this section we let you know about the conferences that have been announced as of the end of April. We will keep this list (which is in chronological order) updated as more conferences are announced during the year.

GSE Nordic Region Conference

The 2024 GSE Nordic Region conference will be held on June 11 to 13 in Copenhagen, Denmark. The conference is dedicated to mainframe topics, offering six tracks, and presentations by a top class mix of customers and vendors. You can find more information about the conference, including the agenda, a link to the registration form, link to the download area, on this website.

SHARE

The SHARE August 2024 conference (4+ days, this time) will be held in the Loews Kansas City Hotel in Kansas City, Missouri, from Sunday, August 4 through Thursday, August 8. Registration is open now, with the $300 ‘early bird’ discount available if you register by July 12, 2024.

The sessions have not been selected yet, but will be available on the conference website when the agenda has been finalized. We will provide more information in *Tuning Letter 2024 No. 2*.

IBM TechXchange 2024

The IBM TechXchange Conference 2024 will be held at the Mandalay Bay Convention Center in Las Vegas, from Tuesday, October 22, to Thursday, October 24, 2024. The announcement doesn't say what the price is, but it does say "The Tier 1 price reflects "early bird" pricing when purchasing 1-4 passes through May 31, 2024. After which time, pricing will increase on June 1, 2024 by ~7% and on August 1, 2024 by ~18%. The purchase price of 5 or more passes in a single transaction will not increase (Tier 2/Tier 3)."

The announcement also doesn't specify which IBM platforms will be included, but given that it mentions "1000 sessions", we expect that all IBM platforms will be participating. There is also no information yet about the agenda. Hopefully that will be available before the deadline for the "early bird" pricing.

GSE UK

The annual in-person GSE UK conference will take place from Monday, November 4 to Thursday, November 7 in Whittlebury Hall, in Towcester, Northamptonshire. Last year's outstanding in-person conference (in the same venue) was sold out, so if you want to be sure of a place at the 2024 conference (and sure of a room in Whittlebury Hall), we recommend signing up for the GSE UK newsletter (so you won't miss the formal announcement) and then registering for the conference as soon as registration is opened up.

With the 2024 Virtual conference just ended, I expect the GSE UK team will soon start looking at the agenda for the in-person conference. You can get the latest status on the conference at https://www.gse.org.uk/events/gse-uk-in-person-conference-2024/.

Also, while on the topic of GSE UK, we want to make sure you are aware that the GSE UK organization make the recordings of their past events available to *everyone* from pre-schoolers to pensioners for *free* on their YouTube channel. They show a commitment to the mainframe platform that is sadly absent from many other companies in our industry, so the least we can do is support them by attending their excellent events.

Acknowledgments

It would not be possible to create a document like this without the kind help of many people. So we offer our sincere thanks to the following people for their help, support, and enthusiasm:

- Scott Barry, SBBTech LLC, USA
- Mario Bezzi, Watson & Walker, Italy
- Scott Compton, IBM, USA
- Todd Havekost, IntelliMagic, an IBM company, USA
- Joe Hyde, IntelliMagic, an IBM company, USA
- Ed Jaffe, Phoenix Software International, USA
- Stephen Kocik, IBM, USA
- Shawn Lundvall, IBM, USA
- Peter Muench, IBM, Germany
- Jim Mulder, IBM, USA
- Peter Relson, IBM, USA
- Dale Riedy, IBM, USA
- Oliver Robinson, Altair, UK
- Andrew Rowley, Black Hill Software
- Rick Schaeffer, IBM, USA
- Brad Snyder, IBM, USA
- Bart Steegmans, IBM, Belgium
- Marna Walle, IBM, USA
- Tom Wasik, IBM, USA
- Mark Wilson, Vertali, UK
- Xiao Zhen (Joey) Zhu, IBM, China

Cheryl's List Blog Posts

We hope that all our readers are familiar with Cheryl's List. This blog is available to anyone at no charge. We use it to inform Tuning Letter subscribers when a new issue is available on our subscriber-only publications site (www.watsonwalkerpublications.com). We also use it to get information out to readers (and anyone that subscribes to the blog) about important news that can't wait until the next Tuning Letter. As you can see in Figure 19, the blog is available under the Blog/Media pulldown on our public website.

If you do not currently subscribe to the blog, we strongly encourage you to do so. You can subscribe to the blog by clicking here. Apart from using the blog to pass along urgent news to our readers, we also use the blog to inform subscribers when we have published a new edition of the Tuning Letter.

Figure 19 - Cheryl's List Blog

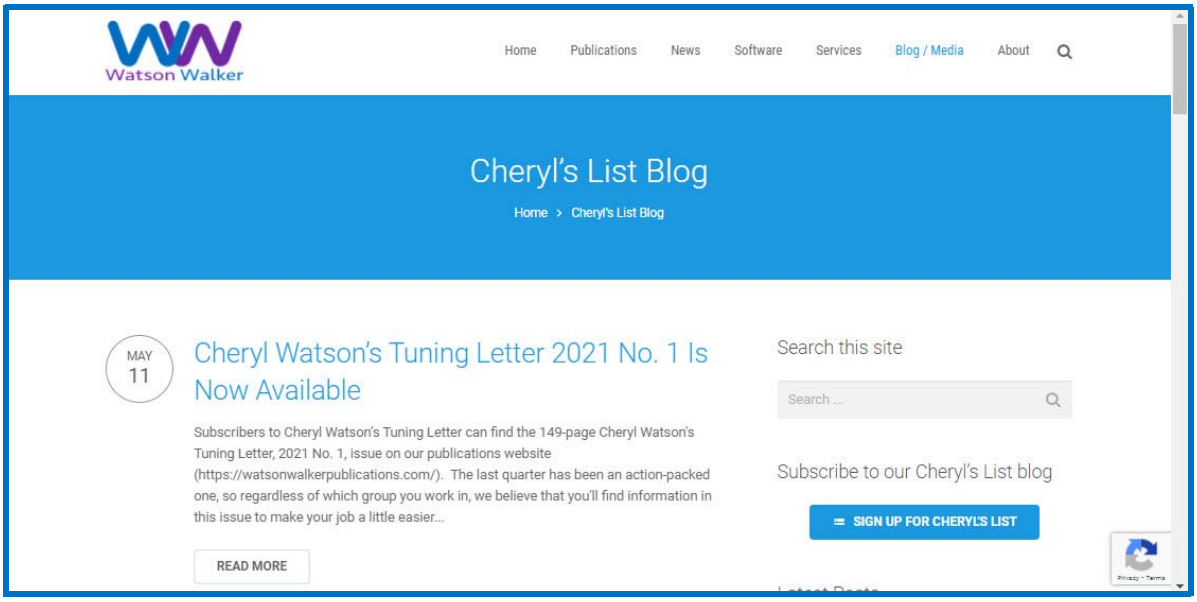

If you have any information or experience that you think would be of interest to your peers, please let us know and we will be happy to pass it on. We are all members of the mainframe family, and we all benefit when we all contribute.

Blog Entries

For the last while, we had moved to just providing a link to the recent blog posts. However, with our new publications site, this meant that the blog posts were no longer included in the index on that site. Therefore, we are reverting back to including the full text from each blog

post in the Tuning Letter. The following are our new posts for January 2024 to the end of March 2024:

Feb122024 - Cheryl Watson's Tuning Letter 2023 No. 4 Available Now

Subscribers to Cheryl Watson's Tuning Letter can find the latest 114-page issue, *2023 No.4*, on our publications website (https://watsonwalkerpublications.com/). In this issue, you will find the following articles:

- *Cheryl and Tom's blog post* announcing their plans to retire and close the Watson & Walker company and cease publication of Cheryl Watson's Tuning Letter. This is truly the end of an era. Many of our friends and readers have pointed out that Cheryl and Watson & Walker have been there for their entire working lives. But no one can deny that after a lifetime devoted to helping customers and the mainframe platform, Tom and Cheryl have earned the chance to hang up the keyboard and enjoy life to the full. We will have more news about their plans and what this means for Tuning Letter subscribers in future issues. For now, everything will continue Business-as-Usual until early 2025, with the normal four Tuning Letter issues (after this one) being published for 2024.

- While mainframe workloads continue to grow to support their businesses, many customers find that what they need is capacity rather than a screaming-fast processor. A growing number of customers are now including sub-capacity processors in their upgrade evaluations. Moving to a CPC in a slower speed range presents both opportunities and challenges. To help our readers understand the considerations, our friend **Todd Havekost** from IntelliMagic addresses this in his latest article – *The Expanding Role of Sub-capacity Processors in Today's Mainframe Configurations*. Whenever you are next considering a CPC upgrade, we highly recommend using Todd's article to help you decide if a sub-cap CPC should be among your list of candidates.

- Huge application databases tend to receive all the attention and love and care in a production system. This is as it should be – performance issues on those files can impact your entire business. However, there are likely thousands of partitioned data sets that are the workhorses of your z/OS system – they hold your software, your application executables, your JCL and utilities, the source for our application programs, and countless other things. But in our experience, they get little TLC. Even performance-enhancing capabilities that have been with us for 20+ years are often not exploited. To address this missed opportunity for system optimization, our *Optimizing PDS and PDSE Performance* article in this issue describes the basics and provides recommendations for providing a sound PDSE environment. In *Tuning Letter 2024 No. 1*, we will dive into the commands and metrics that you can use to identify the best candidates for optimization improvement.

- Possibly the best thing about Parallel Sysplex is its robustness and resiliency – the ability to 'take a licking and keep on ticking'. But that is also possibly the worst thing about Parallel Sysplex. While it will perform great with properly-sized CF structures, it sort of shrugs its shoulders and makes the most of under-sized structures too. As a result, it is not uncommon to see CF structures that are functioning flawlessly, but that could deliver better performance and efficiency if their size was better suited to their workload. The *Best Practices for Sizing CF Structures* article discusses the factors that can affect a structure's size, describes how structure sizes have traditionally been managed (or ignored), and the discusses the latest Structure Sizing enhancement to the z/OS CFRM Policy Editor. To cut a long story short, it can help you achieve better efficiency while also saving you time and effort – sounds like a win-win to me.

- Our friend **Mike FitzGerald** from Fitz Software in Ireland presented a very interesting session about the z/OS IPL process at last year's GSE UK Conference in Whittlebury. This is a topic that many experienced system programmers take for granted, but that could potentially be horrendously complex to someone new to the platform (or even an experienced person from a different site). With so many z/OS infrastructure staff retiring over the coming few years, we felt that this is a very interesting topic for the next generation that will replace them. So Mike kindly offered to create The *Life of a z/OS IPL* article for us based, on his GSE presentation. I've been working with z/OS for nearly 40 years, and *I* learned things from Mike's article. We think this article will be valuable for anyone that could someday be faced with a system that won't IPL.

- Of course, we also have our regular helping of *User Experiences and Tips*. In this issue we describe one customer's experience with adding and removing large amounts of memory to a z/OS LPAR. We have a valuable little tip about a relatively new function in the DEVSUPxx member to protect VTOC, as well as a clarification about the ability to use new parms in the IEAOPTxx member to control the behavior of the new Implicit Long Term CPU Protection function in WLM. There is also some additional information about the z/OS Connect metrics that Todd discussed in his *z/OS Connect EE: Strategic On-ramp to the Mainframe* article in *Tuning Letter 2023 No. 3*. And we have additional information about how sites might integrate the new z/OSMF UUID function into their own software deployment processes.

- And to round things off, the *News* article provides a meaty list of recent New Function APARs for z/OS, CICS, and Db2, plus info about updated CPS tools from IBM, and information about upcoming user conferences.

You can find the full Table of Contents here. We hope you find this information helpful and timely. Please let us know if there are specific topics you would like to see in a future Tuning Letter article.

If you are not a current Tuning Letter subscriber, see our website for information about subscription rates and the ordering process.

Mar012024 - W&W Guide to SHARE in Orlando 2024 (Part 1)

The 2024 Winter SHARE conference will be held in the Rosen Center hotel in Orlando, Florida, starting this coming Sunday, March 3. This one is a 5-day conference, with the last day (Thursday) being a full day. If you are attending, check out the agenda and download the SHARE app for your mobile device.

I don't have the Technical Agenda from a previous SHARE conference to hand, but it seems like there are more sessions in parallel than normal (for example, there are 13 sessions all at 9:15 on the Monday morning); maybe it is my imagination, or maybe it is because the conference ends on Thursday. But in any case, when building my plan for the week I'm finding that I have many instances of scheduling clashes between multiple sessions that I really want to attend. If you are attending with colleagues, I suggest that you sit down in advance and agree who will attend which session.

Our week will kick off with the keynote session at 7:30 on Monday morning. Those that know me will be aghast at the idea of me being up so early in the morning – obviously I am not in control of the session scheduling. Much more importantly, however, is that this will be Cheryl's last SHARE conference, so if you would like to give her your good wishes, maybe grab a selfie, or even snag an autograph, this will be an excellent opportunity to do so.

Given the short week, and my unofficial role as Cheryl's paparazzi blocker, we only have one session at this SHARE:

- Thursday, 13:15, Salon 13-14, my colleague **Mario Bezzi** and I will be co-presenting the *Watson & Walker zRoadshow – Spring 2024 Collection* session. This will be one of the last sessions of the week, and it is a descendant of Cheryl's famous 'Chery's Hot Flashes' sessions, so we are hoping that Cheryl will join us – this gives you another opportunity to wish her and her husband Tom well in their upcoming adventures.

There are many interesting sessions this week, some related to z/OS V3.1, but there are also many other topics – I don't think there is a single slot when there aren't multiple sessions I would love to attend.

Some of the sessions I would love to attend (but often can't because of scheduling clashes) are:

Monday

- Monday, 09:15, Salon 13-14, *z/OS System Programming & Hardware KickOff*. As appealing as the other sessions in this timeslot are, this truly is a cant'-miss session. To quote from the session abstract, "*the legendary **Cheryl Watson** will provide a window into her storied career in mainframes, and more importantly, hear her thoughts on what the future holds for the platform we all love. This promises to be a fascinating and engaging talk, and for many, a once-in-a-lifetime opportunity to meet a true giant in the*

Mainframe Community." Whether this is your first SHARE, or you've known Cheryl for the last 50 years, this will be a session you will remember for the rest of your career.

- Monday, 10:30, Salon 20, *MQ High Availability – What Are The Options*, by **Neil Johnston**.
- Monday, 10:30, Salon 13-14, *What's new in z/OS 3.1: Sunshine Edition*, by **Dave Surman**, **Marna Walle**, and **Steve Warren**.
- Monday, 13:15, Salon 13-14, *Upgrade to z/OS 3.1, Part 1: Planning*, by **Marna Walle** and **Aaron Kippins**.
- Monday, 14:45, Salon 20, *Key Considerations for Managing and Optimizing IMS and Other Subsystems Under Tailored Fit Pricing*, by **Tracy Dean**.
- Monday, 14:45, Salon 8, *Protect Your Business with IBM Z Cyber Vault*, by **John Thompson**.
- Monday, 14:45, Salon 13-14, *Upgrade to z/OS 3.1, Part 2: Technical Actions*, by **Marna Walle**.
- Monday, 14:45, Salon 21, *What's New in z/OSMF?*, by Joey Zhu.
- Monday, 14:45, Salon 12, *z/OS Performance Track Opening and IBM WSC Performance Hot Topics!*, by **Brad Snyder** and **Meral Temel**.
- Monday, 16:00, Salon 21, *JES2 z/OS 3.1 Product Update and Latest Status*, by **Tom Wasik**.
- Monday, 16:00, Salon 15, *Maximize Your z/OS Support Experience*, by **John Shebey** and **Laura Sperling**.
- Monday, 16:00, Salon 16, *Proactive Incident Management with AI/ML for Anomaly Analytics*, by **Tim Brooks** and **Yuk Chan**.
- Monday, 16:00, Salon 13-14, *z/OS 3.1 User Experiences*, by **Ed Jaffe**.
- Monday, 17:00, *SHARE Technical Exchange and Reception (aka Expo)*. Free 'refreshments' and a chance to chat with other mainframe geeks – how often do you find yourself with an opportunity like that?

Tuesday

- Tuesday, 08:00, Salon 20, *BYOD Lab: Choose your own z/OSMF Topic*, by **Alexander Giemsa**, **Marna Walle**, **Stephen Warren**, and **Joey Zhu**. A team like that to answer all my z/OSMF questions is enough to get even a lazy sod like *me* out of bed at an un-Godly hour for the *third day in a row*!

- Tuesday, 08:00, Salon 6, *Db2 for z/OS 101: Buffer Pools and Group Buffer Pools*, by **Anna McKee**.
- Tuesday, 08:00, Salon 12, *Workload Management (WLM) Update for IBM z16 and IBM z/OS 3.1*, by **Andreas Henicke**.
- Tuesday, 09:15, Salon 20, *Acquisition-to-Deploy, Software Management Best Practices using z/OSMF*, by **Mike DuBois** and **Dawn Damore**.
- Tuesday, 09:15, Salon 8, *Cyber Resilience for Enterprise Data – Protection from Logical Corruption*, by **Beth Peterson**.
- Tuesday, 09:15, Salon 21, *What every z/OS System Programmer needs to know about AI*, by **Meral Temel** and **Steve Warren**.
- Tuesday, 09:15, Salon 7, *When is an Anomaly a Problem? Identifying the Meaning Behind Using Machine Learning to Identify Operational Anomalies*, by **Tim Brooks** and **Dan Wiegand**.
- Tuesday, 10:30, Salon 21, *30th Anniversary of Parallel Sysplex: A Retrospective and Lessons Learned*, by **Peter Enrico**.
- Tuesday, 10:30, Salon 19, *Explore all the latest happenings within the Java ecosystem on z/OS*, by **James Tang**.
- Tuesday, 13:00, Salon 6, PSP: *State of the Art for RMF and SMF Performance and Capacity Analysis (IntelliMagic)*, by **Todd Havekost** and **Brent Philips**.
- Tuesday, 13:00, Salon 17-18, *PSP: Observability Revolutionized: A Blueprint for Operational Excellence (Broadcom Mainframe Software)*, by **Nicole Fagen**. I don't know any of the details on this one, but was told that "It is the biggest launch of my career for sure".
- Tuesday, 14:15, Salon 13-14, *Capacity on Demand Technical Deep Dive*, by **Mitchell Bride**.
- Tuesday, 14:15, Salon 17-18, *z/OS AI Framework Overview and Configuration*, by **Steven Partlow** and **Hiren Shah**.
- Tuesday, 15:30, Salon 15, *DFSMS: Adventures in Out of Space*, by **Matthew Barragan** and **Frank McCune**.
- Tuesday 16:30, *STE and Reception*. An opportunity for you to pick up your conversation from when the bar went dry on Monday.

And that's just for the first two days! We'll be back in a day or two with some suggestions for Wednesday and Thursday – and remember that Thursday is a nearly-full day, with sessions running until 15:30, and our zRoadshow session from 13:15 to 14:15.

Many of you will know that we (Watson & Walker) have a strong working relationship with IntelliMagic. Intellimagic's co-founder, **Dr. Gilbert Houtekamer**, recently announced his retirement. Like Cheryl, Gilbert will be at SHARE next week and is looking forward to meeting up with old friends and reminiscing about 'the good old times'. This is going to be a SHARE conference to remember, with two titans (and two good friends) of the mainframe performance world, both in attendance for the last(?) time. If you know Gilbert, make sure to say Hello and wish him a long and happy retirement.

If you are going to be at SHARE, please say Hi if you see us. And don't be shy about wishing Cheryl well with her next big adventures. We hope you have safe travels and an enjoyable conference.

Frank

Mar052024 - W&W Guide to SHARE in Orlando 2024 (Part 2)

Hi All,

It is hard to believe that the 2024 Winter SHARE conference is half-over already. And what a conference it has been – our very own Cheryl Watson was presented with the SHARE Distinguished Service Award at the opening session on Monday morning. This honor recognizes Cheryl's 47 years' of contributions and dedication to the SHARE organization, its members, and the mainframe platform. In addition, SHARE presented my colleague Mario Bezzi and me with the John Ehrman Award for our work on delivering technical education and guidance to SHARE members. I have to say, when I was an eager young System Programmer 40 years ago, I would never have believed that I would someday be working with superstars like Cheryl and Mario, and receiving an award like this – I truly have had a charmed career.

What a fantastic start to a great week. I always say that I learn something new every day. Well, that's not true so far this week. At the moment, I'm running at about TWENTY new things each day, and I expect tomorrow and Thursday to be just as interesting and enlightening.

The conference agenda and abstracts for the sessions are available to anyone here. Conference attendees can access the agenda, the session handouts, and other material on the events.rdmobile.com website – the link to the site was included in an email titled "Download the SHARE Orlando Mobile App" that you should have received after you registered for the conference.

Given the short week, we are only presenting one session at this SHARE:

- Thursday, 13:15, Salon 13-14, my colleague **Mario Bezzi** and I will be co-presenting the Watson & Walker zRoadshow session. We will have our usual eclectic mix of new functions, customer experiences, interesting APARs, and so on – well, neither Mario nor I have any slides yet, but when we do, I expect they will cover all those topics. Much more important, given that the zRoadshow is the descendant of Cheryl's famous 'Chery's Hot Flashes' sessions, we are expecting Cheryl to join us.

In case you haven't heard, Cheryl is planning for this to be her last SHARE conference, so if you would like to give her your thanks for all her great work over the years, or your good wishes for an adventure-filled retirement, maybe grab a selfie, or even snag an autograph, this will be an excellent opportunity to do so.

Our *W&W Guide to SHARE in Orlando 2024 (Part 1)* post included a list of some of the sessions we were hoping to attend yesterday and today (Tuesday). The following are the sessions on Wednesday and Thursday that especially caught our eye:

Wednesday

- Wednesday, 08:00, Salon 13-14, *Introduction to z/OS Containers*, by **Kershaw Mehta**.
- Wednesday, 08:00, Salon 7, *PCI DSS Version 4.0 – How Will it Impact Your Mainframe Cybersecurity?*, by **Jessica Doherty** and **Steve Hosie**. If you think anything to do with PCI DSS is limited to your Security colleagues, you are in for a rude awakening. If you work in the Payments Card Industry, you should attend this session to find out how busy you are going to be over the next 12 months.
- Wednesday, 08:00, Salon 8, *Using Automated Expert Recommendations to Bridge the Mainframe Skill Gap*, by **Harish Vellore Yuvaraj**.
- Wednesday Monday, 08:00, Salon 21, *z16 Concepts and Capacity User Experiences*, by our friends **John Baker** and **Todd Havekost**. This is always a fun-filled and very interesting session from two award-winning presenters, although I must speak to John about the time – I've a sneaking suspicion that he scheduled this at 08:00 so that I'll be too drowsy to ask any hard questions.
- Wednesday, 09:15, Salon 13-14, *End-to-End Demo of z/OS Containers (Parts 1 & 2)*, by **Kershaw Mehta**.
- Wednesday, 09:15, Salon 21, *Pitfalls on Non-traditional Migrations*, by **David Hutton**. If a sub-capacity CPC is among your CPC upgrade options, you have to attend this session by one of IBM's leading and most approachable Z performance experts.
- Wednesday, 10:30, Salon 6, *Run-it-Back Db2 for z/OS SWAT Team All New 2023 SWAT Tales*, by **Anthony Ciabattoni**.

- Wednesday, 10:30, Salon 21, *Sweat The Small stuff Before It Becomes The Big Stuff*, by **Jim Horne** and **Jerry Street**.
- Wednesday, 13:45, Salon 16, *BYOD Lab : Building a Dashboard to Monitor Your Workloads with RMF and Grafana*, by **Anastasiia Didkovska** and **Paul Stewart**.
- Wednesday, 13:45, Salon 17-18, *Exploring Future Directions for AI on IBM Z*, by **Andrew Sica** and **Elpida Tzortzatos**.
- Wednesday, 15:00, Salon 12, *Cyber Intrusion Detection for z Systems Implementation Strategy for Mainframe customers*, by **Justin Bastin**.
- Wednesday, 15:00, Salon 21, *RMF: The Latest and Greatest*, by **Anastasiia Didkovska** and **Paul Stewart**.
- Wednesday, 16:15, Salon 15, *Integrating Cloud Objects in your Batch Processing*, by **Andrew Witt**.
- Wednesday, 16:15, Salon 12, *Look at the Changes in z/OS 3.1 Change Tracker!*, by **Scott Marcella** and **Chris Taylor**.
- Wednesday, 16:15, Salon 13-14, *MVS logger: push your logs up the stream*, by **Mike Shorkend**.
- Wednesday, 16:15, Salon 21, *WSC Performance Experiences with z/OS V3.1*, by **Dave Betten** and **Brad Snyder**.
- Wednesday, 17:30, Salon 12, *Gain Insights from The z/OSMF Guild Community: Advantages, Obstacles, and Importance of z/OSMF Adoption*, by **Fiona King**, **Rolando Perez**, and **Hiren Shah**.

Thursday

- Thursday, 08:00, Salon 12, *Introduction to IBM SMF Explorer*, by **Alexander Giemsa**.
- Thursday, 08:00, Salon 13-14, *ISPF Hints, Tips, and Recent Updates*, by **Sam Reynolds**.
- Thursday, 08:00, Salon 8, *WatchTower's Auto-Discovered Mainframe Topology using SMF with no Magical Genie or Professor Von Drake's Required*, by **Krishna Kasibhatla** and **Michal Kramaric**.
- Thursday, 09:15, Salon 7, *Cyber Resiliency and Regulatory Compliance*, by **Butch Rambish**, **Chris Taylor**, and **Chris Walker**.
- Thursday, 09:15, Salon 19, *Testing from Zero to Hero in 60 mins*, by **William Yates**.
- Thursday, 09:15, Salon 12, *WSC System Programmer Hot Topics*, by **Meral Temel**.

- Thursday, 10:30, Salon 13-14, *Managing resources in JES2*, by **Tom Wasik**.
- Thursday, 10:30, Salon 12, *Processor Sizing with zPCR*, by **Shawn Lundvall** and **Brad Snyder**.
- Thursday, 10:30, Salon 16, *z/OS Parallel Sysplex Update*, by **Mark Brooks**.
- Thursday, 13:15, Salon 21, *BYOD Lab: WSC CPS Tools Lab: Pick 'em – zPCR and zBNA Hands-on Lab*, by **Shawn Lundvall** and **Brad Snyder**.
- Thursday, 13:15, Salon 16, *User Experience – Moving to well defined well managed Sysplex*, by **Jerry Edgington**.
- Thursday, 13:15, Salon 13-14, *Watson & Walker zRoadshow – Spring 2024 Collection*, by **Mario Bezzi** and **Frank Kyne**, with a special guest appearance by the inimitable **Cheryl Watson**.
- Thursday, 14:30, Salon 13-14, *Bit Bucket x'43'*, by **Lionel Dyck**, **Ed Jaffe**, **Henri Kuiper**, and **Mike Shorkend**.
- Thursday, 14:30, Salon 21, *Configuring LPARs to Optimize Performance*, by **Scott Chapman**.

I think that should be more than enough to keep you entertained for the next two days and on your flight home. If you are at SHARE, please say Hi if you see us. And make sure to wish Cheryl well and pass along any paraglider tips you might have. We hope you have a great week and a safe trip home afterwards.

Frank

If you are not a current Tuning Letter subscriber, see our website for information about subscription rates and the ordering process.

That's it for this quarter. Thank you to all our subscribers for all your support, your helpful feedback, and for sharing your very interesting experiences with us.

Stay Tuned!

The Team at Watson & Walker